ESA STUDY CONTRACT REPORT

No ESA Study Contract Report will be accepted unless this sheet is inserted at the beginning of each volume of the Report.

ESA CONTRACT No **AO/1-4469/03/NL/SFe**	SUBJECT **BIONICS & SPACE SYSTEMS DESIGN TECHNICAL NOTE 1**		CONTRACTOR **UNIVERSITY OF SURREY**
* ESA CR( )No	* STAR CODE	No of volumes 1 This is Volume No 1	CONTRACTOR'S REFERENCE

ABSTRACT:
This document is Technical Note 1 for the Bionics & Space Systems Design study jointly authored by the University of Surrey (prime), University of Sussex, University of Bath, EADS Astrium and the British Antarctic Survey It provides a comprehensive survey of the field of biomimetics whereby engineered technology is biologically-inspired. We have emphasised its application to robotics as robotics has particular significance to space technology. This is especially the case for planetary exploration.

The work described in this report was done under ESA contract. Responsibility for the contents resides in the author or organisation that prepared it.

Names of authors: Dr Alex Ellery

** NAME OF ESA STUDY MANAGER Dr Mark Ayre DIV: Advanced Concepts Team DIRECTORATE:	** ESA BUDGET HEADING

BIONICS & SPACE SYSTEMS DESIGN AO/1-4469/03/NL/SFe

Technical Note 1

Version: 1.0

Date: 02 November 2004

Prepared by:

Dr. Alex Ellery
Surrey Space Centre
School of Electronics and Physical Sciences
University of Surrey
Guildford
Surrey GU2 7XH
UK
Tel: +44 1483 683882
Fax: +44 1483 689503
Email: a.ellery@surrey.ac.uk

This page is left intentionally blank

List of Contributing Authors

Dr. Alex Ellery
Surrey Space Centre
School of Electronics and Physical Sciences
University of Surrey
Guildford
Surrey GU2 7XH
UK
Tel: +44 1483 683882
Fax: +44 1483 689503
Email: a.ellery@surrey.ac.uk

Professor Julian Vincent
Centre for Biomimetic and Natural Technologies
Department of Mechanical Engineering
The University
Bath BA2 7AY
UK
Tel: +44 1225 386596
Fax: +44 1225 386928
Email: j.f.v.vincent@bath.ac.uk

Professor Phil Husbands
Department of Infomatics
University of Sussex
Brighton BN1 9QH
UK
Tel: +44 1273 678556
Fax: +44 1273 671320
Email: P.Husbands@sussex.ac.uk

Dr Charles Cockell
British Antarctic Survey
High Cross
Madingley Road
Cambridge CB3 OET
UK
Tel: +44 1223 411722
Fax: +44 1223 411722
Email: csco@bas.ac.uk

Mr. Steven Eckersley
Mission Systems Department
EADS Astrium Ltd.
Gunnels Wood Road
Stevenage
Hertfordshire SG1 2AS

UK
Tel: +44 1438 773301
Email: Steven.ECKERSLEY@astrium.eads.net

Dr Yang Gao
Surrey Space Centre
School of Electronics and Physical Sciences
University of Surrey
Guildford
Surrey GU2 7XH
UK
Tel: +44 1483 689278
Fax: +44 1483 689503
Email: y.gao@surrey.ac.uk

Mr Greg Scott
Surrey Space Centre
School of Electronics and Physical Sciences
University of Surrey
Guildford
Surrey GU2 7XH
UK
Tel: +44 1483 683418
Fax: +44 1483 689503
Email: g.scott@surrey.ac.uk

Mr Matthew Quinn
Centre for Computational Neuroscience & Robotics
John Maynard Smith Building
University of Sussex
Brighton BN1 9QH
UK
Tel: ++44 1273 872948
Fax: +44 1273 671320
Email: matthewq@sussex.ac.uk

Mr Jon Bird
Centre for Computational Neuroscience & Robotics
John Maynard Smith Building
University of Sussex
Brighton BN1 9QH
UK
Tel: ++44 1273 872948
Fax: +44 1273 671320
Email: jonba@sussex.ac.uk

Mr. David Wyatt
Mission Systems Department
EADS Astrium Ltd.
Gunnels Wood Road

Stevenage
Hertfordshire SG1 2AS
UK
Tel: +44 7789 682692
Email: david@wyatts.me.uk

List of Applicable Documents

[AD1]	Invitation to Tender	IMT-CTM//SFe/2003.960
[AD2]	Bionics & Space System Design: Statement of Work (Appendix 1)	GSP-03/L27
[AD3]	Draft contract (Appendix 2)	AO/1-4469/03/NL/Sfe
[AD4]	Special Conditions of Tender (Appendix 3)	AO/1-4469/03/NL/Sfe
[AD5]	General Conditions of Tender for ESA Contracts	ESA/C/290 rev 5
[AD6]	Biomimicry – A Review	ESA pdf publication (Mark Ayre)

List of Reference Documents

Abbattista, F., Di Gioia, G., Di Santo, G. and Fanelli, A. M. (1996) "An associative memory based on the immune networks" *Proceedings of the 1996 IEEE International Conference on Neural Networks* **1**, 519-523, Washington, DC, 3-6 June 1996. IEEE, New York, NY, USA

Abdelhammed M, Pinspon U & Cetinkunt S (2002) "Adaptive learning algorithm for cerebellar model articulation controller" *Mechatronics* **12**, 859-873

Abrahamsson, M.L.A., Baudin, H.B., Tran A., Philouze C., Berg, K.E., Raymond-Johansson, M.K., Sun, L., Akermark, B., Styring, S. & Hammarstrom, L. (2002) "Ruthenium-manganese complexes for artificial photosynthesis: factors controlling intramolecular electron transfer and excited state quenching reactions" *Inorg. Chem.* **41**, 1534-1544.

Abu-Mostafa Y & St Jacques J (1985) "Information capacity of the Hopfield model" *IEEE Trans Inform Theory* **31**(4), 461-464

Abu-Mostafa Y (1989) "Information theory, complexity and neural networks" *IEEE Comm Mag* (Nov), 25-28

Ackley D & Littman M (1991) "Interaction between learning and evolution" *ALife II*, Santa Fe Institute Studies **10**, 487-507

Adams P & Cox K (2002) "Synaptic Darwinism and neocortical function" *Neurocomput* **42**, 197-214

Adams, G. G., Truong, T. Q., Barnes, T. G., Lu, X., Aceros, J. C. (2002) "Biomimetic Flow and Contact / Bending MEMs Sensors" In *Neurotechnology for Biomimetic Robots* (ed. Ayers J, Davis J, Rudolph A), Bradford Books, MIT Press, Cambridge, MA, USA, 13-30

Agrawal A (1999) "Raw computation" *Sci Am* (Aug)

Ahmad Z & Mark J (1998) "Biomimetic materials: recent developments in organic-inorganic hybrids" *Mater Sci & Eng* **6C**, 183-196

Aiello L et al (1986) "Representation and the use of metaknowledge" *Proc IEEE* **74**(10), 1304-1321

Ainsworth W (1967) "On the efficiency of learning machines" *IEEE Trans Systems Science & Cybernetics* **3**(2), 111-116

Aisu, H. & Mizutani H. (1996), "Immunity-based learning – integration of distributed search and constraint relaxation", *Proceedings of the International Workshop on Immunity-Based Systems 1996*, Kyoto, Japan

Akbarzadeh-T M, Kumbla K, Tunstel E and Jamshidi M, "Soft computing for autonomous robotic systems," Computers and Electrical Engineering, Vol. 26, pp. 5-32, 2000

Alaimonas J, Weiss I & Bandyopadhyay A (1988) "Active vision" *Int J Comp Vision*, 333-356

Albus J (1971) "A theory of cerebellar function" *Math Biosci* **10**, 25-61

Albus J (1975) "A new approach to manipulator control: the cerebellar model articulation controller (CMAC)" *Trans ASME J Dynamics Systems Measurement & Control* **97** (Sep), 220-227

Albus J (1979a) "Model of the brain for robot control I" *Byte* (Jun), 10-34

Albus J (1979b) "Model of the brain for robot control II" *Byte* Jul), 54-95

Albus J (1991) "Outline for a theory of intelligence" *IEEE Trans Syst, Man & Cyber* 21(3), 473-509

Albus J (1999) “Engineering of mind” *Info Sci* 17, 1-18

Albus J Lumia R & McCain H (1988) “Hierarchical control of intelligent machines applied to space station telerobotics” *IEEE Trans Aero & Elect Syst* 24 (5), 535-541

Albus J, McCain H & Lumia R (1987) “NASA/NBS Standard Reference Model for Telerobotic Control Systems Architecture” *NASA TN 1235*

Aleksander I & Morton H (1986) “Artificial intelligence: an engineering approach” *IEE Proc* 134D (4), 218-223

Aleksander I (1973) “Brains, minds and machines” *Electron & Power* (Nov), 499-502

Allen J & Littman D (1986) “Plans, goals and language” *Proc IEEE* 74 (7), 939-948

Almqvist N, Thomson N et al (1999) "Methods for fabricating andcharacterizing a new generation of biomimetic materials 1." *Materials Science and Engineering C*7, 37–43

Aloimonos J, Weiss I & Bandyopadhyay A (1988) “Active vision” *Int J Computer Vision* **1** (4), 333-356

Amari S (1971) “Characteristics of randomly connected threshold element networks” *Proc IEEE* 59 (1), 35-47

Amari S (1994) “Mathematical foundations of neurocomputing” *Proc IEEE* 78 (9), 1443-1462

Andersen S (2001) "Matrix proteins from insect pliable cuticles: are they flexible and easily deformed?" *Insect Biochemistry and Molecular Biology*31, 445-452

Andersen S, Peter M et al (1996) "Cuticular sclerotization in insects." *Comparative Biochemistry and Physiology B-Biochemistry &Molecular Biology*113, 689-705

Anderson J (1983) "Retrieval of information from long term memory" *Sci* 220, 25-30

Anderson J (1989) "Theory of the origins of human knowledge" *Artif Intell* 40, 313-351

Anderson J (2002) “Spanning seven orders of magnitude: a challenge for cognitive modelling” *Cog Sci* 26, 85-112

Anderson T & Donath M (1990) “Animal behaviour as a paradigm for developing robot autonomy” *Robot & Auton Syst* 6, 145-168

Anderson, A. W., Nordan, H. C., Cain, R. F., Parrish, G. & Duggan, D. E. (1956) "Studies on a radiation-resistant micrococcus. I. Isolation, morphology, cultural characteristics and resistance of γ-radiation" *Food Technology* **10**, 575-577

Anderson, J.M. & Anderson, B. (1988) “The dynamic photosynthetic membrane and regulation of solar energy conversion” *Trends in Biochemical Sciences* **13**, 351-355

Anderson, J.M. (1986) “Photoregulation of the composition, function, and structure of the thylakoid membranes” *Annual Review of Plant Physiology* **37**, 93-136.

Angell J et al (1983) “Silicon micromechanical devices” *Scientific American* **248** (4), 26-47

Angre P & Chapman D (1990) “What are plans for?” *Robotics & Autonomous Syst* 6, 17-34

Antsaklis P, Passino K & Wang J (1991) “Introduction to autonomous control systems” *EEE Control Syst Mag* (Jun), 5-13

Arbib M & Caplan D (1979) “Neurolinguistics must be computational” *Behaviour & Brain Sci* 2, 449-483

Arbib M (1967) “Automata theory and development” *J Theoretical Biology* 14, 131-156

Arcmiegas J, Eltimseley A, Cios K (1997) “Neural network based adaptive control of flexible robot arms” *Neurocomputing* **17**, 141-157

Ardell D & Andersen S (2001) "Tentative identification of a resilin gene in*Drosophila melanogaster*." *Insect Biochemistry and Molecular Biology*31, 965–970

Arimoto S et al (1984) "Bettering operation of robots by learning" *J Rob Syst* 1(2), 123-140

Arkin R (1987) “Motor schema based mobile robot navigation” *International Journal of Robotics Research* **8** (4), 92-112

Arkin R (1990) “Impact of cybernetics on the design of a mobile robot system: a case study” *IEEE Trans Syst Man & Cyber* 20 (6), 1245-1257

Arkin R (1990) “Integrating behavioural, perceptual and world knowledge in reactive navigation” *Robotics & Autonomous Systems* **6**, 105-122

Arkin R, Ali K, Wetzenfeld A & Cervantes-Perez F (2000) “Behavioural models of the praying mantis as a basis for robotic behaviours” *Robot & Auton Syst* **32**, 39-60

Arzt, E., Gorb, S., & Spolenak, R. (2003) "From micro to nano contacts in biological attachment devices" *Proceedings of the National Academy of Sciences of the USA* **100**(19), 10603-10606.

Asada M, MacDorman K, Ishiguro H & Kuniyoshi Y (2001) “Cognitive developmental robotics as a new paradigm for the design of humanoid robots” *Robot & Auton Syst* **37**, 185-193

Ashley S (2004) “Artificial muscles” *Scientific American* (Oct), 35-41

Ashmore J & Gale J (2000) “The cochlea” *Curr Biol* **10** (9), R325-R327

Atick J & Redlich A (1990) “Towards a theory of early visual processing” *Neural Comp* **2**, 308-320

Atkeson V, Hale J, Kawato M, Kotosaka S, Pollick F, Riley M, Schaal S, Shibata T, Tevatia G, Ude A & Vijayakumar S (2000) “Using humanoid robots to study human behaviour” IEEE Intelligent Systems 15 (4), 46-56

Atkins E (1985) "Conformations in polysaccharides and complex carbohydrates" *Proceeding of the International Symposium of Biomolecular Structural Interactions*, Supplement to the *Journal of Bioscience*.

Autumn, K., Liang, Y. A., Hsieh, S. T., Zesch, W., Chan, W. P., Kenny, T. W., Fearing, R. S. & Full, R. J. (2000) "Adhesive force of a single gecko foot-hair" *Nature* **405**, 681-685

Autumn, K., Sitti, M., Liang, Y. A., Peattie, A. M., Hansen, W. R., Sponberg, S., Kenny, T. W., Fearing, R. S., Israelachvili, J. N. & Full, R. J. (2002) "Evidence for van der Waals adhesion in gecko setae" *Proceedings of the National Academy of Sciences of the USA* **99**(19), 12252-12256

Axelsson, H., Muhammed, A., and Egersted, M. (2003). Autonomous formation switching for multiple, mobile robots. In *IFAC Conf. on Analysis and Design of Hybrid Systems*, Sant Malo, Brittany, France.

Ayara, M., Timmis, J., de Lemos, R., de Castro, L. N. and Duncan, R. (2002) “Negative selection: how to generate detectors” *Proceedings of ICARIS 2002*, 89-98, Editors: J. Timmis and P.J. Bentley

Ayers J (1995) “A reactive ambulatory robot architecture for operation in current and surge” *Proc Auton Vehicles in Mine Countermeasures Symp* 15-31

Ayers J, Davis J & Rudolf A (2002) *Neurotechnology for Biomimetic Robots.* Bradford books, MIT Press, USA

Ayers J, Zavracky P, McGruer N, Massa D, Vorus V, Mukherjee R, Currie S (1998) “A modular behavioral-based architecture for biomimetic autonomous underwater robots” in Proc Asutonomous Vehicles in Mine Countermeasures Symp, Naval Postgraduate School, USA

Baatz M & Wagner G (1997) "Adaptive inertia caused by hidden pleiotropic effects" Theor Popul Biol 51 (1), 49-66

Babloyantz A & Destexhe A (1988) "Is the normal heart a periodic oscillator?" *Biological Cybernetics* 58, 203-211

Babloyantz A & Lourenco C (1996) "Control of spatiotemporal chaos in neuronal networks" *International Journal of Neural Systems* 7(4), 507-517

Babloyantz A, Lourenco C, Sepulchre J (1995) "Control of chaos in delay differential equations in a network of oscillators and in a model cortex" *Physica D* 86, 274-283

Bachelder I & Waxman A (1995) "View-based neurocomputational system for relational map-making and navigation in visual environments" *Robot & Auton Syst* **16**, 267-289

Baerends G (1976) "Functional organisation of behaviour" *Animal Beh* 24, 726-738

Bailey S et al (2000) "Biomimetic robotic mechanisms via shape deposition manufacturing" *Proc SPIE Conf Sensor Fusion & Decentralised Control in Robotic Systems III* (ed. McKee G & Schenker P) 4196

Bailey S, Cham J, Cutkosky M & Full R (2000) "Biomimetic robotic mechanisms in shape deposition manufacturing" *Proc SPIE Conf Sensor Fusion & Decentralised Control in Robotic Systems III* (ed. McKee G & Schenker P) 4196

Bailieul J (1980) "Chaotic motion in nonlinear feedback systems" *IEEE Trans Circuits & Syst* 27 (11), 996-997

Bajcsy R (1988) "Active perception" *Proc IEEE* **76** (8), 996-1005

Balch, T. (1999). The impact of diversity on performance in multi-robot foraging. In *Proc. Agents99*. Seattle, WA. USA.

Balch, T. and Arkin, R. (1994). Communication in reactive multiagent robotic systems. *Autonomous Robots*, (1):1–25.

Balch, T. and Arkin, R. (1998). Behavior-based formation control for multiagent robot teams. *IEEE Transactions on Robotics and Automation*, 14(6):926–939.

Balch, T. and Parker, L. E. (2000). Guest editorial. *Autonomous Robotics*, 8(3), 207–208.

Baldassarre G., Nolfi S., Parisi D. (2003a). Evolution of collective behaviour in a team of physically linked robots. In Raidl G., Guillot A., Meyer J.-A. (eds.), *Applications of Evolutionary Computing - Proceedings of the Second European Workshop on Evolutionary Robotics (Evo*ROB-2003, held at *EvoWorkshops2003), pp. 581-592*. Berlin: Springer Verlag.

Baldassarre G., Nolfi S., Parisi D. (2003a). Evolution of collective behaviour in a team of physically linked robots. In Raidl G., Guillot A., Meyer J.-A. (eds.), *Applications of Evolutionary Computing - Proceedings of the Second European Workshop on Evolutionary Robotics (Evo*ROB-2003, held at *EvoWorkshops2003), pp. 581-592*. Berlin: Springer Verlag.

Baldassarre, G., Nolfi, S., and Parisi, D. (2002). Evolving mobile robots able to display collective behaviour. In Hemelrijk, C. K., editor, *Proceedings of an International Workshop on Self-Organisation and the Evolution of Social Behviour*, pages 11–22. Zurich: Swiss Federal Institute of Technology.

Baldassarre, G., Nolfi, S., and Parisi, D. (2003b). Evolving mobile robots able to display collective behaviour. *Artificial Life*, 9(3), 255–267.

Baldassarre, G., Nolfi, S., and Parisi, D. (2004). Coordination and behaviour integration in cooperating simulated robots. *Proceedings of the 8th International Conference on Simulation of Adaptive Behavior*.

Balkenius C (1994) “Biological learning and artificial intelligence” *Lund University Cognitive Stud (LUCS)* **30**, ISSN 1101-8453

Balkenius C (1994) “Natural intelligence for autonomous agents” *Lund University Cognitive Studies* **29** (1101), 8453

Ball P (2001) "Life's lessons in design." *Nature*409, 413-416

Ballard D (1991) “Animate vision” *Artif Intell* 48, 57-86

Ballard D et al (1983) “Parallel visual computation” *Nature* 306 (6), 21-26

Baltes H & Brand O (2001) “CMOS-based microsenors and packaging” *Sensors & Actuators* **A92**, 1-9

Baltes H (1996) “Future of IC microtransducers” *Sensors & Actuators* **A56**, 179-192

Bandura A (1974) “Behaviour theory and the models of Man” *Amer Psychol* (Dec), 859-869

Baran V (1989) "Associative memory & fractal neural networks" *IEEE Syst, Man & Cyber* 19 (5), 1133-1141

Bar-Cohen Y (2002) “Electroactive polymers as artificial muscles: a review” *J Spacecraft & Rockets* **39** (6), 822-827

Bar-Cohen Y et al (1998) “Flexible, low mass robotic arm actuated by electro-active polymers” *Robotics* 98, 3rd Conf Robotics for Challenging Environments, Albuquerque, New Mexico

Bar-Cohen Y et al (2000) “Ultrasonic/sonic drilling/coring (USDC) for in-situ planetary applications” Proc SPIE Smart Structures Conf, Newport Beach, CA, paper no 3992-101

Bar-Cohen Y, Leary S, Shahinpoor M, Harrison J & Smith J (1999) “Electro-active polymer (EAP) actuators for planetary applications” *Proc SPIE Annual International Symposium on Smart Structures & Materials*, March 1999, Newport Beach, CA, USA, **3669**, 57-63

Bar-Cohen Y, Leary S, Shahinpoor M, Harrison J & Smith J (1999) “Flexible, low-mass devices and mechanisms actuated by electroactive polymers” *Proc SPIE Annual International Symposium on Smart Structures & Materials*, March 1999, Newport Beach, CA, USA, paper no. 3669-38

Bar-Cohen Y, Xue T, Shahinpoor M, Simpson J & Smith J (1998) “Flexible, low mass robotic arm actuated by electro-active polymers and operated equivalently to the human arm and hand” *Robotics* **98**, 3rd ASCE Conf Robotics for Challenging Environments, April 1998, New Mexico, USA

Bar-Cohen Y, Xue T, Shahinpoor M, Simpson J & Smith J (1998) “Low-mass muscle actuators using electroactive polymers (EAP)” *Proc SPIE Conf on Smart Materials Technologies*, Mar 1998, san Diego, CA, USA **3324**, 218-223

Bar-Cohen Y. et al. (2001) “Ultrasonic/sonic drilling/coring (USDC) for planetary applications” *Proc SPIE 8th Ann Symp on Smart Structures & Materials,* 4327-55, 2001

Barlow H (1970) “Definition of intelligence” *Nature* 228, 1008

Barlow H (1983) “Intelligence, guesswork and language” *Nature* 304, 207-209

Barlow J (1964) “Inertial navigation as a basis for animal navigation” *J Theor Biol* **6**, 76-117

Barnes D (1998) “Hexapod robot locomotion over uneven terrain” Proc IEEE Conf on Control Applications, Trieste, Italy, 441-445

Barth F (1973) "Microfiber reinforcement of an arthropod cuticle: laminated composite material in biology." *Zeitschrift fur Zellforschnung*144, 409-433

Barth F (1978) "Slit sense organs: 'strain gauges' in the arachnid exoskeleton" *Symposium of the Zoological Society of London* **42,** 439-448

Barto A & Ananden P (1985) "Pattern recognising stochastic learning automata" *IEEE Trans Systems Man & Cybernetics* 15 (3), 360-375

Barto A, Sutton R & Anderson C (1983) "Neuron-like adaptive elements that can solve difficult learning control problems" *IEEE Trans Syst Man & Cyber* 13 (5), 834-846

Bastin T (1970) "General property of hierarachies" in *Towards a Theoretical Biology* 3 (ed . Waddington C), 252-257

Batista A & Newsome W (2000) "Visuo-motor control: giving the brain a hand" *Curr Biol* **10** (4), R145-R148

Baughman R (1996) "Conducting polymer artificial muscles" *Synthetic Metals* **78**, 339-353

Baz A, Imam K & McCoy J (1990) "Active vibration control of flexible beams using shape memory actuators" *J Sound & Vibration* **140** (3), 437-456

Bechert D, Bruse M et al (1997). "Biological surfaces and their technological application - laboratory and flight experiments on drag reduction and separation control" *AIAA Fluid Dynamics Conference.*

Beckers, R., Holland, O., and Deneubourg, J. (1994). From local actions to global: Stigmergy and collective robotics. In Brooks, R. and Maes, P., editors, *Artificial Live IV: Proceedings of the Fourth International Workshop on the Synthesis and Simulation of Living Systems*. MIT Press, Cambridge, MA.

Bedford F (1999) "Keeping perception accurate" *Trends in Cog Sci* **3** (10), 4-11

Beer R & Chiel H (1997) "Brain has a body: adaptive behaviour emerges from interactions of nervous system, body and environment" *Trends in Neurosci* **20** (12), 553-557

Beer R (1995) "A dynamical systems perspcive on agent-environment interaction" *Artif Intell* **72**, 173-215

Beer R et al (1990) "Biological perspective on autonomous agent design" *Robot & Auton Syst* 6, 169-186

Beer R, Chiel H & Sterling L (1990) "Biological perspectives on autonomous agent design" *Robot & Auton Syst* **6**, 169-186

Beer R, Chiel H, Quinn R, Espanscied K, Lurrson P (1992) "Distributed neural network architecture for hexapod robot locomotion" *Neural Comp* **4**, 356-365

Beer, R. D., Chiel, H. J., Quinn, R. D., Ritzmann, R. E. (1998) "Biorobotic Approaches to the Study of Motor Systems" *Current Opinion in Neurobiology* **8**(6), 777-782.

Beer, R. D., Quinn, R. D., Chiel, H. J., Ritzmann, R. E. (1997) "Biologically Inspired Approaches to Robotics: What can we learn from insects?" *Communications of the ACM* **40**(3), 31-38.

Bejan A (1992) "Surfaces covered with hair: optimal strand diameter and optimal porosity for minimum heat transfer." *Biomimetics* 1, 25-40

Bekey G (1996) "Biologically inspired control of autonomous robots" *Robot & Auton Syst* **18**, 21-31

Belew R (1991) "Artificial life: a constructive lower bound on artificial intelligence" *IEEE Expert* (Feb), 8-15

Benbruhum H & Franklin J (1997) "Biped dynamic walking using reinforcement learning" *Robot & Auton Syst* **22**, 283-302

Bennett A (1996) "Do animals have cognitive maps?" *J Exper Biol* **199**, 219-224

Bennetto, H. P., Stirling, J. L., Tanaka, K. & Vega C. A. (1983) "Anodic reactions in microbial fuel cells" *Biotechnology and Bioengineering* **25**, 559-568

Bergveld P (1996) “The future of biosensors” *Sensors & Actuators A* 56, 65-73

Berlucchi G & Aglioti S (1997) “Body in the brain: neural bases of corporeal awareness” *Trends in Neurosci* **20**, 560-564

Bernardino A Santos-Victor J (1998) “Visual behaviours for binocular tracking” *Robot & Auton Syst* **25**, 137-146

Berns K, Dillman R & Piekenbrock S (1995) “Neural networks for the control of a six-legged walking machine” *Robot & Auton Syst* **14**, 233-244

Bianconi P (1995) "Biomimetic mineralization."*Advance in Chemistry Series*245,509-522

Biederman I (1987) “Recognition by components: a theory of human image understanding” *Psychol Rev* 94 (2), 18-64

Bigi A, Boanini E, Panzavolta S, Roveri N (2000) "Biomimetic growth of hydroxyapatite on gelatin films doped with sodium polyacrylate." *Biomacromolecules* 1 (4), 752-756

Billard A & Dautenhahn K (1998) “Grounding communication in autonomous robots: an experimental study” *Robot & Auton Syst* **24**, 71-79

Billard A & Dautenhahn K (1999) “Experiments in social robotics: grounding and use of communication in autonomous agents” *Adaptive Behav* 7 (3/4), 415-438

Billard A & Mataric M (2001) “A biologically inspired robotic model for learning by imitation” Proc 4th Int Conf on Autonomous Agents, Barcelona, Spain, 373-380

Birk, A., Coradeshi, S., and Tadokoro, S., editors (2002). *Robocup-2001: The Fifth RoboCup Competitions and Conferences*. Springer Verlag, Berlin.

Bledsoe W (1975) "Non resolution theorem proving" *Proc 4th Int J Conf Artif Intell,* 91-108

Blum H (1962) “Origin and evolution of human culture” *Amer Sci* 50, 32-47

Blum H (1967) "A machine independent theory of the complexity of recursive functions" *J Assoc Comp Mach* 14 (2), 322-336

Blynel, J. (2003) "Evolving Reinforcement Learning-Like Abilities for Robots" A. Tyrrell, P.C. Haddow, and J. Torresen: *Evolvable Systems: From Biology to Hardware: 5th International Conference*, ICES 2003.

Blynel, J. and Floreano, D. (2002) "Levels of Dynamics and Adaptive Behavior in Evolutionary Neural Controllers" B. Hallam, D. Floreano, J. Hallam, G. Hayes, and J.-A. Meyer, editors. *From Animals to Animats 7: Proceedings of the Seventh International Conference on Simulation on Adaptive Behavior*, MIT Press.

Blynel, J. and Floreano, D. (2003) "Exploring the T-Maze: Evolving Learning-Like Robot Behaviors using CTRNNs" *Proceedings of EvoROB2003: 2nd European Workshop on Evolutionary Robotics.*

Bock P (1985) “Emergence of AI: learning to learn” *AI Mag* (Fall), 180-190

Bonabeau E, Theraulaz G, Deneubourg J-L, Aron S & Camazine S (1997) “Self-organisation in social insects” *TREE* **12** (5), 188-193

Bond G, Richman R et al (1995) "Mimicry of natural material designs and processes." *Journal of Materials Engineering and Performance* 4, 334-345

Bond, D. R. & Lovley, D. R. (2003) "Electricity production by *Geobacter sulfurreducens* attached to electrodes" *Applied and environmental microbiology* **69**(3), 1548–1555

Bond, D. R., Holmes, D. E., Tender, L. M. & Lovley, D. R. (2002) "Electrode-reducing microorganisms that harvest energy from marine sediments" *Science* **295**, 483-485

Bonser R & Dawson C (1999) "The structural mechanical properties of down feathers and biomimicking natural insulation materials." *Journal of Materials Science Letters*18, 1769-1770

Booker L et al (1989) “Classifier systems and genetic algorithms” *Artif Intell* 40, 235-282

Bos E & Jeannerod M (2002) “Sense of body and sense of action both contribute to self-recognition” *Cog* **85**, 177-187

Bossley K, Mills D, Brown M and Harris C (1995) ``Construction and Design of parsimonious neurofuzzy systems." in Neural Network Engineering in Dynamic Control Systems, K. J. Hunt, G. W. Irwin and K. Warwick (Ed.), Berlin: Spring-Verlag, pp. 153-178

Bouslapda F., and Ichikawa A (1993) “Application of neural networks to fuzzy control”, *Neural Networks* **6,** 791-799

Bower J (1995) “Reverse engineering the nervous system: an in vivo, in vitro and in computo approach to understanding the mammalian olfactory system” in *An Introduction to Neural & Electronic Networks*, Academic Press Inc, 3-29

Brachman R (1983) "What an is-a link is and isn't: analysis of taxonomic links in semantic networks" *IEEE Comp* (Oct), 30-36

Bradley D & Tyrell A (2000) “Hardware fault tolerance: an immunological solution” in Proc IEEE Conf on Systems, Man & Cybernetics, Nashville, USA, 107-112

Bradley D & Tyrell A (2000) “Immunotronics: hardware fault tolerances inspired by the immune system” *Proc 3rd Int Conf on Evolvable Syst*, 11-20

Bradley D, Ortega-Sanchez C & Tyrell A (2000) “Embryonics + Immunotronics: a bio-inspired approach to fault tolerance” Proc 2nd NASA/DOD Workshop on Evolvable Hardware, 215

Bradley, D. W. and Tyrrell, A. M. (2000) “Immunotronics: hardware fault tolerance inspired by the immune system” Proceedings of the 3rd International Conference on Evolvable Systems 1801, 11-20, Lecture Notes in Computer Science, Springer-Verlag

Bradt J-H, Mertig M et al (1999) "Biomimetic mineralization of collagen by combined fibril assembly and calcium phosphate formation." *Chemistry and Materials* 11, 2694-2701

Braham R & Hornblen J (1990) “Design of a neural network with a biologically motivated architecture” *IEEE Trans Neural Networks* 1 (3), 251-262

Braitenburg V (1984) “Vehicles: experiments in synthetic psychology” *MIT Press*, Cambridge, MA, USA

Brecht A & Gauglitz G (1995) “Optical probes and transducers” *Biosensors & Bioelectron* 10, 923-936

Bremer F (1993) "Biomimetic materials: nature's material concepts" *Ingenieur Werkstoffe*5, 28-29

Breulmann M, Colfen H et al (1998) "Elastic magnets: template controlled mineralization of iron oxide colloids in a sponge-like gel matrix." *Advanced Materials*10, 237-241

Bridgeman B, van der Heijden A & Velichkovsky B (1994) “A theory of visual stability across saccadic eye movements” *Behav & Brain Sci* **17** (2), 247-292

Brooks R (1989) "A robot that walks: emergent behaviours from a carefully evolved network" *Neural Comp* 1, 253-262

Brooks R & Stein L (1994) "Building brains for bodies" *Auton Robots* 1, 7-25

Brooks R (1986) "Robust layered control system for a mobile robot" *IEEE J Robot & Autom* 2 (1), 14-23

Brooks R (1990) "Elephants don't play chess" *Robot & Auton Syst* 6, 3-15

Brooks R (1991) "Intelligence without representation" *Artif Intell* 47, 139-159

Brooks R (1991) "New approaches to robotics" *Science* **253**, 1227-1232

Brooks R (1996) "Prospects for human level intelligence for humanoid robots" Proc 1st Int Symp on Humanoid Robots, Tokyo, Japan

Brooks R, Breazeal C, Irie R, Kemp C, Merjunovic M, Scarrellati B, Williamson M (1998) "Alternative essences of intelligence" *AIAA*

Brooks, R. (2001) "The relationship between matter and life" *Nature* **409** (18 January)

Brooks, R. A. (1992). Artificial Life and real robots. In Varela, F. J. and Bourgine, P., editors, *Proc. 1st Euro. Conf. on Artificial Life*, pages 3–10, Paris, France. MIT Press.

Brown, R. G. and Jennings, J. S. (1995). A pusher/steerer model for strongly cooperative mobile robot cooperation. In *Proc. IEEE/RSJ Int. Conf. on Intelligent Robots and Systems*, pages 562–586.

Bruce B (2000) "Chloroplast transit peptides: structure, function and evolution" *Trends in Cell Biol* 10 (Oct), 440-447

Bruner J (1957) "On perceptual readiness" *Psychol Rev* 64, 123-146

Bucher, D., Akay, T., DiCaprio, R. A., Buschges, A. (2003) "Interjoint Coordination in the Stick Insect Leg-Control System: The Role of Positional Signaling" *Journal of Neurophysiology* **89**, 1245-1255.

Buckley J and Hayashi Y, ``Fuzzy Neural Networks A Survey," Fuzzy Sets and Systems, Vol. 66, pp. 1-13, 1994.

Bugman G (1997) "Biologically plausible neural computations" *Biosyst* **40**, 11-19

Bundy A (1999) "Survey of automated deduction" Informatics Research Report, University of Edinburgh, UK

Bundy A, Silver B & Plummer D (1985) "Analytical comparison of some rule-learning programs" *Artif Intell* **27** (2), 137-181

Burdinski, D., Wieghardt, K., Steenken, S. (1999) Intramolecular electron transfer from Mn or ligand phenolate to photochemically generated $\mathrm{Ru^{III}}$ in multinuclear Ru/Mn complexes. Laser flash photolysis and EPR studies on photosystem II models. *J. Am Chem Soc.* **121**, 10781-10787.

Burgos J & Moreno-TovarP (1996) "Zipf scaling behaviour in the immune system" *Biosyst* **39**, 227-232

Caianello E (1961) "Outline of a theory of thought processes and thinking machines" *J Theor Biol* 2, 204-235

Caldwell D, Tsagarakis N & Medrano-Cerda G (2000) "Biomimetic actuators: polymeric pseudo muscular actuators and pneumatic muscle actuators for biological emulation" *Mechatronics* **10**, 499-530

Calvert P, Cesarano J et al (2002) "Toughness in synthetic and biological multilayered systems." *Philosophical Transactions of the Royal Society A*360, 199-209

Campbell D (1956) “Perception as substitute trial and error” *Psychol Rev* **63** (5), 330-341

Campbell, N. and Reece, J. (1997) *Biology* (6th edition), Addison-Wesley, Boston, Massachusetts, USA

Canham, R. O. and Tyrrell, A. M. (2002) “A multilayered immune system for hardware fault tolerance within an embryonic array” *Proceedings of ICARIS 2002*, 3-11, Editors: J. Timmis and P.J. Bentley

Cao, Y. U., Fukunaga, A. S., and Kahng, A. B. (1997). Cooperative mobile robotics: antecendents and directions. *Autonomous Robots*, 4:1–23.

Capaday C (2002) “Special nature of human walking and its neural control” *Trends in Neurosci* 25 (7), 370-376

Capi, G., Nasu, Y., Barolli, L., Mitobe, K. (2003) “Real time gait generation for autonomous humanoid robots: A case study for walking.” *Robotics and Autonomous Systems*. **42**, 107-116.

Carbonell J (1989) "Introduction: paradigm for machine learning" *Artif Intell* 40, 1-9

Carbonell J et al (1982) “Machine learning: a historical and methodological analysis” *AI Mag* (Fall), 69-78

Carlson D & Ram S (1990) "Hyperintelligence: next frontier" *Comm Assoc Comp Mach* 33 (3), 311-321

Carnap R (1951) "Foundations of logic & mathematics" *Int Encyc Unified Sci* 1, 143-171

Carpenter G & Grossberg S (1988) “ART of adaptive pattern recognition by self-organising neural networks” *IEEE Comp* (March), 77-88

Carpenter R (2000) “Neural control of looking” *Curr Biol* **10**(8), R291-R293

Caulfield J et al (1989) “Optical neural networks” *Proc IEEE* 77 (10), 1573-1587

Caveney S (1969) "Muscle attachment related to cuticle architecture in apterygota." *Journal of Cell Science*4, 541-559

Cayzer, S. and Aickelin, U. (2002) “A recommender system based on the immune network” *HP Technical Report HPL-2002-1*

Chaimowicz, L., Sugar, T., Kumar, V., and Campos, M. F. M. (2001). An architecture for tightly coupled multi-robot cooperation. In *Proc. IEEE Int. Conf. Robotics and Automation*, pages 2292–2297, Seoul, South Korea. IEEE Press.

Chambers L, Practical handbook of genetic algorithms: new frontiers. Boca Raton GL: CRC Press, 1995.

Chandrasekaram B & Shen D (1968) “On expediency and convergence in variable stochastic automata” *IEEE Trans Systems Science & Cybernetics* 4 (1), 52-60

Chandrasekaram B & Shen D (1969) “Stochastic automata games” *IEEE Trans Systems Science & Cybernetics* 5 (1), 145-149

Chandrasekaran B (1981) "Natural & social metaphors for distributed problem solving" *IEEE Syst Man & Cyber* 11 (1), 1-5

Chandresekaran B & Goel A (1988) "From numbers to symbols to knowledge structures: AI perspectives on classification tasks" *IEEE Trans Syst Man & Cyber* 18(3), 415-424

Chang S-M, Muramatsu H, Nakaruma C & Miyake J (2000) “Principle and applications of piezoelectric crystal sensors” *Mater Sci & Engin* **C12**, 111-123

Chang-an WC, Huang Y et al (2000) "Biomimetic structure design — a possible approach to change the brittleness of ceramics in nature." *Materials Science and Engineering C*11, 9–12

Chao C, Chen Y and Teng C, Simplification of fuzzy-neural systems using similarity analysis, IEEE Trans. System, Man and Cybernetic, 26 (1996) 344-354.

Chapman D (1987) "Planning for conjunctive goals" *Artif Intell* **32**, 333-337

Charles J, Chihara C et al (1997) "A cluster of cuticle protein genes of *Drosophila melanogaster*at 65a: sequence, structure and evolution." *Genetics*147, 1213-1224

Charniak C (1988) "Undebuggability and cognitive science" *Comm Assoc Comput Machin* 31 (4), 402-412

Chaudhuri, S. K. & Lovley, D. R. (2003) "Electricity generation by direct oxidation of glucose in mediatorless microbial fuel cells" *Nature Biotechnology* **21**, 1229-1232

Chella A, Frixione M & Gaglio S (1997) "Cognitive architecture for artificial vision" *Artif Intell* **89**, 73-111

Chen B, PengX et al (2002) "Research on the microstructure of insect cuticle and the strength of a biomimetic preformed hole composite." *Micron*33

Chen F (1990) "Backpropagation neural networks for nonlinear self-tuning adaptive control" *IEEE Control Syst Mag* (Apr), 44-48

Chiao, M., Lam, K. B. & Lin, L. (2003) "Micromachined microbial fuel cells" *Proceedings of IEEE Micro Electro Mechanical Systems Conference*, 383-386, January 2003, Kyoto, Japan

Chiel H & Beer R (1997) "The brain has a body: adaptive behaviour emerges from interactions of nervous system, body and environment" *Trends in Neurosci* **20**, 553-557

Choi S, Han J, Kim J & Cheong C (2001) "Force tracking control of a flexible gripper featuring shape memory alloy actuators" *Mechatronics* **11**, 677-690

Chomsky N (1958) "Three models for the description of language" *Proc IRE Inf Th* 2 (3), 134-124

Chomsky N (1962) "Logical basis of linguistic theory" *9th Int Cong Linguistics*, Cambridge, Mass

Choset , H. (1996) "Sesor Based Motion Plannin: The Hierarchical Generalized Voronoi Graph" PhD hesis, Carnigie Mellon University

Chu R et al (1990) "Neural networks for system identification" *IEEE Control Syst Mag* (Apr), 31-34

Chua L & Yang L (1988) "Cellular neural networks: applications" *IEEE Trans Circuits & Systems* 35 (10), 1273-1290

Chua L & Yang L (1988) "Cellular neural networks: theory" *IEEE Trans Circuits & Systems* 35 (10), 1257-1272

Clancey W (1993) "Knowledge level reinterpreted: modelling socio-technical systems" in *Knowledge Acquisition in Modelling* (ed. Ford M & Bradshaw J), John Wiley & Sons, New York, 33-50

Clancey W (1993) "Situated action: a neurophysiological interpretation" *Cog Sci* 17 (1), 87-116

Clarac F, Cattaert D & Le Ray D (2000) "Central control components of a simple stretch reflex" *Trends in Neurosci* **23** (5), 199-208

Clark A & Grush R (1999) "Towards a cognitive robotics" Adaptive Behaviour 7 (1), 5-16

Clark J, Cham J et al (2001). "Biomimetic design and fabrication of a hexapedal running robot" *IEEE International Conference on Robotics and Automation.*

Clark R (1994) "Biosystems analysis: the relationship between direct velocity feedback control and the monosynaptic pathway of the central nervous system" *J Intelligent Material Systems & Structures* **5** (Sept), 723-728

Cliff D (1994) “Computational neuroethology” *Univ Sussex Tech Rep CSRP338*

Cliff D, Harvey I & Husbands P (1993) “Explorations in evolutionary robotics” *Adaptive Behaviour* 2 (1), 71-104

Cliff D, Harvey I, Husbands P (1992) “Incremental evolution of neural network architectures for adaptive behaviour” *University of Sussex Cognitive Sci Research Paper CSRP 256*

Coblentz F, Shafer T et al (1998) "Cuticular proteins from the blue crab alter in vitro calcium carbonate mineralization." *Comparative Biochemistry and Physiology B-Biochemistry and Molecular Biology*121, 349-360

Coburn M (1952) "Brain analogy: a discussion" *Psychol Rev,* 453-460

Cogdell R & Lindsay J (1998) “Can photosynthesis provide a biological blueprint for the design of novel solar cells?” *Trends in Biotechnol* **16**, 521-527

Coghlan A (1992) “Smart ways to treat materials” *New Scientist* (Jul), 27-29

Cohen F (1987) "Computer viruses - theory & experiment" *Comp & Sec* 6, 22-35

Cohen P & Levesque H (1990) "Intention is choice with commitment" *Artif Intell* 42, 213-261

Cohen P & Perrault C (1979) “Elements of a plan-based theory of speech acts” *Cog Sci* 3, 177-212

Cohen Y & Knudsen E (1999) “Maps versus clusters: differenr representations of auditory space in the midbrain and forebrain” *Trends in Neurosci* **22** (3), 128-135

Collett T (2000) “Animal navigation: birds as geometers?” *Curr Biol* **10** (19), R718-R721

Collins A & Loftus E (1974) "Spreading activation theory of semantic networks" *Psych Rev* 82, 407-428

Coltheart M (1999) “Modularity and cognition” *Trends in Cog Sci* 3 (3), 115-120

Comilleau-Perez V & Gideon C (1996) “Interaction between self motion and depth perception in the processing of optic flow” *Trends in Neurosci* **19** (5), 196-202

Commuri S & Lewis F (1997) “CMAC neural networks for control of nonlinear dynamical systems: structure, stability and passivity” *Automatica* **33** (4), 635-641

Conant R (1972) “Detecting subsystems of a complex system” *IEEE Trans Syst Man & Cyber* 2 (4), 550-553

Conant R (1974) “Information flows in hierarchical systems” *Int J Gen Syst* 2, 9-18

Conant R (1976) “Laws of information governing systems” *IEEE Trans Syst Man & Cyber* 6 (4), 240-255

Conrad M (1970) “Statistical and hierarchical aspects of biological organisation” in *Towars a Theoretical Biology* 3 (ed Waddington C), 189-221

Corchs S & Deco G (2001) “Selective attention in visual search: a neural network of phase oscillators” **38-40**, 1151-1160

Courtois P (1975) "Decomposability, instabilities and saturation in multiprogramming systems" *Comm Assoc Comp Mach* 18(7), 371-390

Courtois P (1985) "On time and space decomposition of complex structures" *Comm Assoc Comp Mach* 28(6), 590-603

Crabbe J (1994) "Morphogenesis and training of trees: questions from physiology to biomechanics." *Biomimetics*2, 193-198

Crawley E & de Luis J (1987) "Use of piezoelectric actuatora as elements of intelligent structures" *AIAA J* **25** (10), 1373

Crawley E (1994) "Intelligent structures for aerospace: a technology overview and assessment" *AIAA Journal* **32** (8), 1689-1699

Cruse H, Kindermann T, Schumm M, Dean J, Schultz J (1998) "Walknet – a biologically inspired network to control six-legged walking" *Neural Networks* **11**, 1435-1447

Cruse, H., Dean, J., Durr, V., Kindermann, Th., Schmitz, J., Schumm, M. (2002) "A Decentralized, Biologically Based Network for Autonomous Control of (Hexapod) Walking" In *Neurotechnology for Biomimetic Robots* (ed. Ayers J, Davis J, Rudolph A), Bradford Books, MIT Press, Cambridge, MA, USA, 383-400

Cruse, H., Dean, J., Durr, V., Kindermann, Th., Schmitz, J., Schumm, M. (2000) "Control of Hexapod Walking – A Decentralized Solution Based on Biological Data" *RTO AVT Symposium on Unmanned Vehicles (UV) for Aerial, Ground and Naval Military Operations*, 9-13 October, Ankara, Turkey

Cudeiro J & Rivadulla C (1999) "Sight and insight – on the physiological role of nitric oxide in the visual system" *Trends in Neurosci* **22** (3), 109-116

D'haeseleer, P., Forrest, S. and Helman, P. (1996) "An Immunological Approach to Change Detection: Algorithms, Analysis and Implications" *Proceedings of 1996 IEEE Symposium on Research in Security and Privacy*, 110-119

Dahl V (1983) "Logic programming as a representation of knowledge" *IEEE Computer* (Oct), 106-111

Dale K & Collett T (2001) "Using artificial evolution and selection to model insect navigation" *Curr Biology* **11** (17), 1305-1316

Daley S & Gill K (1986) "Design study for self-organising fuzzy logic controller" *Proc IMechE* 200C(1), 59-69

Damasio A (1992) "Brain and language" *Sci Amer* 267 (3), 63-71

Damper R, French R & Scutt T (2000) "ARBIB: an autonomous robot based on inspiration from biology" *Robot & Auton Syst* **31**, 247-274

Dario P, Valleggi R, Carrozza M, Montesi M & Cocci M (1992) "Microactuators for microrobots: a critical survey" *J Micromechanical Microengineering* **2**, 141-157

Das H et al (1999) "Robot manipulator technologies for planetary exploration" *Proc 6th Ann Int Symp on Smart Structures & Materials* (No. 3668-17), Newport Beach, California, USA

Das S (1997) "A biologically motivated neural network architecture for visuomotor control" *Info Sci* **96**, 27-45

Dasgupta, D. (editor) (1999) *Artificial Immune Systems and Their Applications*, Springer-Verlag, Berlin

Dasgupta, D. and Forrest, S. (1995) "Tool breakage detection in milling operations using a negative selection algorithm" *Technical Report No. CS95-5*, Dept. of Computer Science, University of New Mexico, USA

Dasgupta, D., Cao, Y. and Yang, C. (1999) "An Immunogenetic Approach to Spectra Recognition", *Proceedings of the International Genetic and Evolutionary Computation Conference (GECCO) 1999*, 149-155, July 13-17

Dautenkahn K (1995) "Getting to know each other – artificial social intelligence for autonopmous robots" *Robot & Auton Syst* **16**, 333-316

Davidson R (1992) “Smart composites: where are they going?” *Materials & Design* **13** (2), 87-91

Davis G & Patel N (2000) “Origin and evolution of segmentation” *Trends in Biochem Sci* **24** (12), M68-M72

Davis R (1985) “Logic programming and prolog: a tutorial” *IEEE Software* (Sep), 53-62

Daya B & Chauvet G (1999) “On the role of anatomy in learning by the cerebellar cortex” *Math Biosci* **155**, 111-138

De Angelis G (2000) “Seeing in 3 dimensions: the neurophsiology of stereopsis” *Trends in Cog Sci* **4** (3), 80-90

De Aragon M (1994) “Future applications of micro/nanotechnologies in space systems” *ESA Bull* 95, 99-107

de Castro, L. N. & Von Zuben, F. J. (2000b), “The clonal selection algorithm with engineering applications”, *Proceedings of the International Genetic and Evolutionary Computation Conference 2000.*

de Castro, L. N. and Von Zuben, F. J. (2000a) *Artificial immune systems: part II – a survey of applications,* Technical report DCA-RT 02/00, State University of Campinas, Brasil

de Castro, L. N. and Von Zuben, F. J. (2001) “aiNet: An Artificial Immune Network for Data Analysis” *Data Mining: A Heuristic Approach,* Editors: Hussein A. Abbass, Ruhul A. Sarker and Charles S. Newton, Idea Group Publishing, USA

de Castro, L. N., Von Zuben, F. J. and de Deus Jr., G. A. (2003) “The construction of a Boolean competitive neural network using ideas from immunology” *Neurocomputing* **50**, 51-85

De Focatiis D & Guest S (2002) "Deployable membranes designed from folding tree leaves." *Philosophical Transactions of the Royal Society A*360, 227-238

De Garis H (1991) “Genetic programming: building artificial nervous systems with genetically programmed neural network modules” in *Neural & Intelligent Systems Integration* (ed. Soucek B), John Wiley & Sons, 207-234

De Garis H (1996) “CAM-BRAIN: the evolutionary engineering of a billion neuron artificial brain by 2001 which grows/evolves at electronic speeds inside a cellular automata machine (CAM)” *Lecture Notes in Computer Science – Towards Evolvable Hardware* **1062**, 99-117, Springer-Verlag

De Jong K (1980) “Adaptive system design: a genetic approach” *IEEE Trans Systems Man & Cybernetics* 10 (9), 566-574

De Kleer (1986) “Assumption-based truth maintenance systems” *Artif Intell* 28, 124-162

De Kleer J & Brown J (1986) “Theories of causal ordering” *Artif Intell* 29, 33-61

De Rossi D, Domenici C & Chiarelli P (1988) “Analogs of biological tissues for mechanoelectrical transduction: tactile sensors and muscle-like actuators” *NATO ASI Series* **F43** (Sensors & Sensory Systems for Advanced Robots), Springer-Verlag, Berlin, 201-218

Decker K (1987) “Distributed problem-solving techniques: a survey” *IEEE Trans Systems Man & Cybernetics* 17 (5), 729-740

Delcomyn F & Nelson M (2000) “Architecture for a biomimetic hexapod robot” *Robot & Auton Syst* **30**, 5-15

Delcomyn F (1980) “Neural basis of rhythmic behaviour in animals” *Sci* 210, 492-497

Dellaert F. and R.D. Beer (1994). Toward an evolvable model of development for autonomous agent synthesis, in Proceedings of the Forth Conference on Artificial Life, (R. Brooks, and P. Maes, eds.), Cambridge, MA: MIT Press.

Delon J & Germaine R (2000) "Information transfer at the immunological synapse" *Curr Biol* **10**, R923-R933

Denning P & Tichy W (1990) "Highly parallel computation" *Science* 250, 217-1222

Denning P (1988) "Computer viruses" *Amer Sci* 76(May/June), 236-238

Denoyer K, Erwin R & Ninneman R (2000) "Advanced smart structures flight experiments for precision spacecraft" *Acta Astronautica* **47** (2-9), 389-397

Dey S, Hooroo R et al (1998) "Electron microscopy and spectroscopical studies on the coloured patches on the wing of a butterfly, *Graphium sarpedon*(Lepidoptera: Papillionidae) with reference to their photobiological and electrical properties." *Pigment Cell Research*11, 1-11

Dhar V & Pope H (1987) "Rule-based versus structure based methods for explaining and generating expert behaviour" *Comm Assoc Comput Machin* 30 (6), 542

Di Natale C, Salimbeni D, Paolesse R, Macagnano A & D'Amico A (2000) "Porphryins-based opto-electronic nose for volatile compounds detection" *Sensors & Actuators* **B65**, 220-226

Dias J, Parecks C, Fonseca I, Araijo H, Batista J & Almeida A (1998) "Simulating prsuit with machine experiments with robots and artificial vision" *IEEE Trans Robot & Autom* **14** (1), 1-18

Dickinson M (1990) "Linear and nonlinear encoding properties of an identified mechanoreceptor on the fly wing measured with mechanical noise stimuli." *Journal of Experimental Biology*151, 219-244

Dickinson M (1999) "Bionics: Biological insight into mechanical design." *Proc Nation Acad Sci* 96, 14208–14209

Dickinson M, Lehmann F-O, Save S (1999) "Wing rotation and the aerodynamic basis of insect flight" *Sci* **284**, 1454-1960

Dickinson M, Tammero L & Tarstino M (2002) "Sensory fusion in freeflight search behaviour of fruit flies" In *Neurotechnology for Biomimetic Robots* (ed. Ayers J, Davis J, Rudolph A), Bradford Books, MIT Press, Cambridge, MA, USA, 573-592

Dickinson T, White J, Kauer J & Walt D (1998) "Current trends in artificial nose technology" *Trends in Biotechnol* **16** (Jun), 250-258

Dickinson, M. H., Farley, C. T., Full, R. J., Koehl, M. A. R., Kram, R., Lehman, S. (2000) "How Animals Move: An Integrative View" *Science* **288** (April 7), 100-106

Dickmanns E (1998) "Vehicles capable of dynamic vision: a new breed of tehnical beings?" *Artif Intell* **103**, 49-76

Dietz V (2002) "Do human bipeds use quadrupedal coordination?" *Trends in Neurosci* **25** (9), 462-467

Dilger, W. (1997), "Decentralized autonomous organization of the intelligent home according to the principle of the immune system", *Proceedings of the IEEE Conference on Systems, Man and Cybernetics 1997*, 351-356.

Dimcheva, N., Horozova, E. & Jordanova, Z. (2002) "An amperometric xanthine oxidase enzyme electrode based on hydrogen peroxide electroreduction" *Zeitschrift für Naturforschung* **57c**, 883-889

Ditto W (1996) "Applications of chaos in biology and medicine" *Chaos & The Changing Nature Of Science & Medicine - An Introduction* (ed Herbert, D) AIP Press, 175-201

Dodds D (1988) "Fuzziness in knowledge-based robotic systems" *Fuzzy Sets & Syst* 20, 179-193

Don Schuilli (2002) "The PowerZyme vision" http://www.powerzyme.com/PowerZyme-Vision-Web1.pdf (accessed 26/07/2004)

Donald, B., Gariepy, L., and Rus, D. (2000). Distributed manipulation of multiple objects using ropes. In *Proc. IEEE Int. Conf. Robotics and Automation*, pages 450–457.

Donnart J-Y & Meyer J-A (1996) "Learning reactive and planning rules in a motivationally autonomous animat" *IEEE Trans Syst Man & Cyber B* **26** (3), 381-396

Dorigo M (1993) "Genetics-based machine learning and behaviour-based robotics: a new synthesis" *IEEE Trans Systems Man & Cybernetics* 23 (1), 141-154

Dorigo M (1995) "Alecsys and autonoMouse: learning to control a real robot by distributed classifier systems" *Mach Learn J* 19 (3)

Dorigo M, Maniezzo V & Colorni A (1996) "Ant system: optimisation by a colony of cooperating agents" *IEEE Trans Systems Man & Cybernetics B: Cybernetics* **26** (1), 29-41

Doyle J (1979) "Truth maintenance system" *Artif Intell* 12, 231-272

Drabble B (1991) "Spacecraft command and control using AI techniques" *J British Interplanetary Soc* 44, 251-254

Dracopoulos D & Jones A (1993) "Neuromodels of analytic dynamic systems" *Neural Computing & Applications* 1(4), 268-279

Dreyfus H & Dreyfus S (1987) "Why computers can never think like people" *Technol Rev*, 43-61

Dreyfus H (1967) "Alchemy and AI" *RAND Corporation* P-3244

Driver J & Spence C (2000) "Multisensory perception: beyond modularity and convergence" *Curr Biol* **10** (20), R731-R735

Dror G & Tsodyks M (2000) "Chaos in neural networks with dynamic synapses" *Neurocomput* **32-33**, 365-370

Dubowsky S (1999) "Self-transforming robotic planetary explorers" NASA Institute for Advanced Concepts, Phase 1 Study Final Report

Duda R, Hart P & Nilson N (1976) "Subjective Bayesian network for rule-based inference systems" *Proc AFIPS Comput Conf* 45, 1072-1082

Dudek, G., Jenkin, M., Milios, E., and Wilkes, D. (1995). Experiments in sensing and communication for robot convoy navigation. In *Proc. IEEE Int. Conf. on Intelligent Robots and Systems(IROS)*, pages 268–273, Pittsburgh, PA.

Dudek, G., Jenkin, M., Milios, E., and Wilkes, D. (1996). A taxonomy for multi-agent robotics. *Autonomous Robots*, 3(4):375–397.

Dudley R (1998) "Atmospheric oxygen, giant Palaeozoic insects and the evolution of aerial locomotor performance" *J Exp Biol* **201**, 1043-1050

Duke T (2002) "Power of hearing" *Physics World* (May), 29-33

Durfee E & Montgomery T (1991) "Coordination as a distributed search in a hierarachical behaviour space" *IEEE Trans Syst Man & Cyber* 21 (6), 1363-1377

Durfee et al (1989) "Trends in cooperative distributed problem solving" *IEEE Trans Know & Data Eng* 1 (1), 63-83

Duysens J, van der Grommert H, Smits-Engelsman B, van der Helm F (2002) “Walking robot called human: lessons to be learned from neural control of locomotion” *J Biomechanics* **35**, 447-453

Eckhard F et al (2002) “MUSE, a lab-on-a-chip system for in-situ analysis” *53rd International Astronautics Congress (World Space Congress)*, Houston, Texas (Oct), paper IAC-02-IAA.12.1.05

Edelman G & Reeke G (1982) “Selective networks capable of representative transformations, limite generalisations nd associative memory” *Proc Nation Acad Sci* **79**, 2091-2095

Edelman G (1987) *Neural Darwinism: the Theory of Neuronal Group Selection*, Basic Books, New York

Edelman G (1990) The Remembered Present: A Biological Theory of Consciousness.

Edelman S & Intrator N (1999) “Computational models of perceptual learning” *preprint*

Edelman S & Newell F (1998) “On the representation of object structure in human vision: evidence from differential priming of shape and location” Cognitive Science Computer Science Research Report 500, University of Sussex, UK

Edelman S (1996) “Representation is representation of similarities” Technical Report CS96-08, Mathematrics & Computer Science, Weitzmann Institute of Science

Edelman S (1997) “Computational theories of object recognition” *Trends in Cog Sci* **1** (8), 296-304

Edelman S (2002) “Constraining the neural representation of the visual wrld” *Trends in Cog Sci* **6** (3), 125-131

Egelhaaf M, Kern R, Krapp H, Kretzberg J, Kurtz & Warzecha A-K (2002) “Neural encoding of behaviourally relevant visual-motion information in the fly” *Trends in Neurosci* **25** (2), 96-102

Egerstedt, M. and Hu, X. (2001). Formation constrained multi-agent control. *IEEE Transactions on Robotics and Automation*, 17(6):947–951.

Eggenberg, P. Evolving morphologies of simulated 3D organisms based on differential gene expression. In P. Husbands and I. Harvey, editors, *4th European Conference on Artificial Life*, pages 205–213, Cambridge, MA, 1997. The MIT Press/Bradford Books.

Eldredge N & Gould S (1972) “Punctuated equilibria: an alternative to phyletuc gradualism” in *Models in Palaeobiology* (ed. Schopf T), Freeman, San Fransisco

Ellery A (2000) *An Introduction to Space Robotics*. Praxis-Springer series on Astronomy & Space Sciences, Praxis Publishers Ltd, Chichester, UK

Ellington C (1999) "The novel aerodynamics of insect flight: applications to micro-air vehicles." *Journal of Experimental Biology* 202, 3439-3448

Evans, M.C.W. & Bredenkamp, G. (1990) “The structure and function of the photosystem I reaction center” *Physiol. Plant.* **79**, 415-420.

Fahleman S & Hinton G (1987) "Connectionist architectures for AI" *IEEE Comp*(Jan), 100-108

Fakushima K et al (1983) “Neocognitron: neural network model for a mechanism of visual pattern recognition” *IEEE Syst Man & Cyber* 13 (5), 826-834

Falini G, Fermani S et al (1997) "Biomimetic crystallization of calcium carbonate polymorphs by means of collagenous matrices." *Chemistry - A European Journal*3, 1807-1814

Farinelli, A., Grisetti, G., Iocchi, L., and Nardi, D. (2002). Coordination in dynamic environments with constraint on resources. In *IROS Workshop on Cooperative Robotics*, Lausanne, Switzerland

Farmer, J. D., Packard, N. H. and Perelson, A. S. (1986) "The immune system, adaptation, and machine learning" *Physica D* **22**, 187-204

Fatmi M & Young R (1970) "A definition of intelligence" *Nature* 228*,* 1008

Fearing R (1990) "Tactile sensing mechanisms" *Int J Robot Res* **9** (3), 3-28

Fearing, R. S., Avadhanula, S., Campolo, D., Sitti, M., Yan, J., Wood, R. (2002) "A Micromechanical Flying Insect Thorax" In *Neurotechnology for Biomimetic Robots* (ed. Ayers J, Davis J, Rudolph A), Bradford Books, MIT Press, Cambridge, MA, USA, 469-480

Feather N (1967) "Expectancy-value model of information seeking behaviour" *Psychol Rev* 74 (5), 342-360

Feldman D 'Fuzzy network synthesis with genetic algorithms'. Proceedings of International Conference on Genetic algorithms*,* 1993, pp. 312-317

Feldman J & Ballard D (1982) "Connectionist models and their properties" *Cog Sci* 6, 205-254

Feldman J et al (1988) "Computing with structural neural networks" *IEEE Computer* (Mar), 91-102

Fellman R (1988) "Reasoning with worlds and truth maintenance in a knowledge-based programming environment" *Comm Assoc Comp Mach* 31 (4), 382-401

Feng Q, Cui F et al (2000) "Crystal orientation, toughening mechanisms and a mimic of nacre." 11, 19–25

Ferrell C (1995) "A comparison of three insect-inspired locomotion controllers" *Robot & Auton Syst* **16**, 135-159

Fettiplace R, Ricci A & Hackney C (2001) "Clues to the cochlear amplifier from the turtle ear" *Trends in Neurosci* **24** (3), 169-175

Ficici, S. G., Watson, R. A., and Pollack, J. B. (1999). Embodied Evolution: A response to challenges in Evolutionary Robotics. In *Proceedings of the 8th European Workshop on Learning Robots*, pages 14.22, Lausanne, Switzerland.

Fikes R & Kehlar T (1985) "Role of frame based representation in reasoning" *Comm Assoc Comput Machin* 28 (9), 904-920

Fikes R, Hart P & Nilson N (1972) "Learning and executing generalised robot plans" *Artif Intell* 3, 251-288

Fischer G, Cox A, Gogola M, Gordon K, Lobontin N, Monopoli D, Garcia E & Goldfarb M (1999) "Elastodynamic locomotion mesoscale robotic insects" *Proc SPIE Conf on Electroactive Polymer Actuators & Devices* **3669**, Newport Beach, California, 362-368

Fischer K (1980) "Theory of cognitive development: the control and construction of hierachies of skills" *Pyschol Rev* 87 (6), 477-523

Floreano, D. and Mondada, F. (1996) Evolution of Plastic Neurocontrollers for Situated Agents. In P. Maes, M. Mataric, J-A. Meyer, J. Pollack, and S. Wilson. (Eds.), From *Animals to Animats 4. Proceedings of the Fourth International Conference on Simulation of Adaptive Behavior*, Cambridge, MA: MIT Press.

Floreano, D. and Mondada, F. (1998) Evolutionary Neurocontrollers for Autonomous Mobile Robots. *Neural Networks*, 11, 1461-1478.

Floreano, D. and Nolfi, S. (1997a). "Adaptive behavior in competing co-evolving species" Husbands, P. and Harvey, I., editors, *Proceedings of the 4th European Conference on Artificial Life*, MIT Press, Cambridge, Massachusetts, USA, 378-387

Floreano, D. and Nolfi, S. (1997b). "God save the Red Queen! Competition in co-evolutionary robotics" Koza, J. R., Deb, K., Dorigo, M., Foegel, D., Garzon, B., Iba, H., and Riolo, R. L., editors, *Proceedings of the 2nd Annual Conference on Genetic Programming*, pages 398-406. Morgan Kaufmann.

Floreano, D. and Urzelai, J. (1999) "Evolution of Adaptive-Synapse Controllers" D. Floreano et al. (Eds.), *Advances in Artificial Life. Proceedings of the 5th European Conference on Artificial Life*, Berlin: Springer Verlag.

Floreano, D. and Urzelai, J. (2000) "Evolutionary Robots with on-line self-organization and behavioral fitness". *Neural Networks* **13,** 431-443.

Floreano, D., Nolfi, S. and Mondada, F. (2001) "Co-Evolution and Ontogenetic Change in Competing Robots" M. Patel, V. Honavar, and K. Balakrishnan (eds.), *Advances in the Evolutionary Synthesis of Intelligent Agents*, Cambridge (MA): MIT Press.

Floreano, D., Nolfi, S., and Mondada, F. (1998). "Competitive co-evolutionary robotics: From theory to practice" Pfeifer, R., Blumberg, B., Meyer, J.-A., and Wilson, S. W., editors, *Proceedings of the 4th International Conference on Simulation of Adaptive Behavior*, pages 512-524. MIT Press.

Floreano, D., Schoeni, N., Caprari, G. and Blynel, J. (2002) Evolutionary Bits'n'Spikes. In R.K.Standish, M. A. Beadau and H. A. Abbass, editors. Artificial Life VIII: Proceedings of the Eight International Conference on Artificial Life, MIT Press.

Fodor J (1985) “Precis of The Modularity of Mind” *Behav & Brain Sci* **8**, 1-42

Fogel D (1992) "Using evolutionary programming for modelling" *IEEE J Oceanic Eng* 17 (4), 333-339

Font´an, M. S. and Matari´c, M. J. (1998). Territorial multi-robot task division. *IEEE Transactions on Robotics and Automation*, 14(5).

Fontana W & Buss L (1994) “Arrival of the fittest: towards a theory of biological organisation” *Bull Math Biol* 56 (1), 1-64

Forbus K (1984) “Qualitative process theory” *Artif Intell* **24**, 85-168

Forbus K (1987) “Interpreting observations of physical systems” *IEEE Trans Systems Man & Cybernetics* 17 (3), 350-359

Forrest S, Hofmeyer S & Samayaji A (1996) “Computer immunology” *Commun Assoc Comput Machin* **40** (10), 88-96

Forrest, S., Hofmeyr, S. A., Somayaji, A. and Longstaff, T. A. (1996) “A sense of self for Unix processes” *Proceedings of the 1996 IEEE Symposium on Security and Privacy*, 120-128, IEEE Computer Society Press, Los Alamitos, California, USA

Forrest, S., Perelson, A. S., Allen, L. and Cherukuri, R. (1994) “Self-nonself discrimination in a computer” *Proceedings of 1994 IEEE Symposium on Research in Security and Privacy*

Fox M (1981) "Organisational view of distributed systems" *IEEE Trans Syst, Man & Cyber* 11 (1), 70-80

Franceschini N (1996) “Engineering applications of small brains” *FED J* **7** (2), 38-52

Franceschini N, Pichon J & Blanes C (1992) “From insect vision to robot vision” *Phil Trans Roy Soc Lond* **B337**, 283-294

Franks N (1989) “Army ants: a ciollective intelligence” *Amer Sci* **77** (Mar/Apr), 139-145

Franks N, Wilby A, Silverman B & Tofts C (1992) “Self-organising nest construction in ants: sophisticated building by blind bulldozing” *Anim Behav* **44**, 357-375

Franz M & Mallot H (2000) “Biomimetic robot navigation” *Robot & Auton Syst* **30**, 133-153

Fredslund, J. and Matari´c, M. J. (2002). Robots in formation using local information. *Proc 7th Int. Conf. Intelligent Autonomous Systems.*

Freedman P (1991) “Time, Petri nets and robotics” *IEEE Trans Robotics & Automation* 7 (4), 417-433

Freeman W (1991) "The physiology of perception" *Scientific American* (Feb), 34-41

Freeman W (1994) “Qualtitative overview of population neurodynamics” *Neural Modelling & Neural Networks*, 185-215

Freitas A & Gilbreath W (1980) “Advanced automation for space missions” *NASA CP-2255*

Fritz M, Belcher A et al (1994) "Flat pearls from biofabrication of organized composites on inorganic substrates." *Nature* 371, 49-51

Fritz W et al (1989) “The autonomous intelligent system” *Robotics & Auton Syst* 5, 109-125

Fu K (1971) “Learning control systems and intelligent control systems” *IEEE Trans Automatic Control* 16 (Feb), 70-72

Fu L (1993) “Knowledge-based connectionism for revising domain theories” *IEEE Trans Systems Man & Cybernetics* **23** (1), 173-182

Fujita H (1996) “Future of actuators. *Sensors & Actuators A* **56**, 105-111

Fujita H (2000) “What can MEMS do for robotics?” *Preprint*

Fukai T (1999) “Modelling the interplay of short-term memory and the basal ganglia in sequence processing” *Neurocomput* **26-27**, 687-692

Fukuda T & Nakaguma S (1978) “Approach to the dynamically reconfigurable robotic system (DRRS)” *J Intell & Robot Syst* 1, 55-72

Fukuda T, Buss M, Hosokim H & Kawauchi I (1991) “Cell structured robotic system CEBOT: control, planning and communication methods” *Robot & Autonom Syst* 7, 239-248

Fukuda, T., Mori, K. and Tsukiama, M. (1999) “Parallel search for multi-modal function optimization with diversity and learning of immune algorithm” *Artificial Immune Systems and Their Applications*, 210-220, Editor: D. Dasgupta, Springer-Verlag

Fukunaga, A. S., Morookian, J. M., Quillin, K., Stoics, A., Thjakoor, S., (1998) “Earthwormlike Exploratory Robots” NASA TECH BRIEF. 22(6): Item #138) from JPL New Technology Report NPO-20266, June 1998.

Fukushima K et al (1982) “Neocognitron – neural network model for a mechanism of visual pattern recognition” *IEEE Trans Syst, Man & Cyber* 13 (5), 826-834

Full R (2000) “Biological inspiration: lessons from many-legged locomotors” in Robotics Research 9th Int Symp (ed. Hollerbach J, Koditschek D), Springer, London, 337-341

Funes P & Pollack J (1998) “Evolutionary body building: adaptive physical designs for robots” *Artif Life* 4 (4)

Furuta, T., Tawara, T., Okumura, Y., Shimizu, M., Tomiyama, K. (2001) "Design and construction of a series of compact humanoid robots and development of biped walk control strategies" *Robotics and Autonomous Systems* **37**, 81-100.

G. McHale and P. Husbands (2004a) Quadrupedal locomotion: GasNets, CTRNNs and Hybrid CTRNN/PNNs compared. In J. Pollack et al. (Eds*), Proc. Alife IX*, MIT Press, 2004 (in press).

G. McHale and P. Husbands (2004b) GasNets and other evovalble neural networks applied to bipedal locomotion. In S. Schaal et al. (Eds), *Proc. From Animals to Animats 8: Proceedings of the Eigth International Conference on Simulation of Adaptive Behaviour (SAB'2004),* MIT Press, 2004 (in press).

Gadanho S & Hallam J (1998) "The role of emotions: explaining autonomy mechanisms in mobile robotics" *Department of AI Research paper* **851**, University of Edinburgh, UK

Gaines B (1976) "On the complexity of causal models" *IEEE Trans Syst, Man & Cyber* (Jan), 56-59

Galeotti, J., Rhody, S., Nelson, A., Grant, E., and Lee, G. (2002). EvBots. The Design and Construction Of A Mobile Robot Colony for Conducting Evolutionary Robotic Experiments. In *Proceedings of the ISCA 15th International Conference: Computer Applications in Industry and Engineering (CAINE-2002)*, pp. 86-91, San Diego Ca, Nov. 7-9, 2002.

Galipeau D, Story P, Vetelino K & Mileham R (1997) "Surface acoustic wave microsensors and applications" *Smart Mater Struct* **6**, 658-667

Gallagher J, Beer R, Esenschied K & Quinn R (1996) "Application of evolved locomotion controllers for a hexapod robot" *Robot & Auton Syst* **19**, 95-103

Gallant S (1988) "Connectionist expert system" *Comm Assoc Comput Machin* 31 (2), 152-169

Game G (1993) "Application of neural networks to vibration suppression in spacecraft structures" *Proc 2nd Int Conf Dynamics & Control of Structures in Space*, Cranfield Institute of Technology, Bedfordshire, UK, 529-543

Gao Y and Er M, "Online Adaptive Fuzzy Neural Identification and Control of a Class of MIMO Nonlinear Systems," *IEEE Transaction on Fuzzy Systems*, Vol. 11, No. 4, pp. 462-477, 2003.

Garfinkel A, Spano M, Ditto W & Weiss J (1992) "Controlling cardiac chaos" *Science* 257 (Aug), 1230-1235

Gaussier P & Zrehen S (1995) "PerAc: a neural architecture to control artificial animals" *Robot & Auton Syst* **16**, 291-320

Gaussier P, Joulais C, Banquet J, Lepetre S & Revel A (2000) "Visual homing problem: an example of robotics/biology cross fertilisation" *Robot & Auton Syst* **30**, 155-180

Gaussier P, Revel A, Joulian C, Zehren S (1997) "Living in partially structured environments: how to bypass the limitations of classical reinforcement techniques" *Robot & Auton Syst* **20**, 225-250

Gazzaniga M (1989) "Organisation of the human brain" *Sci* 245, 947-952

Gegenfurtner K & Hawken M (1996) "Interaction of motion and colour in the visual pathways" *Trends in Neurosci* **19** (9), 394-401

Geim, A. K., Dubonos, S. V., Grigorieva, I. V., Novoselov, K. S., Zhukov, A. A. & Shapoval, S. Y. (2003) "Microfabricated adhesive mimicking gecko foot-hair" *Nature Materials* **2**, 461-463

Gelenbe E, Schmajuk N, Staddon J & Reif (1997) "Autonomous search by robots and animals: a survey" *Robot & Auton Syst* **22**, 23-34

Gennes P-Gd & Okamura K (2000) "On the toughness of biocomposites." *C. R. Acad. Sci. Paris, t. 1, Série IV,: Solids, fluids: mechanical and thermal properties*, 257–261

Georgeff M & Lansky L (1986) “Procedural knowledge” *Proc IEEE* 74 (10), 1388-1397

Georgopoulos A (1995) “Current issues in directional motor control” *Trends in Neurosci* **18**, 506-510

Gerard M, Chaubey A & Malhotra B (2002) “Application of coducting polymers to biosensors” *Biosensors & Bioelectronics* **17**, 345-359

Gerardin L (1968) *Bionics*. London, Weidenfeld & Nicolson.

Geschwind N (1972) “Language and the brain” *Sci Amer* (Apr), 76-83

Geschwind N (1979) “Specialisations of the human brain” *Sci Amer* 241 (3), 158-168

Ghandi H et al (1991) “Smart materials and structures incorporating hybrid actuator and sensing systems” *ASME Smart Structures & Materials* **24**, 151-156

Ghiselin M (1978) “The economy of the body” *Am Econ Assoc* 68 (2), 233-237

Giachino J (1986) “Smart sensors” *Sensors & Actuators* 10, 239-248

Gilardi G & Fantuzzi A (2001) “Manipulatig redox systems: application to nanotechnology” *Trends in Biotechnol* **19** (11), 468-476

Gilardi G & Fantuzzi A (2001) “Manipulating redox systems: application to nanotechnology” *Trends in Biotechnol* 19 (11), 468-476

Gilbertson R & Busch J (1996) “Survey of micro-actuator technologies for future spacecraft missions” *J British Interplanetary Society* **49**, 129-138

Gillespie P (1995) “Molecular machinery of auditory and vestibular transduction” *Curr Opinion in Neurobiol* **5**, 449-455

Gisztér, S. F., (2002) “Biomechanical Primitives and Heterarchical Control of Limb Motion in Tetrapods” In *Neurotechnology for Biomimetic Robots* (ed. Ayers J, Davis J, Rudolph A), Bradford Books, MIT Press, Cambridge, MA, USA, 223-240

Giszter, S. F., Moxon, K. A., Rybak, I., Chapin, J. K. (2000) “A Neurobiological Perspective on Humanoid Robot Design” *IEEE Intelligent Systems* **15**(4), 64-69.

Glazer, A.N. & Melis, A. (1987) “Photochemical reaction centers : structure, organization and function” *Annual Review of Plant Physiology* **38**, 11-45.

Gluck M, Myers C & Thompson R (1995) “A computational model of the cerebellum and motor-reflex conditioning” in *An Introduction to Neural & Electronic Networks*, Academic Press Inc, 91

Gold J & Shadlen M (2001) “Neural computations that underlie decisions about sensory stimuli” *Trends in Cog Sci* **5** (1), 10-16

Goldberg D & Holland J (1988) “Genetic algorithms and machine learning” *Mach Learn* 3, 95-98

Goldberg D Genetic algorithms in search, optimization and machine learning, Reading, MA: Addison-Wesley, 1989

Goldberg M & Colley C (1992) “Oculomotor control and spatial processing” *Curr Opinions in Neurobiol* **2**, 198-202

Goldberg, D. and Matari´c, M. J. (1997). Interference as a tool for designing and evaluating multirobot controllers. In *AAAI-97*, pages 637–642.

Goldin D, Venneri S et al (November 2000) "The great out of the small" *Mechanical Engineering,* 71-79

Golman-Rakic P (1992) “Working memory and the mind” *Sci Amer* 267 (3), 72-79

Goodale M & Humphrey G (1998) “The objects of action and perception” *Cog* **67**, 181-207

Goodwyn P & Gorb S (2003) "Attachment forces of the hemelytra-locking mechanisms in aquatic bugs (Heteroptera: Belostomatidae)." *Journal of Insect Physiology*49, 753–764

Gopel W (1996) “Ultimate limits in the miniaturisation of chemical sensors” *Sensors & Actuators A* 56, 83-102

Gorb S & Popov V (2002) "Probabilistic fasteners with parabolic elements: biological system, artificial model and theoretical considerations." *Philosophical Transactions of the Royal Society A*360, 211-225

Gorb S & Scherge M (2000) "Biological microtribology: anisotropy in frictional forces of orthopteran attachment pads reflects the ultrastructure of a highly deformable material." *Proceedings of the Royal Society B*267, 1239-1244

Gorb S (1996) "Design of insect unguitractor apparatus." *Journal of Morphology*230, 219-230

Gorb S (1997) "Porous channels in the cuticle of the head-arrester system in dragon/damselflies (Insecta: Odonata)." *Microscopy Research and Technique*37, 583-591

Gorb S (1998) "Frictional surfaces of the elytra-to-body arresting mechanism in Tenebrionid beetles (Coleoptera: Tenebrionidae): design of co-opted fields of microtrichia and cuticle ultrastructure." *International Journal of Insect Morphology and Embryology*27, 205-225

Gorb S (1998) "The design of the fly adhesive pad: distal tenent setae are adapted to the delivery of an adhesive secretion." *Proceedings of the Royal Society B*265, 747-752

Gorb S (2001) *Attachment devices of insect cuticle*. Dordrecht, The Netherlands, Kluwer.

Gorb S, Gorb E et al (2001) "Scale effects on the attachment pads and friction forces in syrphid flies (Diptera, Syrphidae)." *Journal of Experimental Biology*204, 1421–1431

Gordon J & Jeronimidis G (1980) "Composites with high work of fracture." *Philosophical Transactions of the Royal Society A*294, 545-550

Gough M (1993) “Space instrument neural networks for real-time data analysis” *IEEE Trans Geosci & Remote Sensing* 31 (6), 1264-1268

Gower D & Vincent J (1996) "The mechanical design of the cuttlebone and its bathymetric implications." *Biomimetics*4, 37-57

Graham D (1985) “Pattern and control of walking in insects” in *Advances in Insect Physiology* **18**, 31-140, Academic Press Inc (London)

Graham T & Sarikaya M (2000) "Growth dynamics of red abalone shell: a biomimetic model." *Materials Science and Engineering C*11, 145–153

Grasso F, Consi T, Mountain D & Atema J (2000) “Biomimetic robot lobster performs chemo-orientation in turbulence using a pair of spatially separated sensors: progress and challenges” *Robot & Auton Syst* **30**, 115-131

Green C (1969) "Application of theorem-proving to problem-solving" *Int J Conf Artif Intell*, 219-239

Grefenstette J (1986) “Optimisation of control parameters for genetic algorithms” *IEEE Trans Syst Man & Cyber* 16 (1), 122-128

Gregory R (1970) “On how so little information controls so much behaviour” in *Towards A Theoretical Biology* 3, 136-247 (ed. Waddington C)

Griliches Z (1957) "Hybrid corn: exploration in the economics of technological change" *Econometrica* 25 (4), 501-522

Grillner S, Deliangina T, Ekeberg O, El Manira A, Hill R, Lanser A, Orlovsky G & Wallen P (1995) “Neural networks that co-ordinate locomotion and body orientation in lamprey” *Trends in Neurosci* **18** (6), 270-279

Grossberg S (1980) "How does the brain build a cognitive code?" *Psych Rev* 87(1), 1-51

Grossberg S, Mingella E & Ross W (1997) “Visual brain and visual perception: how does the cortex do perceptual groupings?” *Trends in Neurosci* **20** (3), 106-111

Grossberg S, Mingolla E & Ross W (1997) “Visual brain and visual perception: how does the cortex do perceptual grouping?” *Trends in Neurosci* **20** (3), 106-111

Gruau F (1996) “Artificial cellular development in optimisation and compilation” *Lecture Notes in Computer Science – Towards Evolvable Hardware* **1062**, 76-98, Springer-Verlag

Gruau F., 1994. Automatic definition of modular neural networks, *Adaptive Behavior*, 3:151-183.

Grzywacz N & Poggio T (1995) “Computation of motion by real neurons” in *An Introduction to Neural & Electronic Networks*, Academic Press Inc

Guemez J & Matias M (1993) "Control of chaos in unidimensional maps" *Physics Letters A* 181, 29-32

Guldin N & Gruisser O-J (1998) “Is there a vestibular cortex?” *Trends in Neurosci* **21** (6), 254-359

Gunderson S & Whitney J (1992) "Insect cuticle microstructure and Its applications to advanced composites." *Biomimetics* 1, 177-197

Haak S (1979) “Do we need fuzzy logic?” *Int J Man-Machine Stud* **11**, 437-445

Hackett J & Shah M (1990) “Multisensor fusion: a perspective” *Proc IEEE Int Conf Robotics & Autom*, 1324-1329

Hackney C (1987) “Anatomical features of the auditory pathway from cochlea to cortex” *British Medic Bull* **43** (4), 780-801

Haio A (1986) "Syntactic theory of belief and action" *Artif Intell* 28, 245-292

Hajela, P., & Yoo, J. S. (1999) “Immune Network Modelling in Design Optimization” *New Ideas in Optimization*, 203-215, Editors: D. Corne, M. Dorigo & F. Glover, McGraw Hill, London

Hakura J, Miwa K, Tokoi H, Kakazu Y (1999) “Metamorphic internal image of autonomous agents: emergence of behaviour” *Robot & Auton Syst* **28**, 149-161

Hall B (1991) “Adaptive evolution that requires multiple spontaneous mutations: mutation involving base substitutions” *Proc Nat Acad Sci* 88, 5882-5886

Hallam B & Hayes G (2001) “Comparing robot and animal behaviour” *preprint*

Halme, A., Zhang, X. & Ranta, A. (2000) "Study of Biological Fuel Cells" *2nd Annual Advances in R&D, The Commercialization of Small Fuel Cells and Battery Technologies for Use in Portable Applications*, April 26-28 2000, New Orleans, USA

Halpern J (1997) “A logical approach to reasoning about uncertainty: a tutorial” in *Discourse Interaction & Communication* (ed. Arrazola X, Korta K & Pelletier F), Kluwer

Halpert, G. (1997) "The Direct Methanol Liquid-Feed Fuel Cell" NASA technical report, http://techreports.jpl.nasa.gov/1997/97-1156.pdf (accessed 22/07/2004)

Hammer M (1997) "Neural basis of associative reward learning in honeybees" *Trends in Neurosci* **20** (3), 245-252

Handelman D (1990) "Integrating neural networks and knowledge based systems for intelligent robotic control" *IEEE Control Syst Mag* (Apr), 77-86

Handley S (1993) "Genetic planner: the automatic generation of plans for a mobile robot via genetuc programming" *Proc Inte Symp Intelligent Control*, 190-195

Harmon L (1982) "Automated tactile sensing" *Int J Robot Research* **1** (2), 3-33

Harmon S et al (1986) "Sensor data fusion through a distributed blackboard" *Proce IEEE J Robotics & Autom*, 1449-1454

Harnad S (1990) "Symbol grounding problem" *Physica* 42D, 335-346

Harri A et al (1995) "Aspects of atmospheric science and instrumentation for martian missions" *Adv Space Res* 16 (6), 15-22

Hart, E. & Ross, P. (1999), "An immune system approach to scheduling in changing environments", *Proceedings of the International Genetic and Evolutionary Computation Conference 1999*, 1559-1566, July 13-17

Harth E (1983) "Order and chaos in neural systems: approach to dynamics of higher brain functions" *IEEE Trans Systems Man & Cybernetics 13* (5), 782-789

Harvey I & Thompson A (1996) "Through the labyrinth evolution finds a way: a silicon ridge" *Proc 1st Int Conf on Evolvable Systems*, Tsukuba, Japan, 406-422

Harvey I (1992) "Species adaptation genetic algorithms: the basis for a continuing saga" in *Towards a Practice of Autonomous Systems: Proc 1st European Conf on Artificial Life*, (ed. Varela F & Bourgine P), MIT Press, 346-354

Harvey I (1994) "Evolutionary robotics and SAGA: the case for hill crawling and tournament selection" *Proc Artif Life III* (ed. Langton G), Addison-Wesley, Reading, MA

Harvey I (1997) "Artificial evolution and real robots" *Artif Life & Robotics* 1, 35-38

Harvey I, Husbands P & Cliff D (1993) "Issues in evolutionary robotics" *Proc 2nd Int Conf on Simulation of Adaptive Behaviour* (ed. Meyer J-A, Roitblat H & Wilson S), MIT Press, Cambridge, MA, USA

Harvey I, Husbands P & Cliff D (1994) "Seeing the light: artificial evolution, real vision" *From Animals to Animats III: Proc 3rd Int Conf on Simulation of Adaptive Behaviour*, Brighton, UK, 392-401

Harvey I, Husbands P, Cliff D, Thompson A & Jakobi N (1997) "Evolutionary robotics: the Sussex approach" *Robotics & Autonomous Syst* 20, 205-224

Hasslecher B & Tilden M (1995) "Living machines" *Robot & Auton Syst* **15**, 143-169

Hata K, Kokubo T et al (1995) "Growth of a bone-like apatite layer on a substrate by a biomimetic process." *Journal of the American Ceramic Society*78, 1049-1053

Hay J (1965) "Optical motions and space perception" *Psych Rev*, 550-559

Hayes P (1973) "Frame problem and related problems in AI" *Int J Conf Artif Intell*, 223-230

Hayes-Roth B (1985) "Blackboard architecture for control" *Artif Intell* 26, 251-321

Hayes-Roth F & Jacobstein N (1994) "State of knowledge-based systems" *Comm Assoc Comput Machin* 37, 27-39

Hayes-Roth F (1984) "Knowledge-based expert systems: a tutorial" *IEEE Computer* (Sep), 11-28

Hayes-Roth F (1985) "Rule-based systems" *Comm Assoc Comput Machin* 28 (9), 921-932

Hazel J, Fuchigami N et al (2001) "Ultramicrostructure and microthermomechanics of biological IR detectors: materials properties from a biomimetic perspective." *Biomacromolecules* 2, 304-312

He G & Mai K (1999) "Biological macromolecules and molecular recognition in mollusk biomineralization." *Progress In Biochemistry And Biophysics* 26, 310-312

Heathcote P, Fyfe P & Jones M (2002) "reaction centres: the structure and evolution of biological solar power" *Trends in Biochem Sci* **27** (2), 79-87

Hedenstrom A (2002) "Aerodynamics, evolution and ecology of avian flight" *Trends in Ecol & Evol* **17** (9), 415-422

Heller, A. (2004) "Minature Biofuel Cells" *Physical Chemistry, Chemical Physics* **6**, 209-216

Hemmi H, Mizoguchi J & Shimohara (1996) "Development and evolution of hardware behaviours" *Lecture Notes in Computer Science – Towards Evolvable Hardware* **1062** (ed. Sanchez E & Tomassini), Springer-Verlag

Hendler J et al (1990) "AI planning: systems and techniques" *AI Mag* (Summer), 61-77

Hendrix G & Sacerdoti E (1982) "Natural language processing – the field in perspective" *Byte* (Sept), 304-352

Hendry S & Calkins D (1998) "Neuronal chemistry and functional organisation in the primate visual system" *Trends in Neurosci* **21** (8), 344-349

Herbert R, Young P et al (2000) "The hind wing of the desert locust (*Schistocerca gregaria*Forskål). III. A finite element analysis of a deployable structure." *Journal of Experimental Biology*203, 2945-2955

Hewitt (1985) "Challenge of open systems" *Byte* (April), 223-242

Heyes C (2001) "Causes and consequences of imitation" *Trends in Cog Sci* **5** (6), 253-261

Hightower R, Forrest S & Prelson A (1996) "Baldwin effect in the immune system: learning by somatic hypermutation" in Adaptive Individuals in evolving Populations: Models & Algorithms (ed. Belew R & Mitchell M), Addison Wesley, Reading, MA, 159-167

Hightower R, Forrest S, Perelson A (1993) "Evolution of emergent organisation in immune system gene libraries" Proc European Conf on Artificial Life, Brussels

Hill A (1938) "The heat of shortening and the dynamic constants of muscle" *Proc Royal Society of London* **126B**, 136-195

Hill A (1953) "Mechanics of active muscles" *Proc Royal Society f London* **141B**, 104-117

Hill W et al (2000) "Using microtechnologies to build micro-robot systems" *Proc Advanced Space Technologies for Robotics & Automation 2000*, ESTEC, Noordwijk, Holland

Hink R & Woods D (1987) "How humans process uncertain knowledge " *AI Mag*(Dec), 41-51

Hinton G & Nowlan S (1987) "How learning can guide evolution" *Complex Syst* 1, 495-502

Hinton G (1985) "Learning in parallel networks" *Byte* (Nov), 265-273

Hinton G (1991) “Connectionist learning procedures” *Artif Intell* **40**, 185-234

Hinton G (1992) “How neural networks learn from experience” *Sci Amer* 267 (3), 105-109

Hiong F (1997) “Molecular sensors based on the photoelectric effect of bacteriorhodopsin: origin of differential responsivity” *Mater Sci & Eng* **C4**, 267-285

Hirano S, Yamamoto K et al (1997) "Calcium chloride as a biomimetic intermediate for the mineralization of carbonate ions of water as calcium carbonate in gelatinous matrices of chitosan and chitin." *Energy Conversion and Management*38, S517-S521

Hirose S & Morishima A (1990) “Design and control of a mobile robot with an articulated body” *Int J Robot Res* **9** (2), 99-114

Hirose S (1993) “Biologically inspired robots: snake like locomotors and manipulators” *Oxford University Press*, Oxford, UK

Hirsch M & Baird B (1995) “Computing with dynamic attractors in neural networks” *Biosyst* **34**, 173-195

Hirzinger, G., Sporer, N., Schedl, M., Butterfab, J., Grebenstein, M. (2004) “Torque-Controlled Lightweight Arms and Articulated Hands: Do We Reach Technological Limits Now?” *The International Journal of Robotics Research*, April. **23**(4) 331-340(10).

Hisdal E (1988) “Philosophical issues raised by fuzzy set theory” *Fuzzy Sets & Syst* **25**, 349-356

Ho Y & Pong T-C (1996) “Cooperative fusion of stereo and motion” *Patt Recog* **21** (1), 121-130

Ho Y (1980) “Team decision theory and information structure” *Proc IEEE* 68 (6), 644-654

Hoare C (1983) "Axiomatic basis for computer programming" *Comm Assoc Comp Mach* 26(1), 53-56

Hoffmann, G. W. (1986) “A neural network model based on the analogy with the immune system”, *Journal of Theoretical Biology* **122**, 33-67

Hofmeyr, S. and Forrest, S. (2000) "Architecture for an Artificial Immune System" *Evolutionary Computation* **8**(4), 443-473

Hogan J et al (1985) “Impedance control: an approach to manipulation – Parts I-III” *ASME J Dynamic Systems Measurement & Control* **107** (Mar), 1-24

Hohle U & Stout L (1991) “Foundations of fuzzy sets” *Fuzzy Sets & Syst* **40**, 257-296

Holland J (1962) “Outline for a logical theory of adaptive systems” *J Assoc Comp Mach* 9, 297-314

Holland J (1986) “Mathematical framework for studying learning in classifier systems” *Physica* 22D, 307-317

Holland J (1992) “Genetic algorithms” *Sci Amer* 267 (1), 72-80

Holland, O. and Melhuish, C. (1999). Synergy, self-organization, and sorting in collective robotics. *Artificial Life*, 5(2):173–202.

Hollingworth G, Tyrell A & Smith S (2000) “To evolve in a changing environment” *preprint*

Holyoak K & Thagard P (1989) “Analogical mapping by constraint satisfaction” *Cog Sci* **18**, 295-355

Hopfield J & Tank D (1986) “Computing with neural networks: a model” *Sci* 233, 625-633

Hopfield J (1982) “Neural networks and physical systems with emergent collective computational capabilities” *Proc Nat Aca Sci* 79, 2554-2558

Hopfield J (1984) “Neurons with graded response have collective computational proprties like those of two-state neurons” *Proc Nat Acad Sci USA* 31, 3088-3092

Horikawa S, Furuhashi T, Uchikawa Y, and Tagawa T (1991) ‘A study on fuzzy modeling using fuzzy neural networks’. Proceedings of International Symposium on *Fuzzy Engineering*

Horning R & Johnson B (2002) “Polymer-based MEMS actuators for biomimetics” In *Neurotechnology for Biomimetic Robots* (ed. Ayers J, Davis J, Rudolph A), Bradford Books, MIT Press, Cambridge, MA, USA, 173-186

Horning, R., Johnson, B. (2002) “Polymer-Based MEMs Actuator for Biomimetics” In *Neurotechnology for Biomimetic Robots* (ed. Ayers J, Davis J, Rudolph A), Bradford Books, MIT Press, Cambridge, MA, USA, 173-186

Hoshino M, Kashimori Y, Kambara T (1996) “Self-organised phase transitions in neural networks as a neural mechanism of information processing” *Proc Nat Acad Sci USA* (Apr), 3303-3307

Houk J, Buckingham J & Barto A (1996) “Models of the cerebellum and motor learning” *Behav & Brain Sci* **19** (3), 368-383

Houk J, Singh S, Fisher C & Barto A (1990) “Adaptive sensorimotor network inspired by the anatomy and physiology of the cerebellum” in *Neural networks for Control* (ed. Miller W, Sutton R, Werbos P), MIT Press, Cambridge, Mass, 301-348

Hsaio D & Kennel M (1989) "Heterogeneous databases: proliferation, issues and solutions" *IEEE Trans Know & Data Eng* 1,(1),45-62

Huang I, Ohnishi N et al (1995) "A biomimetic system for localization and separation of multiple sound sources." *IEEE Transactions on Instrumentation and Measurement*44, 733-738

Huang Q, Yokoi K, Kajita S, Kaneko K, Arai H, Koyachi N, Tanie K (001) “Planning walking pattens for a bipedal robot” *IEEE Trans Robot & Autom* **17** (3), 280-289

Huang, P., Magnusson, A., Lomoth, R., Abrahamsson, M., Tamm, M., Sun, L., van Rotterdam, B., Park, J., Hammarstrom, L., Akermark, B., Styring, S. (2002) "Photo-induced oxidation of a dinuclear $Mn_2^{II,II}$ complex to the $Mn_2^{III,IV}$ state by inter-and intra-molecular electron transfer to RuIII tris-bipyridine" *J. Inorg Biochem* **91**, 159-172.

Hubel D & Wiesel T (1962) “Receptive fields, binocular interaction and functional arcuitecture in the cat’s visual cortex” *J Physiol* **160**, 106-154

Hubel D & Wiesel T (1977) “Functional architecture of macaque monkey visual cortex” *Proc Roy Soc London* **198**, 1-59

Huber S, Franz M & Buthoff H (1999) “On robots and fleas: modelling the visual orientation behaviour of flies” *Robot & Auton Syst* **29**, 227-242

Hubermann B & Hogg T (1987) “Phase transitions in AI systems” *Artif Intell* 33, 155-171

Humphreys M (1995) “W-learning: computation among self Q-learners” *Univ Cambridge Computer Laboratory Tech Report* **362**

Hunt E (1978) “Intelligence as an information processing concept” *British J Psychol* **71**, 449-474

Hunt E (1989) "Cognitive science: definition, status and questions" *Ann Rev Psych* 40,603-629

Hunt, J. E. and Cooke, D. E. (1996) “Learning Using an Artificial Immune System” *Journal of Network and Computer Applications* **19**, 189-212.

Hunt, J. E. and Fellows, A. (1996) "Introducing an Immune Response into a CBR system for Data Mining", *Proceedings of the BCS ESG 1996 Conference,* published as *Research and Development in Expert Systems XIII*

Hunter P (2004) "Flower power" *N Scient* (May), 28-31

Husbands I, Harvey I, Cliff D & Miller G (1994) "Use of genetic algorithms for the development of sensorimotor control systems" in *Proc of Conf From Perception To Action* (ed. Nicoud J-D & Gaussier P), IEEE Computer Press, Los Alamitos, CA

Husbands P et al (1995) "Circle in the round: state space attractoirs for evolved sight robots" *Robot & Auton Syst* 15, 83-106

Husbands, P. (1998) Evolving Robot Behaviours with Diffusing Gas Networks, In: P. Husbands and J.-A. Meyer (1998), 71-86.

Husbands, P. and J.-A. Meyer (Eds) (1998*) EvoRobot98: Proceedings of 1st European Workshop on Evolutionary Robotics*, Springer-Verlag LNCS 1468.

Husbands, P. and T. Smith and N. Jakobi and M. O'Shea. Better Living through Chemistry: Evolving GasNets for Robot Control, *Connection Science*, **10,** (3&4), 185-210, 1998.

Hutt L et al (1999) "Microfabricated capillary electrophoresis amino acid chirality analyser for extraterrestrial exoloration" *Analyt Chem* 71, 4000-4006

Ijspeert, A. J., Martinoli, A., Billard, A., and Gambardella, L. M. (2001). Collaboration through the exploitation of local interactions in autonomous collective robotics: the stick pulling experiment. *Autonomous Robotics*, 11(2):149-171

Ikoma T, Kobayashi H et al (2003) "Microstructure, mechanical, and biomimetic properties of fish scales from *Pagrus major*." *Journal of Structural Biology*142, 327–333

Ilg W, Berns K, Muhlfreidel Th, Dillman R (1997) "Hybrid learning concepts based on self-organiing neural networks for adaptive control of walking machines" *Robot & Auton Syst* **22**, 317-327

Imahori, H., Mori, Y. & Matano, Y. (2003) "Nanostructured artificial photosynthesis" J. *Photochem. Photobiol.* **4**, 51-83.

Ingber D (1998) "Architecture of life" *Sci Amer* (Jan)

Iocchi, L., Nardi, D., and Salerno, M. (2001). Reactivity and deliberation: a survey on multirobot systems. In Hannebauer, M., Wendler, J., and Pagello, E., editors, *Balancing Reactivity and Social Deliberation in Multi-Agent Systems, From RoboCup to Real-World Applications*, volume 2103 of *Lecture Notes in Computer Science*, pages 35–52. Springer.

Ishida, Y. (1996), "Agent-based architecture of selection principle in the immune system", *Proceedings of the International Workshop on Immunity-Based Systems 1996*, 92-104, Kyoto, Japan

Ito M (1970) "Neurophysiological aspects of the cerebellar motor control system" *Int J Neurol* **7**, 162-176

Iverson L (1979) "Chemistry of the brain" *Sci Amer* 241 (3), 118-129

Iwaniuk A & Whishaw I (2000) "On the origin of skilled forelimb movements" *Trends in Neurosci* **23** (8), 372-376

Iwasaki Y & Simon H (1986) "Causality in device behaviour" *Artif Intell* 29, 3-32

Jabri M et al (2000) "Models of basal ganglia and cerebellum for sensorimotor integration and predictive control" *Proc SPIE Conf Sensor Fusion & Decentralised Control in Robotic Systems III* (ed. McKee G & Schenker P) 4196, 125-137

Jabri M, Huang J, Oliver J-M, Coenen D, Sejnowski T (2000) "Models of basal ganglia and cerebellum for sensorimotor integration and predictive control" *Proc SPIE Sensor Fusion & Decentralised Control in Robotic Syst III*, ed. McKee G & Schenker P, **4196**, 125-137

Jacob F (1977) "Evolution and tinkering" *Sci* 196, 1161-1166

Jacobs R (1999) "Biologically inspired legged robots for space operations" NASA Institute for Advanced Concepts, Phase 1 Study Final Report

Jacobs R (1999) "Computational studies of the development of functionally specialised neural modules" *Trends in Cog Sci* 3 (1), 31-38

Jaeger H (1997) "From continuous dynamics to symbols" *Proc 1st Joint Conf on Complex Syst in Psychology*, Gstaad, Switzerland

Jakobi, N. (1998). *Minimal Simulations for Evolutionary Robotics*. PhD thesis, University of Sussex, U.K.

Jakobi, N. (1998b) Evolutionary Robotics and the Radical Envelope of Noise Hypothesis, *Adaptive Behavior*, **6(2)**: 325-368.

Jakobi, N. and Quinn, M. (1998). Some problems (and a few solutions) for open-ended evolutionary robots. In Husbands, P. and Meyer, J.-A., editors, *Evolutionary Robotics: Proc.* 1*st Euro.Workshop, EvoRobot98*, pages 108–122. Springer.

Jang S, Sun C and Mizutani E (1997) Neuro-Fuzzy and Soft Computing, New Jersey: Prentice Hall

Jaynes E (1968) "Prior probabilities" *IEEE Syst Sci Cyber* 4 (3), 227-244

Jeannerod M, Arbib M, Rissolotti G & Sakata H (1995) "Grasping objects: the cortical mechanisms of visuomotor transformation" *Trends in Neurosci* **18** (7), 314-320

Jennings N & Compos J (1997) "Towards a social level characterisation of socially responsible agents" *preprint*

Jennings N & Wooldridge M (1996) "Software agents" *IEE Rev* (Jan), 17-20

Jeong B & Gutozoska A (2002) "Lessons from nature: stimuli-responsive polymers and their biomedical applications" *Trends in Biotechnol* **20** (7), 305-311

Jeong I & Lee J (1997) "Evolving cooperative mobile robots using a modified genetic algorithm" *Robot & Auton Syst* **21**, 197-205

Jerbic B, Grolinger K & Vranjes B (1999) "Autonomous agent based o reinforcement learning and adaptive shadowed network" *Artif Intell in Engin* **13**, 141-157

Jerne, N. K. (1974) "Towards a network theory of the immune system" *Annals of Immunology* **125C**, 373–389

Jernimidis G & Atkins A (1995) "Mechanics of biological materials and structures: Nature's lessons for the engineer" *Proc Inst Mechan Eng* **209**, 221-235

Jimenez M & Gonzalex de Santos P (1998) "Attitude and position control method for realistic legged vehicles" *Robot & Auton Syst* **18**, 345-354

Johansson S (1995) "One approach towards the fabrication of a microrobot system" *Materials Science & Engineering* **C2**, 141-149

Johnson-Laird P & Steedman P (1978) “Psychology of syllogism” *Cog Psychol* 10, 64-92

Johnson-Laird P (1980) “Mental models in cognitive science” *Cog Sci* **4**, 71-115

Jones A (1991) “Genetic algorithms and their application to the design of neural networks” *preprint*

Jones A (1996) “New ideas in neural networks” *Research Report*, Imperial College London

Jones W & Hoskins J (1987) “Backpropagation” *Byte* (Oct), 155-162

Juang C and Lin C, ``An On-Line Self-Constructing Neural Fuzzy Inference Network and its Applications," IEEE Trans. Fuzzy Systems, Vol. 6, pp.12-32, 1998

Jun, J.-H., Lee, D.-W. and Sim, K.-B. (1999) “Realization of Cooperative and Swarm Behavior in Distributed Autonomous Robotic Systems Using Artificial Immune System” *Proceedings of the IEEE Conference on Systems, Man and Cybernetics 1999* **4**, 614-619.

Kaelbling L (1991) "Foundations of learning in autonomous agents" *Rob & Auton Syst* 8, 131-144

Kaelbling L et al (1996) “Reinforcement learning: a survey” *Journ Artif Intell Res* 4, 237-285

Kaiser M & Dillman R (1997) “Hierarchical refinement of skills and skill application for autonomous robots” *Robot & Auton Syst* **19**, 259-271

Kalman, L., LoBrutto, R., Allen, J.P. & Williams, J.C. (1999) "Modified reaction centers oxidize tyrosine in reactions that mirror photosystem" II. *Nature* **402**, 696-699.

Kalman, L., LoBrutto, R., Allen, J.P. & Williams, J.C. (2003) "Manganese oxidation by modified reaction centers from Rhodobacter spheroids" *Biochemistry* **42**, 11016-11022.

Kaminka, G., Lima, P., and Roja, R., editors (2003). *Proc. RoboCup-2002 International Symposium*. Springer Verlag.

Kandel E & Schwartz J (1982) “Molecular biology of learning: modification of transmitter release” *Science* **218**, 433-442

Kandel E & Squire L (1992) “Cognitive neuroscience” *Current Biology* **2** (2), 143-145

Kandel E (1992) “Biological basis of learning” *Sci Amer* 267 (3), 53-69

Kaplan D (1998) "Mollusc shell structures: novel design strategies for synthetic materials." *Current Opinion in Solid State & Materials Science*3, 232-236

Karam G & Gibson L (1994) "Biomimicking of animal quills and plant stems: cylindrical shells with foam cores." *Materials Science and Engineering C*2, 113-132

Karmiloff-Smith A (1994) “Precis of Beyond modularity: a development al perspective on cognitive science” *Behav & Brain Sci* **17**, 693-745

Karr C and Gentry E (1993) ‘Fuzzy control of pH using genetic algorithms’, *IEEE Trans. Fuzzy Syst.* **1** (l), pp. 4653

Karr C, Belew R and Booker L (1991) ‘Genetic algorithm based fuzzy control of spacecraft autonomous rendezvous’. Proceedings of international conference on *Genetic Algorithms* (Morgan Kauffman)

Kato T & Amamiya T (1999) "A new approach to organic/inorganic composites. Thin film coating of CaCO3 on a chitin fiber in the presence of acid-rich macromolecules." *Chemistry Letters*3, 199-200

Kato, T. and Floreano, D. (2001) **"**An Evolutionary Active-Vision System" *Proceedings of the Congress on Evolutionary Computation*, Piscataway, IEEE-Press

Katti D, Katti K et al (2001) "3D finite element modeling of mechanical response in nacre-based hybrid nanocomposites." *Computational and Theoretical Polymer Science*11, 397-404

Katz J & Fodor J (1963) “Structure of a semantic theory” *Language* **39**, 170-210

Katz, E. (2001) "Biofuel Cells – a Review" http://chem.ch.huji.ac.il/~eugeniik/biofuel/biofuel_cells_contents.html (accessed 20/07/2004)

Katz, E., Willner,, I. & Kotlyar, A. B. (1999) "A non-compartmentalized glucose|O_2 biofuel cell by bioengineered electrode surfaces" *Journal of Electroanalytical Chemistry* **479**, 64–68

Kavli T, ``ASMOD - An Algorithm for Adaptive Spline Modeling of Observation Data," Int. J. Control, Vol. 58, pp. 947-967, 1993

Kawato M (2000) “Robotics as a tool for neuroscience: cerebellar internal models for robotics and cognition” *Proc SPIE Conf Sensor Fusion & Decentralised Control in Robotic Systems III* (ed. McKee G & Schenker P) 4196, 321-328

Kawato M et al (1988) “Hierarchical neural network model for voluntary movement with applications to robotics” *IEEE Control Syst Mag* (Apr), 8-15

Kawato M, Furukawa K & Suzuki R (1987) “Hierarchical neural network model for control and learning of voluntary movement” *Biolog Cybernet* **57**, 169-185

Kelly J (1989) "Intelligent machines - what chance?" in *Advances in AI* (ed. Hallam J & Mollish R), Wiley & Son

Kelly, I. D. and Keating, D. A. (1996). Flocking by the fusion and active infrared sensors on physical autonomous mobile robots. In *Proc. 3rd Int. Conf. Mechatronics and Machine Vision*, volume 1, pages 1–4, Guimaraes, Portugal.

Kemper T (1987) “How many emotions are there? Wedding the social and autonomic components” *Amer J Sociology* 93 (2), 263-289

Kendall K (1981) "Interfacial dislocations in tough adhesive composites." *Philosophical Transactions of the Royal Society A*43, 713-729

Kennedy L & Basu M (1997) “Image enhancement using a human visual system model” *Pattern Recog* **30** (12), 2001-2014

Kent E (1978) “Brains of men and machines: machinery of emotion and choice” *Byte* (April), 66-89

Kephart J (1994) “A biologically inspired immune system for computers” *ALife IV*

Kephart J et al (1994) “Computers and epidemiology” *IEEE Spectrum* (May), 20-26

Kephart J et al (1997) “Fighting computer viruses” *Sci Am* (Nov)

Kephart, J. O. (1994) “A Biologically Inspired Immune System for Computers”, *Artificial Life IV, Proceedings of the Fourth International Workshop on Synthesis and Simulation of Living Systems*, 130-139, Editors: Rodney A. Brooks and Pattie Maes, MIT Press, Cambridge, Massachusetts

Kessler H, Ballarini R et al (1996) "Biomimetic example of brittle toughening .1. Steady-state multiple cracking." *Computational Materials Science*5, 157-166

Khalid M, Omatsu S, and Yusef R (1994) ‘Adaptive fuzzyneuro control with application to a water bath process’, Proceedings of IEEE International Conference on *Control applications,* Strathclyde, UK, pp. 173-178

Khatib O (1985) “Real time obstacle avoidance for manipulators and mobile robots” *Proc IEEE Int Conf Robotics & Autom*, 500-505

Kihlstrom J (1987) “Cognitive unconscious” *Science* 237, 1445-1451

Kim J & Thompson R (1997) “Cerebellar circuits and synaptic mechanisms involved in classical eyeblink conditioning” *Treends in Neurosci* **20**, 177-181

Kim J and Kasabov N (1999) ``HyFIS: Adaptive Neuro-Fuzzy Inference Systems and Their Application to Nonlinear Dynamical Systems," Neural Networks, Vol. 12, pp. 1301-1319

Kim S (1989) “Systematic approach to intelligent systems design” *Robotics & CIM* 6 (2), 143-155

Kim, J. and Bentley, P. J. (2002) "Towards an Artificial Immune System for Network Intrusion Detection: An Investigation of Dynamic Clonal Selection" *Proceedings of the Congress on Evolutionary Computation (CEC-2002),* 1015 – 1020, May 12-17, Honolulu, USA

Kimura H, Fukuota Y & Nakaruma H (2000) “Biologically inspired adaptive walking of the quadruped on irregular terrain” *preprint*

Kirchner, F., Spenneberg, D., Linnemann, R. (2002) “A Biologically Inspired Approach Toward Robust Real-World Locomotion in Legged Robots” In *Neurotechnology for Biomimetic Robots* (ed. Ayers J, Davis J, Rudolph A), Bradford Books, MIT Press, Cambridge, MA, USA, 419-448

Kirkpatrick S, Gellatt C & Vecchi M (1983) “Optimisation by simulated annealing” *Science* 320, 671-680

Kirsch D (1991a) "Foundations of AI: big issues" *Artif Intell* 47, 3-30

Kirsch D (1991b) "Today the earwig, tomorrow the man" *Artif Intell* 47, 161-184

Kirschner F (1998) “Q-learning of complex behaviours on a six-legged walking machine” *Robot & Auton Syst* **25**, 252-262

Kitano H (1990) Designing neural networks using genetic algorithms with graph generation system, *Complex Systems*, 4:461-476.

Klein M (1991) "Supporting conflict resolution in cooperative design systems" *IEEE Syst, Man & Cyber* 21 (6), 1379-1389

Kleinrock L (1985) "Distributed systems" *Comm Assoc Comp Mach* 28 (11), 1200-1213

Klopf A (1988) “Neuronal model of classical conditioning” *Psychobiol* **16** (2), 85-125

Klopf A, Weaver S, Morgan J (1993) “hierarchical network of control systems that learn: modelling nervous system function during classical and instrumental conditioning” *Adapt Behav* **1** (3), 263-319

Knierim J, Kudrimoti H & McNaughton B (1995) “Place cells, head direction cells, and the learning of landmark stability” *J Neurosci* **15** (3), 1648-1659

Knight D & Feng D (1992) "Formation of the dogfish egg capsule, a coextruded, multilayer laminate." *Biomimetics* **1**, 151-176

Kohl M, Krevet B & Just E (2002) “SMA gripper system” *Sensors & Actuators A* **97/98**, 646-652

Kohonen T (1988) “Introduction to neural computing” *Neural Networks* 1, 3-16

Kohonen T (1990) “Self-organising map” *Proc IEEE* 78 (9), 1464-1479

Kokubo T (1995) "Apatite formation on organic polymers by a biomimetic process." *European Journal of Solid State and Inorganic Chemistry*32, 819-827

Kokubo T (1996) "Formation of biologically active bone-like apatite on metals and polymers by a biomimetic process." *Thermochimica Acta*280/1, 479-490

Kolacinski R & Quinn R (1998) "A novel biomimetic actuator system" *Robot & Auton Syst* **25**, 1-18

Koob T & Summers A (2002) "Tendon - bridging the gap." *Comparative Biochemistry and Physiology A-Comparative Physiology*133, 905-909

Koopman B, Grootenboer H & Jough H (1995) "Inverse dynamics model for the analysis, reconstruction and prediction of bipedal walking" *J Biomechanics* **28** (11), 1369-1376

Korkin M, de Garis H, Gers F & Hemmi H (1997) "CBM (CAM-Brain Machine): a hardware tool which evolves a neural net module in a fraction of a second and runs a million neuron artificial brain in real time" *Proc 2nd Annual Conf on Genetuc Programming*

Kornbluh R, Full R, Meijer K, Pelrine R & Shastri S (2002) "Engineering a muscle: an approach to artificial muscle based on field-activated electroactive polymers" *Neurotechnology for Biomimetic Robots* (ed. Ayers J, Davis J, Rudolph A), Bradford Books, MIT Press, Cambridge, MA, USA, 137-172

Kornbluh, R., Full, R. J., Meijer, K., Pelrine, R., Shastri, S. V. (2002) "Engineering a Muscle: An Approach to Artificial Muscle Based on Field-Activated Electroactive Polymers" In *Neurotechnology for Biomimetic Robots* (ed. Ayers J, Davis J, Rudolph A), Bradford Books, MIT Press, Cambridge, MA, USA, 137-172

Kornfield W & Hewitt C (1981) "Scientific community metaphor" *IEEE Trans Syst Man & Cybern* 11, 24-33

Kosko B (1992) Neural Networks and Fuzzy Systems, New Jersey: Prentice Hall

Kosslyn, S. (1994) Image and Brain. Cambridge, MA. MIT Press.

Kotha S, Li Y et al (2001) "Micromechanical model of nacre tested in tension." *Journal of Materials Science*36, 2001-2007

Kouskoff V & Nemazee D (2001) "Role of editing and revision in shaping the B and T lymphocyte repertoire" *Life Sci* **69**, 1105-1113

Koza J (1991) "Genetic evolution and co-evolution of computer programs" *Artificial Life II*, Santa Fe Inst 10 (ed. Langton C), 603-628

Koza J (1993) "Evolution of subsumption using genetic programming" *preprint*

Koza J, Bennett F, Andre D & Keane M (1998) "Evolutionary design of analog electrical circuits using genetic programming" *preprint*

Koza J, Bennett F, Hutchings J, Bade S, Keane M & Andre D (1998) "Evolving computer programs using rapidly reconfigurable field-programmable gate arrays and genetic programming" Proc 6th Int Symp on Field Programmable Gate Arrays

Kraft G & Campagna D (1990) "Comparison of CMAC neural network control and two traditional adaptive control systems" *IEEE Control Syst Mag* (Apr), 26-43

Kraulis R & Stone L (1999) "Tracking with the mind's eye" *Trends in Neurosci* **22** (12), 544-549

Kraus S (1997) "Negotiation and cooperation in multi-agent environments" *Artif Intell* **94**, 79-97

Kresling B (1995) "Plant 'Design' : Mechanical simulations of growth patterns and bionics." *Biomimetics*3, 105-120

KrishnaKumar, K. and Neidhoefer, J. (1999) "Immunized Adaptive Critics for an Autonomous Aircraft Control Application" *Artificial Immune Systems and Their Applications* 221-241, Editor: D. Dasgupta, Springer-Verlag

Kriz D, Ramstrom et al (1995) "Introducing biomimetic sensors based on molecularly imprinted polymers as recognition elements." *Analytical Chemistry* 67, 2142-2144

Krogh B & Thorpe C (1986) "Integrated path planning and dynamic steering control for autonomous vehicles" *Proc IEEE Int Conf Robotics & Autom*, 1664-1667

Krus P (1997) "Natural methods for flight stability in birds" *World Aviation Congress and exposition*, SAE AIAA.

Kube, R. C. and Bonabeau, E. (1999). Cooperative transport by ants and robots. *Robotics and Autonomous Systems*, 30(1/2):85-101

Kube, R. C. and Zhang, H. (1994). Collective robotic intelligence. In Meyer, J., H., R., and Wilson, W., editors, *Proc. 2rd Int. Conf. Simulation of Adaptive Behaviour*. MIT Press, London.

Kuipers B (1986) "Qualitative simulation" *Artif Intell* 29, 289-337

Kuipers B (1987) "Qualitative simulation as causal explanation" *IEEE Trans Systems Man & Cybernetics* 17 (3), 432-443

Kundur S Ravir D (1998) "Vision-based pragmatic strategy for autonomous navigation" *Pattern Recog* **31** (9), 1221-1239

Kuperstein M (1987) "Adaptive visual-motor coordination in multi-joijnted robots using parallel architectures" *Proc IEEE Int Conf Robot & Autom,*1675-1682

Kweon I et al (1992) "Behaviour-based mobile robot using active sensor fusion" *Proc IEEE Int Conf Robotics & Autom*, 1675-1682

Laane, C., Pronk, W., Franssen, M. & Veeger, C. (1984) "Use of a bioelectrochemical cell for the synthesis of (bio)chemicals" *Enzyme and Microbial Technology* **6**(4), 165-168

Lacher R et al (1992) "Backpropagation learning in expert systems" *IEEE Trans Neural Networks* 3 (1), 62-72

Lachinov A, Zherebov A & Kornilov V (1996) "New effects in electroactive polymers: new basics for sensors" *Sensors & Actuators A* **53**, 319-324

Lagoutte, B. & Mathis, P. (1989) "The photosystem I reaction center" *Photochemistry and Photobiology* **49**, 833-844.

Laird J, Newell A & Rosenbloom P (1987) "SOAR: an architecture for general intelligence" *Artif Intell* **33**, 1-64

Lakes R (1993) "Design considerations for materials with negative Poisson ratios" *Trans ASME J Mechan Design* **115** (Dec), 696-700

Lambert Ch, Nusgens B & Lapiere Ch (1998) "Mechano-sensing and mechano-reaction of soft connective tissue cells" *Advances in Space Research* **21** (8/9), 1081-1091

Lambrinos D, Moller R, Lebhart T, Pfeifer R & Wehner R (2000) "Mobile robot employing insect strategies for navigation" *Robot & Auton Syst* **30**, 39-64

Lang H et al (1999) "Artificial nose based on a micromechanical cantilever array" *Anal Chimica Acta* 393, 59-65

Langton C (1988) “Artificial life” in *Artificial Life* (ed. Langton C), Addison-Wesley, Santa Fe Institute, Reading, MA, 1-47

Lanyi J & Pohorille A (2001) “Protin pumps: mechanism of action and applications” *Trends in Biotechnol* **19** (4), 140-144

Lanyi J & Pohorille A (2001) “Proton pumps: mechanism of action and applications” *Trends in Biotechnol* 19 (4), 140-144

Lappe M, Bremmer F & van den Berg A (1999) “Perception of self motion from visual flow” *Trends in Cog Sci* **3** (9), 329-336

Lauder G (2001) “Flight of the robofly” *Nature* **412**, 688-689

Laurent G (1996) “Dynamical representation of odours by oscillating and evolving neural assemblies” *Trends in Neurosci* **19** (11), 489-496

Lawes R (1998) “Microsystems and how to access the technology” *J British Interplanetary Soc* 51, 127-134

Layzell P (1998) “New research tool for intrinsic hardware evolution” In *Proc. ICES'98*, ser. Springer-Verlag Lecture Note in Computer Science, pp. 47—56.

Layzell P (1998) “Evolvable motherboard: a test platform for the research of intrinsic hardware evolution” *University of Sussex Tech Report CSRP479*, University of Sussex, UK

Layzell P (1999) “Reducing hardware evolution’s dependency on FPGAs” Proc 7th Int Conf Microelectronics for Neural, Fuzzy & Bio-Inspired Systems (MicroNeuro 99)

Lazarus R (1982) “Thoughts on relations between emotion and cognition” *Amer Psychol* 37 (9), 1019-1024

Lee C (1990) ``Fuzzy Logic in Control Systems: Fuzzy Logic Controller,'' IEEE Trans. System, Man and Cybernetic, Vol. 20, pp. 404-436

Lee G (1996) “System identification and control of smart structures using neural networks” *Acta Astronautica* **38** (4-8), 269-276

Lee M & Nicholls H (1999) “Tactile sensing for mechatronics – a state of the art survey” *Mechatronics* **9**, 1-31

Lee M and Takagi H (1993) ‘Integrating design stage of fuzzy systems using genetic algorithms’. Proceedings of IEEE International Conference on *Fuzzy Systems,* San Francisco, CA, USA, pp**.** 612-617

Lenat D & Feigenbaum E (1991) "On the threshold of knowledge" *Artif Intell* 47, 185-250

Lenat D (1978) “Ubiquity of discovery” *Artif Intell* 9, 257-285

Lenat D (1983) “Theory formation by heuristic search – nature of heuristics II” *Artif Intell* 21, 31-59

Lenat D (1984) “Why AM and Eurisko appear to work” *Artif Intell* 23, 269-294

Lenat D (1990) "CYC: towards programs with common sense" *Comm Assoc Comp Mach* 33, 30-49

Lenneberg E (1960) "Capacity for language acquisition" in *Culture in History* (ed Diamond S), 869-893

Lenneberg E (1964) "Biological perspective on language" in *New Directions in the Study of Language* (ed. Lenneberg E), MIT Press, 65-88

Lesser V & Corkhill D (1981) "Functionally accurate cooperated distributed systems" *IEEE Trans Syst Man & Cyber* 11 (1), 81-96

Lettvin J, Maturana H, McCullogh W & Pitts W (1959) “What the frog’s eye tells the frog’s brain” *Proc IRE* **47** (Nov), 1940-1951

Levin M (1995) “Evolution of understanding: a genetic algorithm model of the evolution of communication” *Biosyst* **36**, 167-178

Levinson R (1994) “Human frontal lobes and AI planning systems” *Proc 2nd AI Planning Systems Conf* (AIPS 94)

Lewis E (1983) “Elements of a single neuron: a review” *IEEE Trans Systems Man & Cybernetics* 13 (5), 702-709

Li J, Chen J, Lui C (2000) "Micromachined biomimetic sensor using modular artificial hair cells" NanoSpace 2000, Houston, Texas

Li S, Zeng Q et al (1995) "Biomimicry of bamboo bast fiber with engineering composite materials." *Materials Science and Engineering C-Biomimetics Materials Sensors and Systems* 3,125-130

Lichter M, Sujan V & Dubowsky S (2000) “Experimental demonstration of a new design paradigm in space robotics” *Proc 7th Int Symp on Experimental Robotics*, Honolulu, Hawaii

Lieber R (1999) “Skeletal muscle as a biological example of a linear electro-active actuator” *Proc SPIE Conf on Electroactive Polymer Actuators & Devices*, Mar 1999, Newport Beach, CA, USA, **3669**, 19-25

Limburg, J., Vrettos, J.S., Chen, H, de Paula J.C., Crabtree, R.H. and Brudvig, G.W. (2001) Characterization of the O2-evolving reaction catalyzed by $[(terpy)(H_2O)Mn^{III}(O)_2Mn^{IV}(OH_2)(terpy))](NO_3)_3$. *J. Am. Chem. Soc* **123**, 423-430.

Lin C (1994) ``Neural Fuzzy Control Systems with Structure and Parameter Learning," Singapore: World Scientific.

Linkens D and Nyongesa H (1993) ‘Real-time acquisition fuzzy rules using genetic algorithms’. Proceedings of 1992 IFAC symposium on *Artificial Intelligence **in** Real-time Control,* pp. 335-339

Linkens D and Nyongesa H (1996) “Systems in intelligent control : an appraisal of fuzzy, neural and genetic algorithm control applications,” IEE Proc.-Control Theory Appl., Vol. 143, No. 4, pp. 367-386

Linkens D and Nyongesa H, ‘Genetic algorithms for fuzzv control part 2: Online system development and application’, IEE *Proc. Control Theory Appl.,* **142,** (3), pp. 177-185

Linkens D and Nyongesa H, ‘Genetic algorithms for fuzzy control part 1: Offline system development and application’ *IEE Proc. Control Theory Appl.,* **142,** *(3),* pp. 161-176

Linsker R (1986a) “From basic network principles to neural architectures: emergence of orientation cells” *Proc Nat Acad Sci USA* 83, 8779-8783

Linsker R (1986b) “From basic network principles to neural architectures: emergence of orientation-selective cells” *Proc Nat Acad Sci USA* 83, 8390-8394

Linsker R (1986c) “From basic network principles to neural architectures: emergence of spatial opponent cells” *Proc Nat Acad Sci USA* 83, 7506-12

Liu C, Li W et al (1996) "Simulation of nacre with TiC/metal multilayers and a study of their toughness." *Materials science & engineering C-biomimetic materials sensors and systems*4, 139-142

Liu, H. & Logan, B. E. (2004) "Electricity generation using an air-cathode single chamber microbial fuel cell in the presence and absence of a proton exchange membrane" *Environmental Science and Technology* **38**(14), 4040-4046

Liversedge S & Findlay J (2000) “Saccadic eye movements and cognition” *Trends in Cog Sci* **4** (1), 6-13

Llinas R (1979) “Cortex of the cerebellum” *Sci Amer* 241 (3), 56-75

Lourenco C & Babloyantz A (1994) "Control of chaos in networks with delay: a model for synchronisation of cortical tissue" *Neural Computation* 6, 1141-1154

Low B et al (1991) “Connectionist expert system with adaptrive learning capability” *IEEE Trans Knowledge & Data Eng* 3 (2), 200-207

Lumelsky, V. and Harinarayan, K. (1997). Decentralized motion planning for multiple robots: The cocktail-party model. *Autonomous Robots*, (4).

Lumia R & Shahinpoor M (1999) “Microgripper design using electroactive polymers” *Proc SPIE Conf on Electroactive Polymer Actuators & Devices*, Mar 1999, Newport Beach, CA, USA, **3669**, 322-329

Lundstrom I (1996) “Why bother about gas-sensitive field-effect devices?” *Sensors & Actuators A* 56, 75-82

Luo R & McKay M (1989) “Multisensor integration and fusion in intelligent systems” *IEEE Trans Systems Man & Cybernetics* 19 (5), 901-927

Lusted R & Knapp R (1996) “Controlling computers with neural signals” *Sci Am* (Oct)

Lynch G & Boudry M (1984) “Biochemistry of memory: a new and specific hypothesis” *Science* 224, 1057-1063

MacIntosh D et al (1991) "Distributed automated reasoning: issues in coordination, cooperation and performance" *IEEE Trans Syst, Man & Cyber* 21 (6), 1307-1316

MacKay W (1997) “Synchronised neuronal oscillations and their role in motor processes” *Trends in Cog Sci* **1** (5), 176-182

Maes P (1990) “Situated agents can have goals” *Robotics & Autonomous Syst* 6, 49-70

Maestripieri D (1997) “Evolution of communication” *Lang & Comm* 17 (3), 269-277

Magnusson, A., Berglund, H., Korall, P., Hammarstrom, L., Akermark, B., Styring, S., Sun, L. (1997) Mimicking electron transfer reactions in photosystem II: synthesis and photochemical characterisation of a ruthenium(II) tris(bipyridyl) complex with a covalently linked tyrosine. *J. Am. Chem. Soc.* **119**, 10720-10725

Maiers J & Sherif Y (1985) "Applications of fuzzy set theory" *IEEE Trans Syst Man & Cyber* 15(1), 175-189

Malcolm C & Smithers T (1990) “Coordination of robot sharing assembly tasks” *Trans ASME J Dynamic Systems Measurement & Control* 197, 299-307

Malkin, R. (1988) “Structure-function studies of photosynthetic cytochrome b-c$_1$ and b$_6$-f complexes” *ISI Atlas of Science: Biochemistry* **10**, 57-64.

Mallot H (1997) “Behaviour-oriented approaches to cognition: theoretical perspectives” *Theory in Biosciences* **116**, 196-220

Mallot H (1999) “Spatial cognition: behavioural competences, neural mechanisms and evolutionary scaling” *Kognitionswissenschaft* **8**, 40-48

Malone, T. and Crowston, K. (1994). The interdisciplinary study of coordination. *ACM Computer Surveys*, 26(1):81–199.

Mamdani E and Assilian S (1975) ``An Experiment in Linguistic Synthesis with a Fuzzy Logic Controller," Int. J. Man-Machine Studies, Vol. 7, pp. 1-13

Mange D, Goeke M, Madon D (1996) "Embryonics: a new family of coarse-grained field programmable gate array with self-repair and self-reproducing properties" *Lecture Notes in Computer Science – Towards Evolvable Hardware* **1062**, 197-220

Mann D (1999) "Creativity as an exact (biomimetic) science" *TRIZ journal*

Mann S (1993) "Molecular tectonics in biomimeralization and biomimetic materials chemistry." *Nature*365q, 499-505

Mann S (1995) "Biomineralization and biomimetic materials chemistry." *Journal of Materials Chemistry*5, 935-946

Mann S (1996) *Biomimetic materials chemistry*, VCH.

Manna Z & Waldringer R (1975) "Knowledge and reasoning in program synthesis" *Artif Intell* 6, 175-208

Manna Z & Waldringer R (1978) "Logic of computer programming" *IEEE Trans Software Eng* 4 (3), 199-229

Marder E & Bucher D (2001) "Central pattern generators and the control of rhythmic movements" *Curr Biol* **11**, R986-R996

Marder E & Calabrese R (1996) "Principles of rhythmic motor pattern generation" *Physiolog Rev* **76** (3), 687-717

Marefat, M. M., Wu, L. (1996) "Purposeful Gazing and Vergence Conttrol for Active Vision" *Journal of Robotics and Computer Integrated Manufacturing.* **12**(2).

Maris M (2001) "Attention-based navigation in mobile robots using a reconfigurable sensor" *Robot & Auton Syst* **34**, 53-63

Marocco, D. and Floreano, D. (2002) "Active Vision and Feature Selection in Evolutionary Behavioral Systems" Hallam, J., Floreano, D. Hayes, G. and Meyer, J. (Eds) *From Animals to Animats 7.* Cambridge, MA. MIT Press.

Marose S, Lindemann C, Ulber R & Schaper T (1999) "Optical sensor systems for bioprocess monitoring" *Tibtech* 17 (Jan), 30-34

Marr D & Poggio T (1976) "Cooperative computation of stereo disparity" *Science* 194, 283-287

Marr D (1969) "Theory of cerebellar cortex" *J Physiol* **202,** 437-470

Martin P & Millan J (2000) "Robot arm reaching through neural inversions and reinforcement learning" *Robot & Auton Syst* **31**, 227-246

Martinez, Carolina (2002) "Creepy Crawlers May Unravel Web of Planetary Mysteries" *NASA News Release* **232** (18 December)

Martinoli, A. (1999). *Swarm Intelligence in Autonomous Collective Robotics: From Tools to the Analysis and Synthesis of Distributed Collective Strategies.* PhD thesis, DI-EPFL, Lausanne, Switzerland.

Massaro D & Stork D (1998) "Speech recognition and sensory integration" *Amer Sci* (May/June)

Mataric M & Cliff D (1995) "Challenges in evolving controllers for physical robots" *Robotics & Autonomous Syst* **19** (1), 67-83

Mataric M & Cliff D (1997) "Artificial evolution of control systems" *preprint*

Mataric M (1992) “Integration of representation into goal-driven behaviour-based robots” *IEEE Trans Robotics & Autom* 8 (3), 304-312

Mataric M (1994) “Learning to behave socially” in *From Animals to Animats 3: Proc 3rd Int Conf on Simulation of Adaptive Behaviour* (ed. Cliff D et al), MIT Press, Cambridge, MA, 453-462

Mataric M (1995) “Issues and approaches in the design of autonomous agents” *Robotics & Autonomous Syst* 16, 321-331

Mataric M (1996) “Designing and understanding adaptive group behaviour” *Adapt Behav* **4** (1), 51-80

Mataric M (1997) “Learning social behaviour” *Robot & Auton Syst* **20**, 191-204

Mataric, M. J. (1992). Designing emergent behaviours: From local interactions to collective behaviour. In Meyer, J.-A., Roitblat, H., andWilson, W., editors, *Proc. 2rd Int. Conf. Simulation of Adaptive Behaviour*, pages 432–441, Honolulu, Hawaii. MIT Press, London.

Mataric, M. J. (1995). Designing and understanding adaptive group behaviour. *Adaptive Behaviour*, 4(1):51–80.

Mataric, M. J. (1997). Behavior-based control: Examples from navigation, learning and group behaviour. *J. Experimental and Theoretical Artificial Intelligence*, 9(2-3):323–336.

Mataric, M. J., Nilsson, M., and Simsarian, K. (1995). Cooperative multi-robot box-pushing. In *Proc. IEEE/RJS Int. Conf. on Intelligent Robots and Systems*, pages 556–561, Buda, Hungary.

Matias M & Guemez J (1994) “Stabilisation of chaos by proportional pulses in the system variables” *Pysical Rev Lett* 72 (10), 1455-1458

Matzinger, P. (1994) “Tolerance, danger and the extended family” *Annual Review of Immunology* **12**, 991-1045

Matzke C et al (1998) “Microfabricated silicon gas chromatographic micro-channels: fabrication and performance” *Proc SPIE Conf on Micromachining & Microfabrication Process Technology IV,* Santa Clara, Calif (Sept), 3511, 262-268

Mayr E (1974) “Behaviour programs and evolutionary strategies” *Amer Sci* 62, 650-659

McCarthy J & Hayes P (1969) “Some philosophical problems from the standpoint of artificial intelligence” *Machine Intell* **4**, 463-502

McCarthy J (1961) “Programs with common sense” *Meachanisation of Thought Processes* 1, 77-84 (HMSO)

McCarthy J (1969) "Some philosophical problems from the standpoint of AI " *Mach Intell* 4, 463-502

McCarthy J (1977) "Epistemological problems of AI "*Proc Int J Conf Artif Intell,* 459-465

McCarthy J (1980) "Circumscription - a form of non-monotonic reasoning" *Artif Intell* 13, 27-39

McCarthy J (1986) "Application of circumscription to formalising common sense knowledge" *Artif Intell* 28, 89-116

McCarthy J (1987) "Generality in AI " *Comm Assoc Comp Mach* 30 (12), 1030-1035

McCoy, D. F. & Devaralan, V. (1997), “Artificial immune systems and aerial image segmentation”, *Proceedings of the IEEE Conference on Systems, Man and Cybernetics 1997*, 867-872.

McCulloch W & Pitts W (1943) “A logical calculus of the ideas immanent in nervous activity” *Bull Mathem Biophys* **5**, 115-133

McDermott D & Doyle J (1980) "Non-monotonic logic" *Artif Intell* 13, 41-72

McGeer T (1990) "Passive dynamic walking" *Int J Robot Res* **9**, 62-82

McGruer, N. E., Adams, G. G., Truong, T.l Q., Barnes T. G., Lu, X., Aceros, J. C. "Biomimetic Flow and Contaqct / Bending MEMs Sensors" In *Neurotechnology for Biomimetic Robots* (ed. Ayers J, Davis J, Rudolph A), Bradford Books, MIT Press, Cambridge, MA, USA, 481-510

McKay C et al (1998) "Mars oxidant experiment (MOx) for Mars 96" *Planet Space Sci* 46(6/7), 769-777

McLaren R (1966) "Stochastic automaton model for the synthesis of learning systems" *IEEE Trans Systems Science & Cybernetics* 2 (2), 109-114

Mead R & Sedgwick (1992) "Design concepts for self-tunnelling probes" *Proceedings of the International Symposium on Missions, Technologies & Design of Planetary Mobile Vehicles*, Toulouse, (Sept 92) – CNES/Cipaduhs-Iditions, Toulouse, France (ISBN 2854283317)

Melhuish, C. (1999). Exploiting domain physics: Using stigmergy to control cluster building with real robots. In Floreano, D., Nicoud, J.-D., and Mondada, F., editors, *Proc. 5th Euro. Conf. Artificial Life*, pages 585–595. Springer Verlag.

Melhuish, C. (2002) "The EcoBot Project" http://www.ias.uwe.ac.uk/Energy-Autonomy/ecobot_web_page.htm (accessed 22/07/2004)

Melhuish, C., Holland, O., Hoddell, S. (1998) "Collective Sorting and Segregation in Robots with minimal sensing." In Pfeifer, R., Blumberg, B., Meyer, J.-A., and Wilson, S. W., editors, *Proc. 5th Int. Conf. Simulation of Adaptive Behaviour*. MIT Press / Bradford Books, Cambridge, MA. August 17-21.

Menzel R & Giurfa M (2001) "Cognitive architecture of a mini-brain: the honeybee" *Trends in Cog Sci* **5** (2), 62-71

Meredith D, Karr C, and Kumar K (1993) 'The use of genetic algorithms in the design of fuzzy logic controllers'. 3rd Workshop on *Neural Networks* WNN92, Houston, TX, USA, pp. 549-555

Metta G, Manzotti R, Panerai F & Sandini G (2000) "Development: is it the right way towards humanoid robotics?" *Proc 6th Int Conf on Intelligent Autonomous Systems*, Venice, Italy, p. 249

Metta G, Panerai F, Manzotti R & Sandini G (2000) Babybot: an artificial developing robotic agent" in *Proc SPIE,* Boston, USA

Meyer J-A (1997) "From natural to artificial life: biomimetic mechanisms in animat design" *Robot & Auton Syst* **22**, 3-21

Michaud F & Mataric M (1999) "Representation of behavioural history for learning in non-stationary conditions" *Robot & Auton Syst* **29**, 187-200

Michel A & Farwell J (1990) "Associative memories via artificial neural networks" *IEEE Trans Control Syst Mag* (Apr), 6-17

Michel, O. and J. Biondi. (1995) Morphogenesis of neural networks. *Neural Processing Letters*, **2** (1): 9-12.

Michelson, R. C. (2002) "The Entomopter" In *Neurotechnology for Biomimetic Robots* (ed. Ayers J, Davis J, Rudolph A), Bradford Books, MIT Press, Cambridge, MA, USA, 481-510

Michie D (1973) "Machines and the theory of intelligence" *Nature* 241, 507-512

Middlehoek S & Noorlag D (1981) "Silicon micro-transducers" *J Phys E Sci Instrum* 14, 1343-1352

Miessner M, Peter M et al (2001) "Preparation of insect-cuticle-like biomimetic materials." *Biomacromolecules*2, 369-372

Migita, M., Okada, T., Mataga, N., Nishitani, S., Kurata N., Sakata, Y., Misumi, S. (1981) *Chem. Phys. Lett.* 84, 263.

Mihalcz I (2001) "Fundamental characteristics and design method for NiTi shape memory alloy" *Periodica Polytechnic Series on Mechanical Engineering* 45 (1), 75-86

Milanova M & Buker U (2000) "Object recognition in image sequences with cellular neural networks" *Neurocomput* **31**, 125-141

Millan J (1995) "Reinforcement leaning of goal-directed obstacle-avoiding reaction strategies in an autonomous mobile robot" *Robot & Auton Syst* **15**, 275-299

Miller A (1981) "Trends and debates in cognitive psychology" *Cognition* 10, 215-225

Miller G (1956) "Magical number 7 ±2: some limits in our capacity for processing information" *Psychol Rev* 83 (2), 81-97

Miller, G. S. P., (2002) "Snake Robots for Search and Rescue" In *Neurotechnology for Biomimetic Robots* (ed. Ayers J, Davis J, Rudolph A), Bradford Books, MIT Press, Cambridge, MA, USA, 271-284

Minsky M (1961) "Steps towards artificial intelligence" *Proc IRE* (Jun), 8-30

Minsky M (1968) "Matter, mind and models" in *Semantic Informatioon Processing* (ed. Minsky M), MIT Press, USA

Minsky M (1975) "A framework for representing knowledge" in *Psychology of Computer Vision* (ed. Winston P), McGraw-Hill, NY

Minsky M (1980) "K-lines: a theory of memory" *Cog Sci* 4, 117-133

Minsky M (1991) "Society of mind" *Artif Intell* 48, 371-396

Minteer, S. D., Akers, N. L., Thomas, T. J. & Moore, C. M. (2003) "Development and characterization of microbioanodes for alcohol/oxygen biofuel cells" *225th National Meeting of the American Chemical Society*, March 23-27 2003, New Orleans, USA

Minton S et al (1989) "Explanation-based learning: a problem-solving perspective" *Artif Intell* 40, 63-118

Mischel T (1963) "Personal constructs, rules and the logic of clinical activity" *Psychol Rev,* 180-192

Mishkin M, Ungerleider L & Marko K (1983) "Object vision and spatial vision: two cortical pathways" *Trends in Neurosci* (Oct), 411-412

Mitchell T (1983) "Generalisation as search" *Artif Intell* 18 (2), 203-276

Mitchell T (1986) "Explanation-based generalisation: a unifying view" *Mach Learning* 1, 47-80

Mjolsness E & Tavormina A (2000) "Synergy of biology, intelligent systems and space exploration" *IEEE Intell Syst* (Mar/Apr), 20-25

Mobus G (2000) "Adopting robot behaviour to a non-stationary environment: a deeper biologically inspired model of neural processing" *Proc SPIE Conf Sensor Fusion & Decentralisation Control in Robotic Syst III*, 4196, 98-111 (ed. McKee G & Schenker P)

Moed M & Saridis G (1990) "Boltzmann machines for the organisation of intelligent machines" *IEEE Trans Syst Man & Cyber* 20 (5), 1094-1102

Mohl B (2000) "Bionic robot arm with compliant actuators" *Proc SPIE Conf Sensor Fusion & Decentralised Control in Robotic Systems III* (ed. McKee G & Schenker P) **4196**, 82-85

Mohri K, Uchiyama T & Panina L (1997) “Recent advances of micro-magnetic sensors and sensing applications” *Sensors & Actuators A* 59, 1-8

Moller R, Maris M & Lambrinos D (1999) “Neural model of landmark navigation in insects” *Neurocomputing* **26/27**, 801-808

Mondada F & Floreano D (1996) “Evolution and mobile autonomous robotics” *Lecture Notes in Computer Science – Towards Evolvable Hardware* **1062,** 221-249, Springer-Verlag

Montemagno C (2002) “Directed application of nanobiotechnology for the development of autonomous biobots” *Final Report Phase 1 Advanced Aeronautical/Space Concept Study*, NIAC CP 00-02

Moore C, Nelson S & Sur M (1999) “Dynamics of neuronal processing in rat somatosensory cortex” *Trends in Neurosci* **22** (11), 513-520

Moravec H (1984) "Locomotion, vision and intelligence" *Robot Res* 21, 215-224

Morefat M & Wu L (1996) “Purposeful gazing and vergence control for active vision” *Robot & CIM* **12** (2), 135-155

Mori K, Nagao H & Sasaki Y (1998) “Computation of molecular information in mammalian olfactory systems” *Comput Neural Syst* **9**, R79-R102

Moriarty D, Schultz A & Grefenstette J (1999) “Evolutionary algorithms for reinforcement learning” *J Artif Intell* **11**, 199-229

Morrison, T. and Aickelin, U. (2002) “An artificial immune system as a recommender for web sites” *Proceedings of ICARIS 2002*, 161-169, Editors: J. Timmis and P.J. Bentley

Muller P (1970) “Definition of intelligence” *Nature* 228, 1008

Munakata T & Jani Y (1994) “Fuzzy systems: an overview” *Comm Assoc Comp Mach* 37 (3), 69-84

Nabili R, Mammano F & Ashmore J (1998) “How well do we understand the cochlea” *Trends in Neural Sci* **21** (4), 159-167

Najum K & Oppenheim G (1991) "Learning systems: theory and application” *IEE Proc* E138 (4), 183-192

Nakashima, M., Ono, K. (2002) “Development of a Two-Joint Dolphin Robot” In *Neurotechnology for Biomimetic Robots* (ed. Ayers J, Davis J, Rudolph A), Bradford Books, MIT Press, Cambridge, MA, USA, 309-324

Narendra K & Thatchachar M (1974) “Learning automata – a survey” *IEEE Trans Systems Man & Cybernetics* 4 (4), 232-333

NASA (1988) "Space Shuttle Reference Manual" http://www.spaceflight.nasa.gov/shuttle/reference/shutref/orbiter/eps/pwrplants.html (accessed 22/07/2004)

Natale L, Metta G & Sandini G (2002) “Development of auditory-evoked reflexes: visuo-acoustic cues integration in a binocular head” *Robot & Auton Syst* **39**, 87-106

Nau D (1983) “Expert computer systems” *IEEE Computer* (Feb), 63-73

Nauta W & Freitag M (1979) “Organisation of the human brain” *Sci Amer* 241 (3), 78-105

Navon D & Gopher D (1979) “On the economy of the human processing system” *Psychol Rev* **86** (3), 214-255

Nehmzow U & Walker K (2003) "Is the behaviour of a mobile robot chaotic?" *Proc AISB Convention*, Aberystwyth, UK

Nehmzow U (1995) "Animal and robot navigation" *Robot & Auton Syst* **15**, 71-81

Nelson, A.L., Grant, E. and Henderson, T.C. (2002). Competitive relative performance evaluation of neural controllers for competitive game playing with teams of real mobile robots. In *Measuring the Performance and Intelligence of Systems, Proceedings of the 2002 PerMIS Workshop, NIST Special Publication 990*, pp43-50.

Nelson, A.L., Grant, E. and Henderson, T.C. (2004). Evolution of neural controllers for competitive game playing with teams of mobile robots. *Robotics and Autonomous Systems*. **46**(3), pp. 135-150.

Nevins A (1974) "Human-oriented logic for automatic theorem-proving" *J Assoc Comput Machin* 21 (4), 606-621

Newell A & Simon H (1956) "Logic theory machine" *IRE Trans Information Theory* 2 (3), 61

Newell A & Simon H (1976) "Computer science as empirical enquiry: symbols and search" *Comm Assoc Comput Machin* 19, 113-126

Newell A (1982) "Knowledge level" *Artif Intell* 18, 87-127

Newell A, Shaw J & Simon H (1958) "Elements of a theory of human problem solving" *Psychol Rev* 65, 151-166

Nicoud, J.D. and Zufferey, J.C. (2002) Toward Indoor Flying Robots. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, pp. 787-792

Nie J and Linkens D (1993) ``Learning Control Using Fuzzified Self-Organizing Radial Basis Function Network," IEEE Trans. Fuzzy Systems, Vol. 1, pp. 280-287

Nikolaev, N., Iba, H. and Slavov, V. (1999) "Inductive Genetic Programming with Immune Network Dynamics" *Advances In Genetic Programming* **3**, Chapter 15, 355-376, Editors: L. Spector, W. B. Langdon, U.-M. O'Reilly and P. J. Angeline, MIT Press, Cambridge, Massachusetts, USA

Nisbet & Wilson (1977) "Telling more than we can know: verbal reports of mental processes" *Psychol Rev* **84**, 231-259

Nolfi S & Floreano D (2001) *Evolutionary Robotics: the Biology, Intelligence & Technology of Self-Organising Machines* Bradford books, MIT Press, USA

Nolfi S (2002) "Power and limits of reactive agents" *Neurocomputing* **42**, 119-145

Nolfi S, Floreano D, Miglio O & Mondana F (1994) "How to evolve autonomous robots: different approaches in evolutionary robotics" *Proc of 4th Workshop on Artificial Life* (ed. Brooks R & Maes P), Boston, MA, MIT Press

Nolfi, S. and Floreano, D. (1998). Co-evolving predator and prey robots: Do 'arm races' arise in artificial evolution? *Artificial Life*, **4**(4):311–335.

Nolfi, S. and Floreano, D. and Miglino, O. and Mondada, F. (1994b) How to Evolve Autonomous Robots: Different Approaches to Evolutionary Robotics, In. Proc. Alife IV, R. Brooks and P. Maes (eds), MIT Press.

Nolfi, S., O. Miglino, and D. Parisi, 1994. Phenotypic Plasticity in Evolving Neural Networks, in *Proceedings of the Intl. Conf. From Perception to Action*, (D. P. Gaussier and J-D. Nicoud eds.), Los Alamitos, CA: IEEE Press, pp. 146-157.

Nordin P, Banzhaf W & Brameier M (1998) "Evolution of a world model for a miniature rbot using genetic programming" *Robot & Auton Syst* **25**, 105-116

Norman D (1980) "Twelve issues for cognitive science" *Cog Science* 4, 1-33

Norman D (1991) "Approaches to the study of intelligence" *Artif Intell* 47, 327-346

North D (1968) "Tutorial introduction to decision theory" *IEEE Trans Systems Science & Cybernetics* 4 (3), 200-210

O'Sullivan C & Guilbault G (1999) "Commercial quartz crystal microbalances – theory and applications" *Biosensors & Bioelectronics* **14**, 661-670

Ogbonna J & Tanaka H (2000) "Production of pure photosynthetic cell mass for environmental biosensors" *Mater Sci & Eng* **C12**, 9-15

Ogbonna J & Tanaka H (2000) "Production of pure photosynthetic cell biomass for environmental biosensors" *Materials Sci & Eng* 12, 9-15

Okamoto H & Fukai T (2000) "A model for the neural representation of temporal duration" *Biosystems* **55**, 59-64

Oliphant M (1996) "Dilemma of saussarean communication" *BioSyst* **37** (1-2), 31-38

Oliveira A & Jones A (1998) "Synchronisation of chaotic maps by feedback control and application to secure communications using chaotic neural networks" to be published

Oliviera A & Jones A (1997) "Synchronisation of chaotic trajectories using parameter control" *Proc 1st Int Conf on Control of Oscillations & Chaos* (COC 97), 1, 46-49 (ed. Chernousko F & Fradkov A), St Petersburg, Russia

Oliviera E et al (1991) "Multiagent environments in robotics" *Robotica* 9, 431-440

Olson D (1970) "Language & thought: aspects of a cognitive theory of semantics" *Psych Rev*, 257-273

Ootsuki, J. T. and Sekiguchi, T. (1999) "Application of the Immune System Network Concept to Sequential Control" *Proceedings of the IEEE Conference on Systems, Man and Cybernetics 1999* **3**, 869-874

Oprea M & Forrest S (1999) "How the immune system generates diversity: pathogen space coverage with random and evolved antibody libraries" Proc Evolutionary & Genetic Computation Conf

Ortega C & Tyrell A (1998) "Evolvable hardware for fault-tolerant applications" *IEE Colloqium on Evolvable Hardware Syst*, London, Digest No 98/233

Ortega C & Tyrell A (1999) "Biologically inspired fault tolerant architectures for real-time control applications" *Control Eng Practise*, 673-678

Ortega C & Tyrell A (1999) "Reliability analysis in self-repairing embryonic systems" *Proc 1st NASA/DOD Workshop on Evolvable Hardware*, Pasadena, California, 120-128

Ortega C & Tyrell A (1999) "Reliability analysis of self-repairing bio-inspired cellular hardware" *IEE Colloqium on Evolutionary Hardware Systems*, London, Digest No 99/033

Ortega C & Tyrell A (1999) "Self-repairing multicellular hardware: a rliability analysis" *Proc 5th European Conf*, Lausanne, Switzerland, Lecture Notes in Artificial Intelligence 1674, 442-446

Ortega C & Tyrell A (2000) "Hardware implementation of an embryonic architecture using virtex FPGAs" Proc 3rd Int Conf Evolvable Systems: From Biology to Hardware, 155-164

Ortega-Sanchez C & Tyrell A (1997) "Fault tolerant systems: the way biology does it!" *Proc Euromicro 97*, Budapest, 146-151

Osada Y & Gong J (1999) "Intelligent gels – their dynamism and functions" *Proc SPIE Conf on Electroactive Polymer Actuators & Devices*, Mar 1999, Newport Beach, CA, USA, **3669**, 12-18

Osborne, B.A. & Geider, R.J. (1987) "The minimum photon requirements for photosynthesis" *New Phytologist* **106**, 631-644.

Østergaard, E. H., Matari´c, M. J., and Sukhatme, G. S. (2002). Multi-robot task allocation in the light of uncertainty. In *Proc. IEEE Int. Conf. Robotics and Automation*, Washington, D.C.

Ota J, Arai T, Yoshida E, Kurabayashi D & Saki J (1996) "Motion skills in multiple mobile robot system" *Robot & Auton Syst* **19**, 57-65

Ota T, Takahashi M et al (1995) "Biomimetic process for producing SiC wood." *Journal of the American Ceramic Society*78, 3409-3411

Otani M & Jones A (1997) "Guiding chaotic orbits" *Research Report*, University of Wales, Cardiff

Otero T, Cantero I & Villaneuva S (1999) "EAP as multifunctional and biomimetic materials" *Proc SPIE Conf on Electroactive Polymer Actuators & Devices*, Mar 1999, Newport Beach, CA, USA, **3669**, 26-33

Ott E, Grebogi C and Yorke J (1990) "Controlling chaos" *Physical Review Letters* 64 (11), 1196-1199

Paajen M, Lekkala J & Kirjavainen K (2000) "Electromechanical film (EMFi) – a new multipurpose electret material" *Sensors & Actuators* **84**, 95-102

Paap K, Dehlwisch M & Klassen B (1998) "GMD-Snake: a semi-autonomous snake-like robot" *Proc 3rd Int symp on distributed Autonomous Robotic Systems (DARS 96)*, Riken, Saitama, Japan

Palmore, G. T. R. & Kim, H.-H. (1999) "Electro-enzymatic reduction of dioxygen to water in the cathode compartment of a biofuel cell" *Journal of Electroanalytical Chemistry* **464**, 110–117

Palsula S (1993) "Molecular composites: third generation polymers for aerospace applications" *ESA J* **17**, 133-145

Panerai F & Sandini G (1998) "Oculo-motor stabilisation reflex: integration of inertial and visual information" *Neural Networks* **11**, 1191-1204

Panerai F, Metta G & Sandini G (2000) "Adaptive image stabilisation: a need for vision-based active robotic agents" SAB 2000, Paris, France

Panerai F, Metta G & Sandini G (2000) "Learning VOR-like stabilisation reflexes in robots" *Proc 8th European Symp Artificial Neural Networks*, Bruges, Belgium

Panerai F, Metta G, Sandini G (2000) "Visuo-inertial stabilisation in space-variant binocular systems" *Robot & Auton Syst* **30**, 195-214

Pang D et al (1987) "Reasoning with uncertain information" *Proc IEE* 134D(4),231-237

Panzeri S, Schultz S Treves A & Rolls E (1999) "Correlations and the encoding of information in the nervous system" *Proc Roy Soc London* **B266**, 1001-1012

Park D, Kandel A and Langholz G, ``Genetic-Based New Fuzzy reasoning models with Application to Fuzzy Control," IEEE Trans. on System, Man, and Cybernetics, Vol. 24, pp. 39-47, 1994

Park R (1936) "Human nature and collective behaviour" *Amer J Sociol*, 773-741

Park, H. S., Kim, B. H., Kim, H. S., Kim, H. J., Kim, G. T., Kim, M., Chang, I. S., Park, Y. K. & Chang, H. I. (2001) "A novel electrochemically active and Fe(III)-reducing bacterium phylogenetically related to *Clostridium butyricum* isolated from a microbial fuel cell" *Anaerobe* **7**, 297-306

Park, J and Sandberg I, "Universal approximation using radial basis function networks", *Neural Comput.,* 1991, **3,** pp. 245-257

Parker, L. E. (1993). Designing control laws for cooperative agent teams. In *Proc. 1993 IEEE Int'l Conf. on Robotics and Automation*.

Parker, L. E. (1996). On the design of behaviour-based multi-robot systems. *Advanced Robotics*, **10**(6):547–578.

Parker, L. E. (1998). ALLIANCE: An architecture for fault tolerant multi-robot cooperation. *IEEE Transactions on Robotics and Automation*, **14**(2):220–240.

Parker, L. E. (2002). Distributed algorithms for multi-robot observation of multiple moving targets. *Autonomous Robots*, **12**(3):231-255

Partridge D (1981) "Computational theorising as the tool for resolving wicked problems" *IEEE Trans Systems Man & Cybernetics* 11 (4), 318-322

Parvaga C, Torsaianko T & Tolhurst D (2002) "Spatiochromatic properties of natural images and human vision" *Curr Biol* **12**, 483-487

Paseman F & Stollenwork N (1998) "Attractor switching by neural control of chaotic neurodynamics" *Comput Neural Syst* **9**, 549-561

Passino K et al (1989) "Neural computing for mnemonic to symbol conversion in control systems" *IEEE Control Systems Mag* (Apr), 44-51

Patella L & Pinz A (2000) "Active object recognition by view integration and reinforcement learning" *Robot & Auton Syst* **31**, 71-86

Pattee H (1968) "Physical basis of coding and reliability" in *Towards A Theoretical Biology* 3, 63-93 (ed. Waddington C)

Paturi F (1974) *Nature, Mother of Invention - the engineering of plant life*. London, Thames & Hudson.

Payton D et al (1990) "Plan guided reaction" *IEEE Trans Systems Man & Cybernetics* 20 (6), 1370-1382

Pearce T (1997) "Computational parallels between the biological olfactory pathway and its analogue 'the electronic nose': part I. Biological olfaction" *Biosystems* 41, 43-67

Pearce T (1997) "Computational parallels between the biological olfactory pathway and its analogue 'the electronic nose': part II. Sensor-based machine olfaction" *Biosystems* 41, 69-90

Pederson A, Rubert J et al (2003) "Thermal assembly of a biomimetic mineral/collagen composite." *Biomaterials* 24,4881-4890

Penfield W & Perot P (1963) "Brain's record of auditory and visual experience" *Brain* **86**, 595-696

Pennycuick C (1996) "Stresss and strain in the flight muscles as constraints on the eevolution of flying animals" *J Biomechanics* **29** (5), 577-581

Percus, J. K., Percus, O. E. and Perelson, A. S. (1993) "Predicting the size of the T-cell receptor and antibody combining region from consideration of efficient self-nonself discrimination" *Proceedings of the National Academy of Sciences of the United States of America* **90**(5), 1691–1695.

Pereira, G. A. S., Das, A. K., Kumar, V., and Campos, M. F. M. (2003a). Formation control with configuration space constraints. In *Proc. IEEE/RJS Int. Conf. on Intelligent Robots and Systems*, Las Vegas, U.S.A. IEEE Press.

Pereira, G. A. S., Kumar, V., and Campos, M. F. M. (2003b). Localization and tracking in robot networks. In *Proc. 11th Int. Conf. on Advanced Robotics*, pages 465–470, Coimbra, Portugal.

Pereira, G. A. S., Pimentel, B., Chaimowicz, L., and Campos, M. F. M. (2002b). Coordination of multiple mobile robots in an object carrying task using implicit communication. In *Proc. 2002 IEEE Int. Conf. on Robotics and Automation*, pages 281–286, Washington D.C., U.S.A. IEEE Press.

Pereira, G., Kumar, V., and Campos, M. F. M. (2002a). Decentralized algorithms for multirobot manipulation via caging. In *Proc. 5th Int. Workshop on Algorithmic Foundations of Robotics*, pages 242–258, Nice, France.

Perelson A & Oster G (1979) “Theoretical studies of clonal selection: minimal antibody reperoire size and reliability of self/non-self discrimination” *J Theor Biol* **81**, 645-670

Perelson A, Hightower R & Forrest S (1996) “Evolution and somatic learning in V-region genes” Research in Immunology 147, 202-208

Perez-Uribe, A., Floreano, D. and Keller, L. (2003) Effects of group composition and level of selection in the evolution of cooperation in artificial ants. In *Proceedings of the 7th European Conference on Artificial Life ECAL'2003*, Springer, LNAI 2801, pp. 128-137.

Perlis D (1985) "Languages with self-reference I: foundations" *Artif Intell* 21, 301-322

Peter M, Andersen S et al (1992) "Catecholamine-protein conjugates - isolation of 4-phenylphenoxazin-2-ones from oxidative coupling of n-acetyldopamine with aliphatic amino acids." *Tetrahedron* **48**, 8927-8934

Petersen K (1982) “Silicon as a mechanical material” *Proc IEEE* **70**(5), 420-457

Petersen K (1996) “From microsensors to microinstruments” *Sensors & Actuators* A56, 143-149

Pfeifer R & Scheier C (1997) “Sensory-motor coordination: the metaphor and beyond” *Robot & Auton Syst* **20**, 157-178

Pfeifer R & Scheier C (2001) *Understanding Intelligence.* Bradford books, MIT Press, USA

Pfeifer R (1995) “Cognition – perspectives from autonomous agents” *Robot & Autin Syst* **15**, 47-70

Pfeifer R (1996) “Building ‘Fungus Eaters’: Design principles of autonomous agents” *Proc 4th Int Conf on Simulation of Adaptive Behaviour*, Cape Cod, MA, USA

Pfeifer R (1997) “Symbols, patterns and begaviour: a new understanding of intelligence” *preprint*

Pfeiffer F , Eltze J & Weidemaron H (1995) “Six legged walking considering biological principles” *Robot & Auton Syst* **14**, 223-232

Philippides, A. and P. Husbands and T. Smith and M. O’Shea. Flexible couplings: diffusing neuromodulators and adaptive robotics. *Artificial Life*, 2004 (in press)

Poggio T et al (1985) “Computer vision and regularisation theory” *Nature* 317, 314-319

Pojdestvenski I, Cottain M, Park Y, Oguist G (1999) “Robustness and timescale hierarchy in biological systems” *Biosyst* **50**, 71-82

Polanyi M (1968) “Life’s irreducible structure” *Science* 160, 1308-1312

Pollatschek M (1977) "Hierarchical systems and fuzzy set theory" *Kybernetes* 6, 147-151

Post S & Sage S (1990) “Overview of automated reasoning” *IEEE Trans Systems Man & Cybernetics* 20 (1), 202-224

Potter, M. A. & De Jong K. A. (1998), “The Coevolution of Antibodies for Concept Learning”, *Proceedings of the 5th International Conference on Parallel Problem Solving from Nature*, 530-539, Editors: A. E. Eiben, T. Bck, M. Schoenauer, and H. Schwefel, Springer-Verlag, Amsterdam, 1998

Prased B & Reed R (1999) “Chemosensation: molecular mechanisms in worms and mammals” *Trends in Genetics* **15** (4), 150-153

Prescott T, Redgrave P & Gurney K (1999) “Layered control architectures in robots and vertebrates” *Adaptive Behav* **7**, 99-127

Primio F, Muller B & Lengeler J (2001) “Minimal cognition in unicellular organisms” *reprint*

Procyk T & Mamdani E (1979) "Linguistic self organising process controller" *Automatica* 15, 15-30

Pschorn, R., Rukle, W. & Wild, A. (1988) “Structure and function of NADP$^+$ oxidoreductase” *Photosynthesis Research* **17**, 217-229

Pylyshyn Z (1973) “What the mind’s eye tells the mind’s brain: a critique of mental imagery” *Psychol Bull* 80 (1), 1-23

Pylyshyn Z (1999) “Is vision continuous with cognition? The case for cognitive impenetrability of visual perception” *Behav & Brain Sci* 22(3), 341-365

Pylyshyn Z (2000) “Situating vision in the world” *Trends in Cog Sci* **4** (5), 197-207

Pylyshyn Z (2001) “Visual indexes, preconceptual objects and situated vision” *Cog* **80**, 127-158

Pyragas K (1992) "Continuous control of chaos by self-controlling feedback" *Physics Letters A* 170, 421-428

Qian C, Luquan R et al (1999) "Using characteristics of burrowing animals to reduce soil-tool adhesion." *Transactions of the ASAE* **42**, 1549-1556

Quay M, Leroque P & Gaussier P (2002) “Learning and motivational couplings promote smartwer behaviours of an animat in an unknown world” *Robot & Auton Syst* **38**, 149-156

Quinn, M. (2000). Evolving cooperative homogeneous multi-robot teams. In *Proceedings of the IEEE / RSJ International Conference on Intelligent Robots and Systems (IROS 2000)*, pages 1798-1803. Takamatsu, Japan. IEEE Press.

Quinn, M. (2001). A comparison of approaches to the evolution of homogeneous multi-robot teams. In *Proceedings of the Congress on Evolutionary Computation (CEC01)* pages 128-135. Seoul, S. Korea. IEEE Press.

Quinn, M. (2004). *The Evolutionary Design of Controllers for Minimally-Equipped Homogeneous Multi-Robot Systems.* PhD Thesis, University of Sussex, Brighton, U.K.

Quinn, M., Smith, L., Mayley, G. and Husbands, P. (2002). Evolving formation movement for a homogeneous multi-robot system: Teamwork and role-allocation with real robots. *Cognitive Science Research Paper 515*. School of Cognitive and Computing Sciences, University of Sussex, Brighton, BN1 9QG. ISSN 1350-3162.

Quinn, M., Smith, L., Mayley, G. and Husbands, P. (2003). Evolving controllers for a homogeneous system of physical robots: Structured cooperation with minimal sensors. *Philosophical Transactions of* *the Royal Society of London, Series A: Mathematical, Physical and Engineering Sciences* **361**, pages 2321-2344. October 2003.

Quinn, R. D., Nelson, G. M., Ritzmann, R. E. (2002) "Towards the Development of Agile and Mission-Capable Legged Robots" In *Neurotechnology for Biomimetic Robots* (ed. Ayers J, Davis J, Rudolph A), Bradford Books, MIT Press, Cambridge, MA, USA, 401-418

Rabaey, K., Lissens, G., Siciliano, S. D. & Verstraete, W. (2003) "A microbial fuel cell capable of converting glucose to electricity at high rate and efficiency" *Biotechnology Letters* **25**(18), 1531-1535

Raibert M & Craig J (1981) "Hybrid position/force control of manipulators" *Trans ASME J Energy Research Technology* **102** (Jun), 126-133

Ranang, M. T. (2002) "An artificial immune system approach to preserving security in computer networks" *Graduate thesis*, Norwegian University of Science and Technology

Ratnieks, F. L.W. and Anderson, C. (2000). Task partitioning in insect societies. *Insectes Sociaux*, (46):95–108.

Raulin F et al (1999) "Chromatographic instrumentation in space: past, present and future developments for exobiological studies" *Adv Space Res* 23 (2), 361-366

Rebers J & Willis J (2001) "A conserved domain in arthropod cuticular proteins binds chitin." *Insect Biochemistry and Molecular Biology*31, 1083-1093

Rechenberg, I. (1973). *Evolutionsstrategie*, Friedrich Frommann Verlag, Stuttgart.

Reeder T & Cullen D (1976) "Surface-acoustic wave pressure and temperature sensors" *Proc IEEE* 64 (5), 754-755

Reeke G (1997) "Darwin II and Darwin IV" in *Handbook of Evolutionary Computation*, Oxford University/Institute of Physics Press, G8.4:1-G8.4:11

Reeke G, Finkel L & Edelman G (1990) "Selective recognition automata" in *An Introduction to Neural & Electronic Networks*, Academic Press Inc

Regan D & Gray R (2000) "Visually guided collision avoidance and collision achievement" *Trends in Cog Sci* **4** (3), 99-107

Regnier F et al (1999) "Chromatography and electrophoresis on chips: critical elements of future integrated, microfluidic analytical systems for life science" *Trends in Biotech* 17 (Mar), 101-106

Reid A & Staddon J (1997) "A reader for the cognitive map" *Info Sci* **100**, 217-228

Rein K, Zockler M, Mader M, Guibel C & Heisenberg H (2002) "Drosophila standard brains" *Curr Biol* **12**, 227-231

Reiter R (1980) "Logic for default reasoning" *Artif Intell* 13, 81-132

Rescorla R & Solomon R (1967) "Two process learning theory: relationships between Pavlovian conditioning and instrumental learning" *Psychol Rev* 74 (3), 151-181

Reynolds C (1987) "Flocks, herds and schools: a distributed behavioural model" *Comput Graphics* **21** (4), 25-34

Richard K Belew (1993). Interposing an ontogenic model between genetic algorithms and neural networks. In Advances in neural Information Processing Systems (NIPS) SJ Hanson and JD Cowan and CL Giles, Morgan Kauffman.

Richards C & Papaikopoloulos N (1997) “Detection and tracking for robotic visual servoing systems” *Robotics & Comput Integ Manuf* **13** (2), 101-120

Richards D (1985) “Biological strategies for communication” *IEEE Comm Syst Mag* 23 (6), 10-18

Rittersma Z (2002) “Recent achievements in miniaturised humidity sensors – a review of transduction techniques” *Sensors & Actuators* A96, 196-210

Robinson D, Pratt J et al (1999) "Series elastic actuator development for a biomimetic walking robot". *IEEE/ASME International Conference on Advanced Intelligent Mechantronics*

Robinson R (1976) “Adaptive gain control of VOR by cerebellum” *J Neurophys* **39**, 954-969

Rocha L (1997) “Evidence sets: contextual categories” *Proc Meeting on Control Mechanisms for Complex Systems*, Physical Science Lab, New Mexico State University, Las Cruces, New Mexico (ed. Coombs M), NMSU Press, 339-357

Roitt, I. M., and Delves, P. J. (2001) *Roitt’s Essential Immunology*, 10th ed., 68-69, Editors: Nick Morgan, Meg Barton, and Fiona Goodgame, Blackwell Science Commerce Place, 350 Main Street, Malden, MA 02148-5018, USA

Rosa M & Krubitzer L (1999) “Evolution of visual cortex: where is V2?” *Trends in Neurosci* **22** (6), 242-248

Rosenblatt F (1958) “The perceptron: a probabilistic model for information storage and organisation in the brain” *Psychol Rev* 65, 386-407

Rossi Dd & Chiarelli P (1994) "Biomimetic macromolecular actuators." *ACS Symposium Series*548, 517-530

Rossi Dd (1989). "Biomimetic approaches to the design of materials for artificial tactile perception". *Proc Int Workshop Intell Mats*, 251-258, Soc Non-Trad Tech, Tokyo

Rothstein J (1988) “Bus automata, brains and mental models” *IEEE Trans Systems Man & Cybernetics* 18 (4), 522-530

Rozenfeld A & Weska J (1976) “Picture recognition and scene analysis” *IEEE Computer* (May), 28-33

Ruck D et al (1992) “Comparative analysis of backpropagation and the extended Kalman filter for training multi-layer perceptrons” *IEEE Trans Pattern Analysis & Machine Intelligence* 14 (6), 686-691

Rus, D., Donald, B., and Jennings, J. (1995). Moving furniture with teams of autonomous robots. In *Proc. IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*, pages 235–242.

Russel S (1991) “Prior knowledge and autonomous learning” *Robotics & Autonomous Syst* 8, 145-159

Rutherford, A.W. (1989) "Photosystem II, the water splitting enzyme" *Trends in Biochemical Sciences* **14**, 227-242

Saari H et al (2000) “Miniaturised gas sensor using a micromachined Fabry-Perot interferometer” *Preparing for the Future* **10** (3), 4-5

Sage A (1981) “Behavioural and organisational considerations in the design of information systems and processes for planning and decision support” *IEEE Trans Systems Man & Cybernetics* **11** (9), 640-678

Saito Y, Kumagai H et al (2002) "Thermally reversible hydration of beta-chitin: ultramicrostructure and microthermomechanics of biological IR detectors: materials properties from a biomimetic perspective." *Biomacromolecules*3, 407-410

Salisbury J (1988) “Issues in human/computer control of dextrous remote hands” *IEEE Trans Aerospace & Electronic Systems* **24** (5), 591-596

Salomon R (1998) “Achieving robust behaviour by using proprioceptive activity patterns” *Biosyst* **47**, 193-206

Samuel A (1959) “Some studies in machine learning using the game of chequers” *IBM J Res & Dev* **3**, 210-229

Samuel A (1976) “Some studies in machine learning using the game of checkers II – recent progress” *IBM J Res & Devel* 11, 601-617

Santoli S (1997) “Nano to micro integrated single-electron biomacromolecular electronics for miniaturised robotic ‘untethered flying observer’” *Acta Astron* 41 (4-10), 279-287

Santoli S (2002) “Lab-on-a-chip design for miniature autonomous bio/chemoprospecting planetary rovers” *Jour Brit Interplan Soc* 55 115-130

Santos-Victor J & Sandini G (1997) “Embedded visual behaviours for navigation” *Robot & Auton Syst* **19**, 299-313

Saranli, U., Buehler, M., Koditschek, D. E., (2002) “RHex – A simple and Highly Mobile Hexapod Robot.” Department of Electrical Engineering and Comptuer Science. Unpublished.

Saridis G (1979) “Towards the realisation of intelligent control” *Proc IEEE* **67** (8), 1115-1132

Saridis G (1983) “Intelligent robotic control” *IEEE Trans Automatic Control* 28 (5), 547-557

Saridis G (1988) “Analytical formulation of the principle of increasing precision with decreasing intelligence for intelligent machines” *IFAC Robot Control*, Karlsruhe, 529-557

Sarikaya M, Fong H et al. (2002). "Biomimetics: nanomechanical design of materials through biology". *15th ASCE Engineering Mechanics Conference*, Columbia University, New York, NY.

Sasaki S & Karube I (1999) “Development of micro-fabricated biocatalytic fuel cells” *Trends in Biotechnol* **17**, 50-52

Sathyanath, S. and Sahin, F. (2001) “An AIS approach to a color image classification problem in a real time industrial application” *Proceedings of the IEEE Conference on Systems, Man and Cybernetics 2001*

Sathyanath, S. and Sahin, F. (2002) “*AISIMAM* – An artificial immune system based intelligent multi agent model and its application to a mine detection problem” *Proceedings of ICARIS 2002*, 22-31, Editors: J. Timmis and P.J. Bentley

Schaal S (1999) “Is imitation learning the route to humanoid robots?” *Trends in Cog Sci* **3** (6), 233-242

Schaal S, Sternad D, Kotosaka S, Osa R, Kawato M (2000) “Reciprocal excitation between biological and robotic research” *Proc SPIE Conf Sensor Fusion & Decentralised Control in Robotic Syst III*, **4196**, 1-11, (ed. McKee G & Schenker P)

Schank R (1972) “Conceptual dependency: a theory of natural language understanding” *Cog Psychol* **3**, 552-631

Scheutz M, Sloman A & Logan B (2001) “Emotional states and realistic agent behaviour” *preprint*

Schinner G, Peterlik H et al (1995) "Structural design and mechanical properties in spines of Spatangoid sea urchins." *Biomimetics* 3, 13-30

Schmajuk N & DiCarlo J (1992) “Stimulus configuration, classical conditioning and hippocampal function” *Psychol Rev* **99** (2), 268-305

Schmajuk N, Scymanski W, Weaver E (1999) “Adaptive communication in animals and robots” *Signal Processing* **74**, 71-87

Schmidt C et al (1978) “Plan recognition problem: an intersection of psychology and AI” *Artif Intell* 11, 45-83

Schmitt O (1969). "Some interesting and useful biomimetic transforms" *Third International Biophysics Congress.*

Schmitz H & Bleckmann H (1997) "Fine structure and physiology of the infrared receptor of beetles of the genus *Melanophila*(Coleoptera: Buprestidae)." *International Journal of Insect Morphology and Embryology*26, 205-215

Schmitz H & Bleckmann H (1998) "The photomechanic infrared receptor for the detection of forest fires in the beetle *Melanophila acuminata*(Coleoptera: Buprestidae)." *J. Comp. Physiol. A-Sensory, Neural and Behavioural Physiol.*182, 647-657

Scholl B (2001) “Objects and attention: the state of the art” *Cog* **80**, 1-46

Schoner G & Kelso J (1988) “Dynamic pattern generation in behavioural and neural systems” *Science* 239, 1513-1519

Searle J (1980) "Minds, brains & programs" *Beh & Brain Sci* 3(3), 417-428

Searle J (1990) "Is the brain's mind a computer program?" *Sci Am* (Jan), 20-25

Sebeok J (1965) "Animal communication" *Sci* 147, 1006-1014

Segre A (1991) “Learning how to plan” *Robotics & Automous Syst* 8, 93-111

Seligman M (1970) “On the generality of the laws of learning” *Psychol Rev* **77** (5) 406-418

Sellinger A, Weiss P et al (1998) "Continuous self-assembly of organic-inorganic nanocomposite coatings that mimic nacre." *Nature*394, 256-260

Sepulchre J & Babloyantz A (1993) "Control of chaos in network of oscillators" *Physical Review E* 48 (2), 945-950

Shahinpoor M & Thompson M (1995) “Venus flytrap as a model for a biomimetic material with built-in sensors and actuators” *Mater Sci Engin* **C2**, 229-233

Shahinpoor M (1994) “Continuum electromechanics of ionic polymeric gels as artificial muscles for robotic application” *Smart Mater Struct* **3**, 367-372

Shahinpoor M et al (1998) “Ionic polymer-metal composites (IPMC) as biomimetic sensors, actuators and artificial muscles – a review” *Int J Smart Materials & Structures* **7**, R15-R30

Shallice T (1973) "Dual function of consciousness" *Psych Rev* 79, 383-393

Shannon C (1948) “Mathematical theory of communication” *Bell Syst Tech J* 27, 379-423

Shapiro I & Narendra K (1969) “Use of stochastic automata for parameter self-optimisation with multimodal performance criteria” *IEEE Trans Systems Science & Cybernetics* 5 (4), 352-360

Sharkey N & Heemskerk J (1997) “Neural mind and the robot” *preprint*

Sharkey N & Ziamke T (1998) “Consideration of the biological and psychological foundations of autonomous robotics” *Connection Sci* **10** (3-4), 1-28

Sharkey N (1997) “New wave in robot learning” *Robot & Auton Syst* **22**, 179-185

Shepard R & Judd S (1976) “Perceptual illusion of rotation of 3D objects” *Sci* **191**, 952-954

Shepard R & Metzler J (1971) “Mental rotation of 3 dimensional objects” *Science* **171**, 701-703

Shephard R (1984) “Ecological constraints on internal representation: resonant kinematics of perceiving, imagining, thinking and dreaming” *Psychol Rev* 91 (4), 417-447

Shiffrin R & Schneider W (1977) “Controlled and automatic human infprmation processing II: perceptual learning, automatic attending and a general theory” *Psychol Rev* 84 (2), 127-190

Shin H, Jo S et al (2003) "Biomimetic materials for tissue engineering." *Biomaterials*24, 4353-4364

Shoch J & Hupp J (1982) "Worm programs" *Comm Assoc Comp Mach* 25(3), 172-180

Shu S, Lagoudes D, Hughes D & Wen J (1997) “Modelling of a flexible beam activated by shape memory alloy wires” *Smart Materials & Structures* **6**, 265-277

Silvermann B et al (1989) “Blackboard system generator: alternative distributed proboem solving paradigm” *IEEE Trans Systems Man & Cybernetics* 19 (2), 334-354

Simmons J (1989) “View of the world through the bat’s ear: the formation of acoustic images by echolocation” *Cog* **33**, 155-199

Simões, E. D. V. and Dimond, K. R. (1999). An evolutionary controller for autonomous multirobot systems. In *Proceedngs of the IEEE International Conference on Systems, Man and Cybernetics*, pages 596-601, Tokyo, Japan. IEEE Press.

Simões, E. D. V. and Dimond, K. R. (2001). Embedding a distributed evolutionary system into a population of autonomous mobile robots. In *Proceedings of the IEEE International Conference on Systems, Man and Cybernetics*, pages 1069-1074, Tucson, Arizona U.S.A. IEEE Press.

Simon H (1963) “Discovery, invention and development: human creatuive thinking” *Proc National Academy Sci* 80, 4569-4571

Simon H (1978) “Rationality as a process and as the product of thought” *Amer Econ Rev* 68 (2), 1-16

Simon H (1991) “AI: where it has been and where is it going?” *IEEE Trans Knowledge & Data Eng* 3 (2), 128-136

Singh C & Singh M (1978) “Explanatory analysis of organisational hierarchies from an engineering point of view” *IEEE Trans Syst Man & Cyber* 8 (3), 205-208

Singh M & Hoffman D (1997) “Constructing objects and representing visual objects” *Trends in Cog Sci* **1** (3), 98-102

Sirovich L & Karlsson S (1997) "Turbulent drag reduction by passive mechanisms." *Nature*388, 753-755

Sitti, M. & Fearing, R. S. (2002) "Nanomolding Based Fabrication of Synthetic Gecko Foot-Hairs" *2nd IEEE Conf. on Nanotechnology*, August 26-28 2002, Washington, D.C., USA

Sitti, M. (2003) "High Aspect Ratio Polymer Micro/Nano-Structure Manufacturing using Nanoembossing, Nanomolding and Directed Self-Assembly" *IEEE/ASME Advanced Mechatronics Conference,* July 2003, Kobe, Japan

Skarda C & Freeman W (1987) “How brains make chaos in order to make sense of the world” *Behav & Brain Sci* **10**, 161-195

Skarda C & Freeman W (1987) “How brains make chaos in order to make sense of the world” *Behav & Brain Sci* **10**, 161-195

Skullestal A & Hallingstad O (1998) “Vibration parameters identification in a spacecraft subjected to active vibration damping” *Mechatronics* **8**, 691-705

Slack J (1989) “D309 – Cognitive Architecture” *Open University Course Publication*

Slansky J (1964) “Adaptation, learning, self-repair and feedback” *IEEE Spectrum* (May). 172-174

Slansky J (1966) “Learning systems for automatic control” *IEEE Trans Automatic Control* 11 (11), 6-19

Slavov, V. and Nikolaev, N. I. (1998), “Immune Network Dynamics for Inductive Problem Solving”, *Proceedings of the 5th International Conference on Parallel Problem Solving from Nature*, 712-721, Editors: A. E. Eiben, T. Bck, M. Schoenauer, and H. Schwefel, Springer-Verlag, Amsterdam

Sloman A & Croucher T (1981) “Why robots will have emotions” *Int J Conf Artif Intell*, Vancouver, 197-202

Smallwood J (1998) “Introduction to assumption-based truth maintenance” *IEE Comp & Control Eng J* (Oct), 233-246

Smeekens, S., Weisbeek, P. & Robinson, C. (1990) “Protein transport into and within chloroplasts” *Trends in Biochemical Sciences* **15**, 73-76.

Smith B (1991) "The owl and the electric encyclopedia" *Artif Intell* 47, 251-288

Smith C, Herbert R et al (2000) "The hind wing of the desert locust (*Schistocerca gregaria*Forskål). II. Mechanical properties and functioning of the membrane." *Journal of Experimental Biology*203, 2933-2943

Smith R & Davies R (1981) "Framework for cooperation in distributed problem solving" *IEEE Trans Syst Man & Cyber* 11(1), 61-70

Smith R (1980) “Contract net protocol: high level communication and control in a distributed problem solver” *IEEE Trans Computing* 29 (2), 1104-1113

Smithsonian Institution (2001) "Fuel Cells: Origins of the Technology" http://fuelcells.si.edu/origins/origins.htm (accessed 20/07/2004)

Smolensky P (1988) “On the proper treatment of connectionism” *Behav & Brain Sci* 11,1-74

Snyder L (1982) “Introduction to the configurable highly parallel computer” *IEEE Computer* (Jan), 47-56

Sole R & de la Prode L (1995) “Controlling chaos in discrete neural networks” *Physics Lett A* 199, 65-69

Soong R.K. et al. (2000) "Powering an inorganic nanodevice with a molecular motor". *Science* **290**, 1555-1558.

Spaan, M. T. J. and Groen, F. C. A. (2002). Team coordination among robotic soccer players. In Kaminka, G., Lima, P. U., and Rojas, R., editors, *RoboCup 2002*, pages 356–363. Springer-Verlag.

Spletzer, J., Das, A. K., Fierro, R., Taylor, C. J., Kumar, V., and Ostrowski, J. P. (2001). Cooperative localization and control for multi-robot manipulation. In *Proc. IEEE / RSJ Int. Conf. Intelligent Robots and Systems*.

Sporns O, Gally J, Reeke G & Edelman G (1989) “reentrant signalling among simulated neuronal groups leads to coherency in their oscillatory activity” *Proc Nation Acad Sci* **86**, 7265-7269

Springer A et al (1999) “Wireless identification and sensing using surface acoustic wave devices” *Mechatronics* 9, 745-756

Srinivasan A (1996) “Smart biological systems as models for engineered structures” *Mater Sci & Eng* C4, 19-26

Srinivasan M (1992) “How bees exploit optic flow: behavioural experiments and neural models” *Phil trans Roy Soc* **B337**, 253-259

Srinivasan M, Chahl J, Weber K, Ventatesh S, Nagke M & Zhang S (1999) “Robot navigation inspired by principles of insect vision” *Robot & Auton Syst* **26**, 203-216

Stanfill C & Waltz D (1986) "Towards memory based reasoning" *Comm Assoc Comp Mach* 29 (12), 1213-1228

Stanley S (1933) “An explanation for Cope’s rule” *Evol* **27** (1), 1-22

Steels L (1995) “Intelligence – dynamics and representations” in The Biology & Technology of Intelligent Autonomous Agents (ed. Steels L), Springer-Verlag, Berlin

Steels L (1995) “When are robots intelligent autonomous agents?” *Rob & Auton Syst* 15, 3-9

Steels L (1998) “Origins of syntax in visually grounded robotic agents” *Artif Intell* **103**, 133-156

Stein P (1978) “Motor systems, with specific reference to the control of locomotion” *Ann Rev Neurosci* **1**, 61-81

Stein R (1982) “What muscle variable(s) does the nervous system control in limb movements?” *Behavioural & Brain Sciences* **5**, 535-577

Sternberg R (1985) "Human intelligence: the model is the message" *Sci* 230, 1111-1117

Stevens C (1979) “The neuron” *Sci Amer* 241 (3), 49-59

Stevens J (1985) “Reverse engineering the brain” *Bte* (Apr), 287-299

Stewart W. Wilson (1994). ZCS: A Zeroth Level Classifier System, *Evolutionary Computation*, Volume 2, Issue 1.

Stoica A (2001) “Robot fostering techniques for sensory-motor development of humanoid robots” *Robot & Auton Syst* **37**, 127-143

Stone, P. and Veloso, M. (1999). Task decomposition, dynamic role assignment, and lowbandwidth communication for real-time strategic teamwork. *Artificial Intelligence*, **110**(2):241–273.

Stone, P. and Veloso, M. (2000). Multiagent systems: A survey from a machine learning perspective. *Autonomous Robots*, **8**(3):345–383.

Street G & Clarke T (1981) “Conducting polymers: a review of recent work” *IBM J Res & Devel* **25** (1), 51-57

Stroeve S (1998) “Neuromuscular control model of the arm including feedback and feedforward components” *Acta Psychol* **100**, 117-131

Suga N, Yan J & Zhang Y (1997) “Cortical maps for hearing and egocentric selection for self-organisation” *Trends in Cog Sci* **1** (1), 13-20

Sun C (1994) ``Rule-Based Structure Identification in an Adaptive-Network-Based Fuzzy Inference System," IEEE Trans. Fuzzy Systems, Vol. 2, pp. 64-73

Sun, L., Berglund, H., Davydov, R., Norrby T., Hammarstrom, L., Korall, P., Borje, A., Philouze, C., Berg, K. & Tran, A. (1997) "Binuclear ruthenium-manganese complexes as simple artificial models for photosystem II in green plants" *J Am Chem Soc*. **199**, 6996-7004.

Sutherland J (1986) “Assessing AI’s contribution to decision technology” *IEEE Trans Systems Man & Cybernetics* 16 (1), 3-20

Sutton R & Barto A (1981) “Towards a modern theory of adaptive networks: expectation and prediction” *Psychol Rev* **88**, 135-170

Suzuki, J. and Yamamoto, Y. (2001) "Biologically-inspired autonomous adaptability in a communication endsystem: an approach using an artificial immune network" *IEICE Transactions on Information and Systems* **E84-D**(12), 1782--1789

Suzuki, S. K., Matsunaga, I. & T. (1979) *Proceedings of the Biotechnology and Bioenergetics Symposium*, 501–511

Svinin M, Yamada K & Ueda K (2001) "Emergent synthesis of motion patterns for locomotion robots" *Artif Intell in Engin* **15**, 353-363

Swets J et al (1962) "Decision processes in perception" *Psychol Rev*, 301-317

Szentagothai J (1978) "Neuron networks of the cerebral cortex: a functional interpretation" *Proc Roy Soc* B201, 219-248

Tabacco M & DiGuiseppe T (1996) "Optical chemical sensors for environmental control and system management" *Adv Space Res* 18 (4/5), 125-134

Tabib-Azar M & Garcia-Valenzuela A (1995) "Sensing means and sensor shells: a new method of comparative study of piezoelectric, piezoresistive, electrostatic, magnetic, and optical sensors" *Sensors & Actuators* A48,87-100

Tanahashi M, Kokubo T et al (1996) "Quantitative assessment of apatite formation via a biomimetic method using quartz crystal microbalance." *Journal of Biomedical Materials Research*31, 243-249

Tanahashi M, Kokubo T et al (1996) "Ultrastructural study of an apatite layer formed by a biomimetic process and its bonding to bone." *Biomaterials*17, 47-51

Tanaka K et al (1995) "Micromachined vibrating gyroscope" *Sensors & Actuators* A50, 111-115

Tangwongstan S & Fu K (1979) "Application of learning to robotic planning" *Int J Comp & Info Sci* 8 (4), 303-333

Tank D & Hopfield J (1989) "Neural computation by concentrating information in time" *Proc Nat Acad Sci USA* 84, 1896-1900

Tarakanov, A. O., Goncharova, L. B., Gupalova, T. B., Kvachev, S. V. and Sukhorukov, A. V. (2002) "Immunocomputing for Bioarrays" *Proceedings of ICARIS 2002*, 32-40, Editors: J. Timmis and P.J. Bentley

Tarr M & Cutthoff H (1998) "Image-based object recognition in man, monkey and machine" *Cog* **67**, 1-20

Tarski A (1969) "Truth & proof" *Sci Am* 220 (Jun), 43-77

Tatylor G (2001) "Mechanics and aerodynamics of insect flight control" *Biol Rev* **76**, 449-471

TenHuisen K & Brown P (1992) "Microstructural development and formation kinetics in a mineralizing system: gelatin gypsum." *Biomimetics*1, 131-150

Thagard P & Shelley C (1997) "Abductive reasoning: logic, visual thinking and coherence" in *Logic & Scientific Methods*, ed. Della Chiara M-L et al, Kluwer, Dordrecht, 413-427

Thakoor, S. (2002) "Earthworm-Like Burrowing Robot" *NASA TECH BRIEF* **22**(6), Item #138

Thakoor, S., Kennedy, B., Thakoor, A. P. (1999) "Insectile and Vermiform Exploratory Robots" *NASA TECH BRIEF* 23(11), 1-22. November 1999.

Thompson A & Harvey I (1997) "Artificial evolution for real problems" in Evolutionary Robotics: From Intelligent Robots to Artificial Life (ed. Gomi T), AAI Books

Thompson A & Layzell P (1999) “Analysis of unconventional evolved electronics” *Commun Assoc Comput Machin*

Thompson A (1995) “Evolving electronic robot controllers that exploit hardware resources” *Proc 3rd European Conf on Artificial Life* , Springer-Verlag, 640-656

Thompson A (1995) “Evolving electronic robot controllers that exploit hardware resources” *Advances in Artificial Life: Proc 3rd European Conf on Artif Life*, 640-656

Thompson A (1995) “Evolving fault tolerant systems” *Proc 1st IEE/IEE Int Conf on Genetic Algorithms in Engineering Systems*, IEE Conf Publ No 414, 524-529

Thompson A (1996) “An evolved circuit, intrinsic in silicon, entwined with physics” *Proc 1st Int Conf Evolvable Systems*, Tsukuba, Japan, 390-405

Thompson A (1996) “Evolutionary techniques for fault tolerance” *UKACC Int Conf on Control*, 693-698

Thompson A (1996) “Silicon evolution” *Proc Genetic Programming* (ed. Koza J et al), MIT Press, 444-452

Thompson A (1997) “Artificial evolution in the physical world” in Gomi I, AAAI Books, 101-129

Thompson A (1998) “On the automatic design of robust electronics through artificial evolution” *Proc 2nd Int Conf on Evolvable Systems*, Springer-Verlag

Thompson A, Harvey I & Husbands P (1996) “Unconstrained evolution and hard consequences” in *Towards Evolvable Hardware: the Evolutionary Engineering Approach* (ed. Sanchez & Tomassini), Springer, 136-165

Thompson A, Layzell P & Zebulum R (1999) “Explorations in design space: unconventional electronics design through artificial evolution” *IEEE Trans Evol Comp* 3 (3), 167-196

Thompson B et al (1992) “Introduction to smart materials and structures” *Materials & Design* **13** (1), 3-9

Thomson R (1986) “Neurobiology of learning and memory” *Science* 233, 941-947

Thrun S (1995) “Approach to learning mobile robot navigation” *Robot & Auton Syst* **15**, 301-319

Thrun, S., Burgard, W., Fox, D. (1998). “A Probabilistic Approach to Concurrent Mapping and Localization for Mobile Robots.” Machine Learning **31** 29-53.

Thrun, Sebastian, Wolfram Burgard, Dieter Fox (1998). “A Probabilistic Approach to Concurrent Mapping and Localization for Mobile Robots” *Machine Learning* **31**, 29-53.

Ting et al (2002) “Cyanobacterial photosynthesis in the oceans: the origins and significance of divergent light-harvesting strategies” *Trends in Microbiol* 10 (3), 134-142

Toko K (1998) “Electronic tongue” *Biosensors & Bioelectron* 13, 701-709

Tokuda I, Yanai T & Aihara K (2001) “Reconstruction of chaotic dynamics via a network of stochastic resonance neurons and its application to speech” *Artif Life & Robotics* **5**, 33-39

Tolman E (1948) “Cognitive maps in rats and man” *Psychol Rev* **55**, 189-208

Toma, N., Endo, S. & Yamada, K. (1999), “Immune algorithm with immune network and MHC for adaptive problem solving”, *Proceedings of the IEEE Conference on Systems, Man and Cybernetics 1999* **4**, 271-276.

Tootell R, Dale A, Sereno M & Malach R (1996) “New images from human visual cortex” *Trends in Neurosci* **19** (11), 481-489

Tovec M (1995) “Ultraviolet photoreceptors in the animal kingdom: their distribution and function” *Trends in Ecol* **10** (11), 455-459

Tradepalli P & Ok D (1998) “Model-based average reward reinforcement learning” *Artif Intell* **100**, 177-224

Trebi-Ollennu, Ashitey (webmaster), http://prl.jpl.nasa.gov/projects/lemur1/lemur_index.html, 27-July-2004

Treisman A & Gelade G (1980) “Feature integration theory of attention” *Cog Psychol* 12, 97-136

Tresilan J (1999) “Visually timed action: time-out for tau?” *Trends in Cog Sci* **3** (8), 301-310

Trianni, V. (2003). *Evolution of coordinated motion behaviors in a group of self-assembled robots.* Diplome d’itudes approfondies, Université Libre de Bruxelles, Brussels, Belgium.

Trianni, V., Nolfi, S. and Dorigo, M. (2004). Hole Avoidance: Experiments in Coordinated Motion on Rough Terrain. In F. Groen, N. Amato, A. Bonarini, E. Yoshida, and B. Krose, editors, Proceedings of the 8th Conference on Intelligent Autonomous Systems (IAS8), pages 29-36, Amsterdam, The Netherlands, 2004. IOS Press.

Trullier O, Weiner S, Berthoz A & Meyer J-A (1997) “Biologically based artificial navigation systems: review and prospects” *Progress in Neurobiol* **51**, 483-544

Tsotsos J (1995) “Behaviourist intelligence and the scaling problem” *Artif Intell* **75**, 135-160

Tsui A & Jones A (1997) "Controlling chaotic neural networks" *International Work-Conference on Artificial and Natural Networks 97 (IWANN '97),* June 4

Tsujimura, S., Wadano, A., Kano, K. & Ikeda., T (2001) "Photosynthetic bioelectrochemical cell utilizing cyanobacteria and water-generating oxidase" *Enzyme and Microbial Technology* **29**, 225–231

Turakanov A &Dasgupta D (2000) “A formal model of an artificial immune system” *Biosyst* **55**, 151-158

Turing A (1950) "Computing machinery & intelligence" *Mind* 59, 433-460

Tversky A & Kahneman D (1981) “Framing of decisions and the psychology of choice” *Science* 211, 453-458

Tyrell A (1999) “Computer know thyself: a biological way to look at fault tolerance” *Proc 2nd IEEE Workshop on Dependable Computing Systems,* Milan

Urzelai, J. and Floreano, D. (2000) "Evolutionary Robotics: Coping with Environmental Change" *Proceedings of the Genetic and Evolutionary Computation Conference*, San Mateo, California, USA. Morgan Kaufmann, USA

Urzelai, J. and Floreano, D. (2000) "Evolutionary Robots with Fast Adaptive Behavior in New Environments" A. Thompson et al. (Eds.), *Proceedings of the Third International Conference on Evolvable Systems: From Biology to Hardware*, Berlin. Springer Verlag, USA

Urzelai, J. and Floreano, D. (2000) Evolutionary Robots with Fast Adaptive Behavior in New Environments. In A. Thompson et al. (Eds.), *Proceedings of the Third International Conference on Evolvable Systems: From Biology to Hardware*, Berlin: Springer Verlag

Usui, Y. and Arita, T. (2003). Situated and embodied evolution in collective evolutionary robotics. In *Proceedings of the 8th International Symposium on Artificial Life and Robotics*, pp. 212-215.

Valavanis K & Saridis G (1988) “Information theoretic modelling of intelligent robotic systems” *IEEE Trans Systems Man & Cybernetics* 18 (6), 852-872

Valavanis K & Yuan P (1989) “Hardware and software for intelligent robotic systems” *J Intell & Robot Syst* 1, 343-373

Valiant L (1984) “A theory of the learnable” *Comm Assoc Comput Machin* 27 (1), 1134-1142

Vamos T (1983) "Cooperative systems" *IEEE Control Syst Mag* (Aug), 9-13

Van der Smagt P (2000) “Benchmarking cerebellar control” *Robot & Auton Syst* **32**, 237-251

Van Essen D & Anderson C (1995) “Information processing strategies and pathways in the primate visual system” in *An Introduction to Neural & Electronic Networks*, Academic Press Inc, 45-76

van Gelder T (1995) “What might cognition be if not computation?” *J Philosophy* 91, 345-381

Van Leewen M & Vreeken J (2003) “A philosophy of minimalist biomimetic robotics” Institute for Information & Computing Sciences, Utrecht University

Vaughan, R. T., Støy, K., Sukhatme, G. S., and Matari´c, M. J. (2000). Go ahead, make my day: robot conflict resolution by aggressive competition. In *Proc. 6th Int. Conf. Simulation of Adaptive Behaviour*.

Venthatavaman S (1997) “Simple legged locomotion gait model” *Robot & Auton Syst* **22**, 75-85

Verschure P, et al Krose B, Pfeifer R (1982) “Distributed adaptive control: self-organisation of structured behaviour” *Robotics & Autonom Syst* 9, 181-196

Vertosick, F. T. and Kelly, R. H. (1991) “The immune system as a neural network: a multi-epitope approach” *Journal of Theoretical Biology* **150**(2), 225-37

Vico F et al (2001) “Animal-like adaptive behaviour” *Artif Intell in Engin* 15, 5-12

Villasenov J & Mangione-Smith W (1997) “Configurable computing” *Sci Am* (Aug)

Vincent J & Ablett S (1988) "Hydration and tanning in insect cuticle." *Journal of Insect Physiology*33, 973-979

Vincent J & King M (1996) "The mechanism of drilling by wood wasp ovipositors."*Biomimetics*3, 187-201

Vincent J & Mann D (2002) “Systematic technology transfer from biology to engineering” *Phil Trans Roy Soc A* **360**, 159-173

Vincent J (1990) “Structural Biomaterials” Princeton University Press

Vincent J (1992) “The future: towards intelligent materials and structures” *Metals & Materials* (Jun), 12-15

Vincent J (2000) “Deployable structures in nature: potential for biomimicking” *Proc Inst Mechan Eng* **214C**, 1-10

Vincent J (2002) "Arthropod cuticle: a natural composite shell system." *Composites*33, 1311-1315

Vincent J, Jeronimidis G et al (1992) "Biomimetics of flexible composites: towards the development of new materials." *Biomimetics* 4, 251-263

Viswanathan R & Narendra K (1973) “Stochastic automata models with application to learning systems” *IEEE Trans Systems Man & Cybernetics* 3, 107-111

Viswanathan R & Narendra K (1974) “Games of stochastic automata” *IEEE Trans Systems Man & Cybernetics* (Jan), 431-435

Vogel S (1998) *Cats' Paws and Catapults*. New York and London, WW Norton & Co.

Voinnett O (2001) “RNA silencing as a plant immune system against viruses” *Trends in Genetics* **17** (8), 449-459

Voogd J & Glickstein M (1998) “Anatomy of the cerebellum” *Trends in Neurosci* **21** (9), 370-375

Vose M (1983) “Generalising the notion of schemas in genetic algorithms” *Artif Intell* 50 385-396

Vukusic P & Sambles J (2003) "Photonic structures in biology." *Nature*424, 852-855

Vukusic P (2004) “Natural photonics” *Physics World* (Feb), 35-39

Vukusic P, Sambles J et al (1999) "Quantified interference and diffraction in single morpho butterfly scales." *Proceedings of the Royal Society* B266, 1403-1411

Wada B (1990) “Adaptive structures: an overview” *J Spacecraft & Rockets* **27** (3), 330-336

Wagner G & Altenberg L (1995) “Complex adaptations and the evolution of evolvability” *preprint*

Wagner G (1995) “Homologues, natural kinds, and the evolution of modularity” *preprint*

Wagner G (1996) “Adaptation and the modular design of organisms” *preprint*

Wah B et al (1989) “Computation for symbolic processing” *Proc IEEE* 77 (4), 509-539

Waldron K, Vohnout V, Pery A & McGhee R (1984) “Configuration design of the Adaptive suspension Vehicle” *Int J Robot Res* **3** (2), 37-48

Walker M (1987) “How feasible is automated discovery?” *IEEE Expert* (Spring), 69-82

Wang C, Huang Y, Zan Q, Guo H & Cai S (2000) “Biomimetic structure design – a pssible approach to change the brittleness of ceramics in nature” *Mater Sci & Eng* **C11**, 9-12

Wang L, Adaptive Fuzzy Systems and Control: Design and Stability Analysis, New Jersey: Prentice Hall, 1994.

Wang X (1991) "Period-doublings to chaos in a simple neural network" *IEES INNS International Joint Conference on Neural Networks* 2, Seattle, 333-339

Watanabe S (1974) "Paradigmatic symbol - a comparative study of human and AI" *IEEE Trans Syst Man & Cyber* (Jan), 100-103

Watanabe, Y., Ishiguro, A. and Uchikawa, Y. (1999) “Decentralized behavior arbitration mechanism for autonomous mobile robots using immune network” *Artificial Immune Systems and Their Applications*, 187-209, Editor: D. Dasgupta, Springer-Verlag

Watkins C & Dayan P (1992) “Q-learning” *Machine Learning* **8** (3/4), 279-292

Watson R & Pollack J (2000) “How symbiosis can guide evolution” *Proc 5th European Conf on Artificial Life*

Watson, R. A., Ficici, S. G., and Pollack, J. B. (1999). Embodied evolution: Embodying an evolutionary algorithm in a population of robots. In *Proceedings of the Congress on Evolutionary Computation*, pages 335-342, Washington D.C., U.S.A. IEEE Press.

Watt R & Philips W (2000) “Function of dynamic grouping in vision” *Trends in Cog Sci* **4** (12), 447-454

Wax S & Sands R (1999) “Electroactive polymer actuators and devices” *Proc SPIE Conf on Electroactive Polymer Actuators & Devices*, Mar 1999, Newport Beach, CA, USA, **3669**, 2-10

Weaver W (1949) “Mathematics of communication” *Sci Amer* 181 (1), 11-15

Webber B (1983) "Logic and natural language" *IEEE Comp* (Oct), 3-28

Wegst U (1996) The mechanical performance of natural materials. PhD Thesis, University of Cambridge

Weiller B (1999) “Chemical microsensors for gas detection and applications to space systems” in Microengineering for Aerospace Systems, AIAA

Weiner S, Addadi L & Wagner H (2000) “Materials design in biology” *Mater Sci & Eng* **C11**, 1-8

Weischedel R (1986) "Knowledge representation and natural language processing" *Proc IEEE* 74 (7), 905-920

Weiss G (1995) “Distributed reinforcement learning” *Robot & Auton Syst* **15**, 135-142

Weiss L, Merz R, Pruiz F, Neplotnik G, Padmanabhan P, Schultz L & Ramaswanni K (1997) “Shape deposition manufacturing of heterogeneous structures” *J Manufact Syst* **16** (4), 239-248

Weizenbaum J (1966) "ELIZA - a computer program for the study of natural language communication between man and machine" *Comm Assoc Comp Mach* 9(2), 36-45

Weizenbaum J (1967) "Contextual understanding by computers" *Comm Assoc Comp Mach* 10 (8), 474-480

Welsch W et al (1997) “Immunosensing with surface acoustic wave sensors” *Sensors & Actuators A* 62, 562-564

Wendler, G. (2002) “The Organization of Insect Locomotion Systems and Computer-Based Flight Control in the Tobacco Hawkmoth *Manduca sexta*” In *Neurotechnology for Biomimetic Robots* (ed. Ayers J, Davis J, Rudolph A), Bradford Books, MIT Press, Cambridge, MA, USA, 451-468

Weng J (2002) “A theory for mentally developing robots” *2nd Int Conf on Development & Learning*, Cambridge, MA, IEE Computer Society

Weng J, McClelland J, Pentland A, Sporns O, Stockman I, Sur M, Thelen E (2001) “Autonomous mental development by robots and animals” *Sci* **291**, 599-600

Wenning A (1999) “Sensing effectors make sense” *Trends in Neurosci* **22** (12), 550-659

Werger, B. (1999). Cooperation without deliberation: A minimal behaviour-based approach to multi-robot systems. *Artificial Intelligence*, (110):293–320.

Wesson R et al (1981) "Network structures for distributed situation environments" *IEEE Trans Syst, Man & Cyber* 11 (1), 5-23

Wexler M, Kosslyn S, Berthoz A (1998) “Motor processes in mental rotation” *Cog* **68**, 77-94

Weybrauch R (1980) "Promulga to a theory of mechanised formal reasoning" *Artif Intell* 13, 133-170

Wheeler M (1996) “Embodied induction: learning external representations” *Proc AAAI Fall Symp on Embodied Cognition & Action*, ed. Mataric M, AAAI Press, USA

White G (1976) “Speech recognition: an overview” *IEEE Computer* (May), 40-53

White G (1990) “Natural language understanding and speech recognition” *Comm Assoc Comput Machin* 33 (8), 72-82

White R (1959) “Motivation reconsidered: the concept of competence” *Psychol Rev* 66, 249-333

Whitney D (1987) “Historical perspective and state of the art in robot force control” *International J Robotics Research* **6** (1), 3-14

Whitting, J., Safak, K., Adams, G. G., (2002) “SMA Actuators Applied to Biomimetic Underwater Robots” In *Neurotechnology for Biomimetic Robots* (ed. Ayers J, Davis J, Rudolph A), Bradford Books, MIT Press, Cambridge, MA, USA, 117-136

Widrow B & Lehr M (1990) “Thirty years of adaptive neural networks: perceptron, madaline and backpropagation” *Proc IEEE* 78 (9), 1415-1441

Widrow B & Winter R (1988) “Neural nets for adaptive filtering and adaptive pattern recognition” *IEEE Comp* (March), 25-29

Wilkinson, S. (2000) "'*Gastrobots*' - benefits and challenges of microbial fuel cells in food powered robot applications" *Autonomous Robots* **9**, 99–111

Wilson J (1984) “Robotic mechanics and animal morphology” *NATO ASI Series* F11 *Robotics & Artificial Intelligence* (ed. Brady M), Springer-Verlag, Berlin

Wilson M, King C & Hunt J (1997) “Evolving hierarchical robot behaviours” *Robot & Auton Syst* **22**, 215-230

Wilson S (1984) “Robotic mechanisms and animal morphology” *NATO ASI Series* **F11**, Robotics & Artif Intell (ed. Brady M et al), Springer-Verlag, Berlin

Wilson S (1991) “Animat path to AI” *Proc 1st Int Conf Simulation of Adaptive Behaviour (From Animals to Animats),* ed. Meyer J & Wildon S, MIT Press, Cambridge, USA, 15-21

Wilson S (1995). "Classifier fitness based on accuracy", *Evolutionary Computation*, 3(2), 149-175 .

Wilson S (1996) “Explore/exploit strategies in autonomy” *Proc 4th Int Conf Simulation of Adaptive Behaviour (From Animals to Animats 4)*, Cambridge, Mass, USA

Winograd T (1980) "What does it mean to understand language?" *Cog Sci* 4, 209-241

Winston P (1980) “Learning and reasoning by analogy” *Comm Assoc Comput Machin* 23 (12), 689-703

Winters J & Stark L (1985) “Analysis of fundamental human movement patterns through the use of in-depth antagonistic muscle models” *IEEE Trans Biomedical Engineering* **32** (10), 826-839

Wise K, Gillespie N, Stuart J, Krebs M & Birge R (2002) “Optimisation of bacteriorhodopsin for bioelectronic devices” *Trends in Biotechnol* 20 (9), 387-394

Wise S & Murray E (2000) “Arbitaray associations between antecedents and actions” *Trends in Neurosci* **23**, 271- 276

Wise, K.J., Gillespie, N.B., Stuart, J.A., Krebs, M.P. & Birge, R.R. (2002) Optimization of bacteriorhodopsin for bioelectronic devices. *Trends in Biotechnology* **20**, 387-394.

Wittwer, S.H. (1980) “The shape of things to come” In P.S. Carlson (ed.), The biology of crop productivity . Academic Press, New York.

Wolpert D (1997) “Computational approaches to motor control” *Trends in Cog Sci* **1** (6), 208-216

Wolpert D, Gluhramani Z & Flanagan J (2001) “Perspectives and problems in motor learning” *Trends in Cog Sci* **5** (11), 487-494

Woodridge M & Jennings N (1995) “Agent theories, architectures and languages: a survey” in Intelligent Agents, ed. Woodridge & Jennings N, Springer-Verlag, Berlin, 1-22

Woods W (1983) "What's important about knowledge representation" *IEEE Comp* (Oct), 22-27

Woods W (1986) "Important issues in knowledge representation" *Proc IEEE* 74(10), 1322-1334

Wootton R, Evans K et al (2000) "The hind wing of the desert locust (*Schistocerca gregaria*Forskål). I. Functional morphology and mode of operation." *Journal of Experimental Biology*203, 2921-2931

Wright J & Liley D (1996) “Dynamics of the brasin at global and microscopic scales: neural networks and the EEG” *Behav & Brain Sci* **19**, 285-320

Xu G, Yao N et al (1998) "Biomimetic synthesis of macroscopic-scale calcium carbonate thin films. Evidence for a multistep assembly process." *Journal of the American Chemical Society*120, 11977-11985

Yachandra, V.K., Sauer, K & Klein, M.P. (1996) "Manganese cluster in photosystem II: where plants oxidise water to dioxygen" *Chem. Rev.* **96**, 2927-2950

Yagishita, T., Sawayama, S., Tsukahara, K. & Ogi, T. (1996) "Photosynthetic bio-fuel cells using cyanobacteria" *Renewable Energy* **9**, 958-961.

Yagishita, T., Sawayama, S., Tsukahara, K. & Ogi, T. (1997) "Effects of intensity of incident light and concentrations of *Synechococcus* sp. and 2-hydroxy-1,4-naphthoquinone on the current output of photosynthetic electrochemical cell" *Solar Energy* **61**(5), 347-353

Yahiro, A. T., Lee, S. M. & Kimble, D. O. (1964) "Enzyme utilizing bio-fuel cell studies" *Biochimica et Biophysica Acta* **88**, 375–383

Yamaguchi G, Moran D & Si J (1995) “Computationally efficient method for solving the redundant problem in biomechanics” *J Biomechanics* **28** (8), 999-1005

Yen J (1990) “Generalising Dempster-Schafer theory to fuzzy sets” *IEEE Trans Systems Man & Cybernetics* 20 (3), 559-570

Yoshida K et al (2002) “Fabrication of micro electrorheological valves (ER valves) by micromachining and experiments” *Sensors & Actuators* A95, 227-233

Young J (1979) “Learning as a process of selection and amplification” *J Roy Soc Medicine* **72**, 801-814

Zadeh L (1973) "Outline of a new approach to the analysis of complex systems & decision processes" *IEEE Trans Syst Man & Cyber* 3(1), 28-44

Zadeh L (1978) "Fuzzy sets as a basis for a theory of possibility" *Fuzzy Sets & Syst* 1, 3-28

Zadeh L (1983) “Common sense knowledge representation based on fuzzy logic” *IEEE Computer* (Oct), 61-65

Zadeh L (1989) “Knowledge representation in fuzzy logic” *IEEE Trans Knowledge & Data Eng* 1 (1), 89-100

Zajonc R (1980) “Feeling and thinking: preference needs no inferences” *Amer Psychol* 35 (2), 151-175

Zak M (1991) “Unpredictable dynamics approach to neural intelligence” *IEEE Expert* (Aug), 4-10

Zebulum R, Pacheco M & Vellasco M (1997) “Evolvable systems in hardware design: taxonomy, survey and applications” *Proc 1st Int Conf on Evolvable Systems: From Biology to Hardware*, Springer-Velag

Zebulum R, Pacheco M & Vellasco M (1997) “Increasing length genotypes in evolutionary electronics” *Proc Int Conf on Genetic Algorithms*

Zeigler B (1989) “DEVS representation of dynamical systems: event based intelligent control” *Proc IEEE* 77 (1), 72-80

Zent A, Quinn R & Madou M (1998) “Thermo-acoustic gas sensor array for photochemically critical species in the martian atmosphere” *Planet Space Sci* 46 (6/7), 795-803

Zhang H, Dai Z & Gorb S (2004) “Study on adhesive properties of materials bassed on biomimetic results of Gecko’s feet” *J Bionics Engin* **1** (1), 66-71

Zhang S & Gonsalves K (1995) "Chitosan calcium carbonate composites by a biomimetic process." *Materials Science and Engineering C-Biomimetic Materials Sensors and Systems*3, 117-124

Zhang S, Wright G & Yang Y (2000) "Materials and techniques for electrochemical biosensor design and construction" *Beisensors & Bioelectronics* **15**, 273-282

Zhang, X. & Halme, A. (1994) "A summary of the study of bioelectrochemical fuel cell using *Saccharomyces cerevisiae*" *Research reports of Automation Technology Laboratory of Helsinki University of Technology* **10**

Zhao X, Zhou B et al (1996) "A biomimetic model of short fibre reinforced composite materials. Iii intensitizing effect of dumbbell fibre-reinforced composite with brittle matrix." *Jinshu Xuebao*32, 438-442

Zheng Y & Paul R (1985) "Hybrid control of robot manipulators" *Proc IEEE Int Conf Robotics & Automation*, 602-606

Zhou H (1989) "CSM: a computational model of cumulative learning" *Cognition,* 383-406

Zhu B (2000) "Bio-inspired study of structural materials" *Mater Sci & Eng* **C11**, 33-18

Ziegler Ch et al (1998) "Bioelectronic noses: a status report. Part II" *Biosensors & Bioelectron* 13, 539-571

Zielinska T & Hang J (2002) "Development of a walking machine: mechanical design and control problems" *Mechatronics* **12**, 737-754

Ziou B (1995) "Biomimetic design of worst bonding interface for ceramic matrix composites."*Composites Engineering* 5, 1261-1273

Zioupos P (1998) "Recent developments in the study of failure of solid biomaterials and bone: 'fracture' and 'pre-fracture' toughness." *Materials Science & Engineering C-Biomimetic Materials Sensors and Systems* 6, 33-40

Zipser D (1985) "Computational model of hippocampal place fields" *Behav Neurosci* **99** (5), 1006-1018

Zrehen S & Gaussier P (1997) "Building grounded symbols for localisation using motivations" *4th European Conf on Artificial Life*, ed. Husbands P & Harvey I, MIT Press

Zufferey, J.C., Beyeler, A. and Floreano, D. (2003) Vision-based Navigation from Wheels to Wings. In Proceedings of the IEEE/RSJ *International Conference on Intelligent Robots and Systems*, pp. 2968-2973.

Zufferey, J.C., Floreano, D., van Leeuwen, M. and Merenda, T. (2002) "Evolving Vision-based Flying Robots" Bülthoff, Lee, Poggio, Wallraven (Eds.), *Proceedings of the 2nd International Workshop on Biologically Motivated Computer Vision*, LNCS 2525, 592-600, Berlin, Springer-Verlag

List of Acronyms, Abbreviations, and Terminology

ACT	Advanced Concepts Team
AI	artificial intelligence
AO	announcement of opportunity
AOCS	attitude and orbit control system
BAS	British Antarctic Survey
CBNT	Centre for Biomimetic & Natural Technologies
CCNR	Centre for Computational Neuroscience & Robotics
CDH	command and data handling
EPFL	Ecole Polytechnique Federale Lausanne
ESTEC	European Space Research & Technology Centre
GOFAI	good old fashioned AI
ITT	invitation to tender
OBDH	onboard data handling
PI	principal investigator
PM	project manager
PRR	preliminary requirements review
SMP	study management plan
SoW	statement of work
SSC	Surrey Space Centre
SSTL	Surrey Satellite Technology, Ltd
TBC	to be confirmed (by the Agency)
TBD	to be determined (by the contractor)
TBS	to be specified (by the Agency)
TN	technical note
TT&C	tracking telemetry and command
WP	work package
WS	work shop

List of Figures

List of Tables

Table of Contents

0. EXECUTIVE SUMMARY

In this report, we provide a snapshot survey of the field of biomimetics around the world. We have emphasised application with respect to robotics in anticipation of future potential application for robotic space exploration missions. We anticipate that the next generation of space explorers will emphasise in-situ exploration with mobility capability. This is particularly the case as space exploration missions give way from global reconnaissance to more focussed *in-situ* investigation requiring enhanced capabilities. To that end, we have emphasised those aspects of biomimetics which have application to robotic rovers.

There is thus much that engineers can learn from biology to emulate in their design of engineering systems – this is biomimicry (or biomimetics, or bionics, or bio-inspiration). The success of biological organisms is attributed to the process of natural selection whose primary metric is such success – failure implies extinction. No wonder engineers have begun to examine biological systems to learn how organisms solve problems. It seems appropriate therefore to examine biological solutions to common biological and engineering problems. The reverse engineering of biological organisms has recently become a significant research effort to solve engineering problems by learning from biological solutions to similar problems imposed by the natural environment. A spacecraft is coupled to a hostile and highly variable physical world. Biomimicry applied to space engineering promises the prospect of high miniaturisation, high integration and packaging efficiency, high energy efficiency, and high autonomy and robustness. Such robustness and adaptability is particularly critical in space exploration as the environments to be explored are typically unknown with unknown dynamics and variability. In particular, as space exploration missions demand greater sophistication and functionality, so the requirements for miniaturisation, autonomy, robustness, and survivability become paramount. A critical issue for space systems is that of autonomy – the distances involved, particularly once spacecraft venture beyond Earth orbit, preclude real-time control, necessitating high degrees of onboard autonomy. Biological systems exhibit autonomy *par excellence*. We suggest that the primary *raison d'etre* for adapting biological solutions to the design of spacecraft and space missions is to enhance autonomy, i.e. relieve the ground-based (for robotic spaceflight) or spacecraft-based (for human spaceflight) human personnel from interacting with the space mission. This implies the need for extensive adaptability, robustness and reliability. We highlight five main broad biological principles that are particularly relevant to space systems:

1. The principle of autonomous behaviour to minimise reliance on ground systems
2. The principle of integration of mechanics, control, software and electronic systems
3. The principle of robustness for survival under widely variable conditions
4. The principle of compliant, multi-functional structures as part of a control system
5. The principle of neurally-inspired control systems with selectionism (e.g. neural Darwinism)

Natural selection is imposed on organisms by their interaction with the physical world. This requirement for interaction with the real world is also imposed on engineering structures that must adapt to their surroundings such as robotic spacecraft. This interaction with the real world is the key to biological evolution – evolution effectively incorporates aspects of the environment into the structure and behaviour of the organism which enhance survival (or more properly, inclusive reproductive fitness). This is how organisms are "designed" by natural selection. The artificial carving of the physical world into electronics, mechanics, structures and control systems, as is traditional in engineering, is one which is not adopted in biological evolution. Unlike conventional engineering approaches, biological solutions are integrative involving simultaneous co-evolution of mechanical structures with control systems to provide their functionality.

We have somewhat artificially divided our approach into a number of mutually overlapping arenas:

Biological modularity and pleiotropy: The first aspect of biomimetics in its potential application to spacecraft engineering that we consider is the issue of systems design. Modularity and hierarchical organisation is a fundamental approach in both science and engineering as the means for problem reduction, necessary for complex systems. Modularity is characteristic of biological organisms – The primary reason for the evolution of modularity in biological organisms is to suppress pleiotropic effects between different modules. This is the basis of the building block hypothesis underlying the operation of genetic algorithms which ensures that sexual reproduction does not destabilise the gene string. Pleiotropy is primarily confined to morphogenesis, which involves complex interactions of genes in the form of gene networks. Modularity is central to cognitive psychology models of human information processing (Fodor's thesis) - specialised neural modules are attributed to certain functions of the human mind. However, the behaviour-based robotics paradigm suggests that modularity should follow evolutionary lines based on competence – this is the basis of the hierarchical subsumption architecture for mobile robot control systems. This approaches bases its modularisation on behavioural lines rather than functional lines. The subsumption architecture may be represented by the IF conditional to implement a decreasing priority of behaviours.

There is also the related issue of hierarchical structure - hierarchical structures are common in natural, biological and artificial systems. Indeed, the hierarchical structure is universal. The hierarchy is a method of decomposing complex systems into smaller interrelated subsystems nested into levels until some lowest level is reached. Hierarchies are based on successive delegation of functions from higher to lower levels, so that the number of distinct tasks increases down the hierarchy, each task at each level characterised by more limited contextual horizons. The hierarchy is essential when the number of behaviours becomes large on the basis of division of labour. The hierarchy is characterised by pyramid-like structure of decision and goal levels varying in complexity. Each level is coordinated by the level above and higher levels are distinguished by longer time constants. If the potential for interference occurs between functions, division of labour through hierarchical partitioning is necessary. This requires that coupling between subsystems is low, so that the system is decomposable into independent subsystems to minimise internal coordination between subsystems. Near decomposability depends on hierarchical partitioning to separate the short term dynamics of the subsystems from the long term dynamics of their interactions. Hierarchical structures provide maximum efficiency of information processing systems. Human physiology is based on multiple homeostatic feedback systems, the major ones being the cardiovascular system, the respiratory system, the digestive system, the excretory system, the nervous system and the endocrine system – however, these physiological control systems interact in complex ways which defy simple hierarchical/modularity analysis except at the most general level.

In conclusion, modularity and hierarchical organisation are ubiquitous in biological systems where they are feasible. This suggests that the reductive engineering approach is sound as long as undue simplification is unwarranted. Furthermore, biological modularisation may not reflect engineering functionality. In cases where high degrees of interaction occur, complex systems analysis is necessary.

Biomimetic materials, structures and mechanisms: Biological materials and structures are characterised by multiple levels of organisation which impart toughness and multi-functionality in comparison with more traditional engineering materials. Biological materials possess their unique characteristics by virtue of their organisation rather than their material substrate. Cellular compartmentalism within structures is the approach exploited in honeycomb panels for spacecraft primary structures. Most biological materials are soft with stiffness imparted through mechanical rigidisation by fluid pressure or through viscosity changes in gel materials – the biological analogue to electrorheological fluids. They are generally compliant with low Young's modulus and provide strength for resistance to breakage rather than stiffness for resistance to deformation. Composite materials are ubiquitous in biological systems in which the structure incorporates sensors, processors and actuators. Most biological materials are composites comprising polymer fibres embedded within matrix materials – like engineered composites, their physical properties are highly anisotropic. The outer covering of arthropods (insects, spiders, crabs, shrimps, etc.) is known as cuticle, which is a composite material consisting of arrangements of highly crystalline chitin nanofibres embedded in a matrix of protein, polyphenols and water, with small amounts of lipid. Reinforcement of fibres with minerals such as in bone to impart high stiffness. Bone offers structural adaptability to imposed loads by virtue of their dynamic equilibrium between bone loss and growth. The chief defining characteristic of biological materials is their integrative multi-functionality by virtue of microstructure. Wood is able to store elastic strain energy and absorb shocks due to the presence of holes. Holes can be used as strain gauges. Arthropods have strain sensors which are basically holes through their exo-skeletal sheath. These organs allow the animal to measure displacements of molecular size, and they do so by introducing compliance. Fibres also tend to carry the load around the hole, so that although the hole concentrates strain, the fibres dissipate stress.

The manufacturing technique most applicable to such structures is shape deposition manufacturing (SDM), a rapid prototyping manufacturing technology in which 3D parts manufacture and assembly occur simultaneously. Material is deposited and shaped in layers with components embedded into the workpiece when required. SDM offers the capability of creating multiple materials, sensors, actuators and control within a single monolithic structure. The development of multi-functional structures is a step in this direction which incorporate energy distribution – electrical wiring, thermal piping, optical fibres, etc.

Geckos and several other animals such as species of beetles and spiders have the remarkable ability to climb on smooth surfaces, even supporting their own weight on the underside of such surfaces. Such animals are noted for their possession of pads of small hairs (setae) on their feet which are the source of their climbing ability. The main mechanism of setal adhesion in geckos is van der Waals attraction between the spatulae and the surface with which they are in contact. The forces are due to the close contact between the surfaces rather than any specific chemical properties.

Smart/adaptive materials modelled on biological systems are of great interest – they typically incorporate sensors which act as perceptual processing systems, actuators which act like muscles, and a real-time control systems acting like as a nervous system. Actuation mechanisms vary but they generally involve changes in stiffness in response to environmental conditions without the use of moving parts: shape memory alloys, electrorheological fluids and piezoelectric solids. Sensors that may be embedded into the composite material during manufacture may be optical waveguides or microelectronic devices of which fibre optic systems appear to be the most promising due to their capability to withstand the strains of the materials manufacturing process. The polyMEMS actuator (PMA) is composed of micro- to milliscale unit cells configured in a three-dimensional array. A large 3D array could be built to generate large forces and large displacements simultaneously. The polyMEMS actuator is fabricated using a simple process adapted from flexible printed circuit manufacturing. Shape memory alloys exhibit thermoelastic behaviour changes such that they can exist in two temperature-dependent crystal phases separated by phase change transitions. The commonest shape memory alloy is the NiTi alloy known as Nitinol. Shape memory alloys are limited to ~2 Hz frequency oscillations and ~4-8% extensions but are capable of generating high forces (but are limited to 2% recovery strains for deformation/recovery cycling). Generally, Joule effect heating by the application of voltage is exploited for controlling SMA devices with control implemented through the measurement of electrical resistance (rather than temperature). Shape memory alloys may be embedded as 200-400μ plastically elongated wires or springs constrained from recovering their normal memorised length during fabrication. The fibres may be arranged such that they induce a state of strain and the structure thereby altering the structure's stiffness, natural frequency and modal responses. Piezoelectric materials are insulators that produce charges on their surface when mechanically deformed. Piezoelectric ceramics and polymers become electrically polarised due to changes in their resistivity when subjected to mechanical stress. Piezoelectrics are 10-20 times more sensitive than metal film strain gauges and are regarded as part of MEMS technology in which they are exploited in the form of pressure transducers. Generally piezoelectric ceramics such as PZT are more efficient than piezoelectric polymers such as PVDF (which is also limited to low temperature and pressure environments). Electrostrictive and magnetostrictive materials appear to give enhanced performance for active damping over piezoelectric materials. An electrorheological fluid comprises a dielectric carrier fluid in which semiconducting particles are dispersed. Electrorheological and magnetorheological fluids alter their viscosity in response to applied electric or magnetic fields due to polarisation of the particles causing them to form chains with a given orientation to the applied field (Winslow effect). Particle abrasion can cause wear and ER fluids tend to have high sensitivity to temperature. Polymeric gels comprise a tangled network of cross-linked polymer chains immersed in a liquid – when stimulated by acids and alkalis, they can swell or contract as "pH muscles". They have limited performance as they are dependent on ion diffusion time and are not well suited to space application due to their requirement for aqueous solution and high temperature sensitivity. Electro-active polymers (EAP) are electromechanical systems and are used to convert mechanical strain to electric current thereby acting as sensors. Electroactive polymers (EAP) which are lightweight strips of flexible plastic that bend or stretch in the presence of an applied voltage – these offer the potential basis for the reciprocating contractile linear actuator properties of artificial muscles. They provide fast response and good controllability. There are two types of electroactive polymer – ionic polymers and electronic polymers. Electronic electro-active polymers deform by virtue of electronic carrier motion but this requires high voltages ~100 V/μm close to breakdown level. Ionic electro-active polymers are current driven rather than voltage driven allowing the use of low voltages. Ionic polymers, conducting polymers and polymeric gels perform well at low frequencies but cannot operate at high frequencies due to the inherent slowness of ionic diffusion. However, ionic polymers generally require moisture to operate as they are hydrophilic but may be sealed within flexible coatings. EAPs can be rolled into ropes to act like muscle fibres. Force amplification may be achieved through series and parallel configurations of such devices. Such materials can operate at low temperatures ~-140^{o}C and in a vacuum. Conducting polymers convert electrical energy directly into mechanical energy with high force generation capabilities and high power output densities yet require only low voltage input with much greater superiority over polymer gels. However, they have limited lifecycles ~10^6 and limited energy conversion efficiencies. Their key differentiating characteristic from electrostatic, piezoelectric and electrostrictive actuators is that the electrically-generated strain is dependent on the applied voltage rather than the applied field. They have the advantage that they can be maintained at any actuator response position without energy input, though they can suffer drift. Their response times are ~0.3 ms, much faster than biological muscles ~1 ms, and can exert a maximum force of 80-100 times that of muscles.

Smart materials offer promise of artificial muscles to replace electric motors in space systems, eliminating the problem of susceptibility of moving parts to single point failure.

Biological material is typically "soft" with structural stiffness commonly imparted through mechanical rigidisation of fluid pressure (turgor). Biological materials are generally highly compliant with low Young's modulus and provide toughness for resistance to fracture rather than stiffness for resistance to deformation. In addition, structural configuration is exploited in biological systems. The cellular cytoskeleton provides tensegrity (tensional integrity) in which protein chains provide both compression (as in rods) and tension (as in wires) for structural integrity within a stable but flexible structure. Compression-bearing struts resist opposing tension-bearing cables in a mutually stabilised network. Leaf structures may be packaged and then deployed which have applications in the deployment of antennas and solar panels (eg. solar power satellites). Osmotic pressure in plant cell walls provides the basis for such actuation mechanisms. Although the osmotic pressure is low, it is amplified by the storage of elastic energy in the cellulose structure. Insects adopt a similar strategy to store energy in resilin within the cuticle to drive their wingbeats during flight. Spiders transmit hydraulic pressure in their blood by compressing their bodies to extend their legs. Such techniques may be exploited for deployment mechanisms, eg. RolaTube.

Biomimetic methods of propulsion: Movement is a necessary part of exploration and animal locomotion provides the basis for animal survival – the avoidance of threats and the acquisition of resources. The ability to move is the key to intelligence – animals which are mobile (by virtue of their heterotrophic energy assimilation) possess nervous systems, while plants which are static (by virtue of their autotrophic energy assimilation), do not. The key to biomimetic locomotion of all types is the approach to control – robust locomotion requires the implementation of central pattern generation, sensory feedback and preflexes enabled through structural compliance. Bipedal robots require active balancing as they lack static stability and bipedal walking is currently still not robust.

The most common configuration is the insect-inspired six legs which always ensure that three legs are in contact with the ground in motion. The animal minimises the lift height of the leg which is determined by the roughness of the terrain. Legged locomotion in animals relies heavily on the compliance of the body provided by the spine in the case of vertebrates, and in the body segment joints in the case of arthropods. An integrative approach to locomotion focuses on the interactions between the muscular, skeletal, nervous, respiratory, and circulatory systems. Walking locomotion has a number of advantages including greater robustness – discrete footholds rather than complete free pathways are required. This is essential for rough terrain typical of planetary environments, such as Mars, which are covered in rocks. Walking overcomes irregular terrain negotiation by dynamically decoupling the body of the vehicle from the ruggedness of the terrain. This is enabled by a large number of degrees of freedom in the legs allowing for simple adaptability to changes in terrain. Terrain contour following is additionally enhanced by body articulation. Central pattern generators are neural circuits which generate rhythmic dynamics and may be constructed from either pacemaker neurons or oscillatory circuits such as the nervous net, eg. modelled on the cerebellorubrospinal pattern generator. Coordinated control of the legs emerges from interactions between the individual leg controllers. The most common walking patterns in insects are the tripod and tetrapod gaits. Insect locomotion is controlled by neurally-based central pattern generators sometimes augmented by feedback mechanisms. Proprioceptive feedback from stretch receptors is essential for maintaining the frequency and duration of rhythms of central pattern generators. The transition points between gaits or when a decision to sway from the standard gait are called extreme points. Walknet is able to control walking at different speeds with an emerging tripod or tetrapod gait pattern, curve walking without any explicit calculation of leg trajectories, walking uphill and downhill, climbing over obstacles, and walking with partly amputated legs. Insect walking is not controlled by a hardwired central system such as fixed action patterns or central pattern generators organized in a hierarchical structure, rather than by a decentralized architecture consisting numerous different and somewhat independent modules that have to be coupled to cooperate. There is a hierarchical mechanism such that 1) each leg has its own control system which generates rhythmic step movements and 2) several coordinating mechanisms couple the movement of the legs to produce a proper gait. Gait oscillations result from the cooperation of different subsystems and the physical world. The cockroach offers sophisticated leg design with increased agility. They are among the fastest runners of all insects. One of the simplest approaches – RHex – offers robust locomotion over hostile terrain. The legs rotate in a full circle and the three legs of a tripod are driven simultaneously through a slow "retraction" phase, as the opposing tripod begins its "protraction" by rotating away from the ground contact.

Body compliance is taken to its extreme in the case of snakes and centipedes. The snakebot concept is the ultimate form of body articulation and has good adaptability to rough terrain, stiffening into a bridge to cross crevasses, access to orifices, winding around obstacles, low path cross-section requirement, redundant structure in case of joint failure and modular construction. The motion comprises consecutive folding/unfolding of the body with high degrees of friction between the body and the ground. The snake moves by passing a sine wave laterally through its body while simultaneously curving the body so that the peaks are lifted off the ground. A number of different snakebot "gaits" are possible such as side-winding, coiling or flipping. The active cord mechanism (ACM) has been proposed as a general form of snake locomotion and/or manipulation, applicable also to aquatic motion and ciliated/flagellate motion. Ground-penetrating moles/drills may exploit similar motions in order to access sub-surface sites. We consider that snake-like devices have limited utility in that they introduce difficulties in instrument packing and limited force exertion.

Of the arthropods, only insects evolved wings and they were the first animals to evolve flight. Flying is 10 times as energy intensive in time as ground locomotion (power intensive), but is 4 times more efficient in ground coverage (energy per km). Hence, the flight muscles of insects have the highest metabolic rate of any biological tissue. Dragonflies and damselflies (Odonata) were among the first winged flying insects. They have two similar pairs of membranous wings. The two sets of wings are flapped 180^{o} out of phase. This ensures that a lift component is maintained through the wingbeat cycle. They beat their wings at 20-30 Hz. Unlike bats and birds, insect wings are not permeated by muscle but act as levers. Vertebrate wings however have considerable *controllable* flexibility to act as control surfaces. Most insects possess two pairs of wings attached to muscles hinged with resilin to the thorax which flex the wing to manipulate the leading and trailing edges. Insects employ vortex generation to gain lift and are characterised by high wingbeat frequency for a high coefficient of lift. Insect wings are thin with sharp leading edges which generate leading edge vortices which generate most of the lift. Insect wings' leading edge tilts down during the downstroke and upwards during the upstroke. The wing is thus rotated at the end of each stroke. Due to this rotation, the tip of each wing traces out a thin oval tilted at a steep angle. The delayed stall causes a leading edge vortex of high flow velocity to form just above and behind the leading edge on the downstroke which initially increases lift and spirals to the wing tip and shed before the upstroke. The thorax aids wingbeats by elastically deforming such that the wing is stable only in the extreme positions and are dynamically unstable in mid-position – this is enabled by the distortion of the thorax and the hinge articulation. The Entomopter project has the goal of developing a fully functional autonomous insect capable of crawling, swimming and flying. RoboFly is a flapping wing insect analogue that is currently being built.

For fine manipulation, the human hand is a single 22-degree-of-freedom tool of unprecedented multi-tasking capability by virtue of different grips. The most sophisticated robotic hand currently under development is the Robonaut hand for performing EVA-equivalent functions. Two fingered grips must have the fingers arranged parallel to each other, while three fingers provide greater stability, typically arranged 120^{o} from each other. The precision (prehensile) grip involves contact with the object by the fingertips while the power grip uses the fingers and palm to clamp the object - the parallel jaw gripper cannot perform a power grip. More complex grippers involve a minimum of three fingers, each with three or four degrees of freedom (the fourth degree of freedom emulates that of the human palm) to emulate the diversity of grips exhibited by the human hand. The DLR dextrous Hand II which is capable of precision and power grasping - power grasping requires parallel alignment of second, third and fourth fingers while precision grasp requires opposition of first finger and thumb with intersection of other fingers.

The problem of structural stiffness/compliance is fundamental in locomotion and manipulation. Biological structures tend to be close to critically damped unlike artificial mechanical structures which are heavily under-damped generating the problem of vibration. Closed loop control cycles – particularly those associated with force control – limit the rapidity of response so structural compliance is essential for rapid adaptation to environmental conditions, especially as force control is essentially chaotic as small changes in applied torques generates divergent sensed forces on contact with stiff environments ~10^{4}–10^{6} N/m which cannot in general be compensated for by closed loop force control. Passive compliance in the structure absorbs the energy of impact and effectively reduces the stiffness of the environment. Structural compliance provides the basis for "preflexes" within a control system for dealing with interaction forces. Pure feedback control cannot be employed in biological limb movement as neuronal feedback loops would be too slow. The control of stiffness provides the basis for using impedance to control externally applied forces. Muscles behave as tunable springs to regulate the timing and amplitude of the viscoelastic properties of the joints to move between positions. Muscle length-extension defines the equilibrium position and stiffness of each joint – this may be modelled as a potential function of the joint angle (the derivative of which is force). This method is consistent with Hill's model of muscular action.

Biomimetic energy generation: One of the most promising potential developments in biomimetics for space applications is artificial photosynthesis and the development of related redox reactions. It is a much more refined method of gathering solar energy that is much more efficient per unit area than current solar panel technology. The principle behind the chloroplast is important because a similar system might be used to contain artificial photosynthetic apparatus in space. Within the thylakoids are chlorophyll a and b, which are the light harvesting compounds responsible for gathering sunlight. The Stark Einstein Law states that light harvesting pigments can only take up one photon at a time. In natural photosynthesis, light is first gathered by chlorophylls and carotenoids. Chlorophylls absorb in the blue and red regions and carotenoids only absorb in the blue region of the spectrum. One of the major challenges in artificial photosynthesis is the transfer of electrons from the pigment harvesting the light to the reaction centre. In natural photosynthesis this charge transfer occurs from the chlorophylls or the carotenoids to the chlorophyll a in the reaction centre. The process of converting the light harvested by chlorophyll into the photosynthetic reactions is carried out by two photosystems I and II. PSII is based on the P680 reaction centre (which oxidises water) while PSI is based on the P700 reaction centre (which reduces NADPH+ to NADPH). One promising avenue of artificial photosynthesis research is the use of porphyrins and fullerenes as building blocks for artificial photosynthesis systems. Porphyrin-based materials are already exploited as conducting polymers. The key requirement is to achieve multiple electron transfer so that each charge separation prolongs the life time of the final charge-separated state. The superior performance of the fullerenes is explained by the fact that the charge in the fullerene anion radical is distributed over the entire C_{60} structure, rendering the reorganization energy smaller than for the corresponding quinones, where the charge is localized mainly to two oxygens. The fullerene and porphyrin systems and other systems like them can act as the basis for the electron transfer reactions. A further challenge in artificial photosynthesis is to replicate the reactions on the photosystem II side of photosynthesis – that is the reactions involved in the oxidation of water and the use of light driven charge separation to drive water oxidation and thus the production of oxygen. Efforts to replicate water oxidation to date have involved the linking of manganese to a photosensitizer such as rubidium complexes. Proton pumps based on bacteriorhodopsin or ATPase will be essential for designing artificial energy systems for space application.

Fuel cells are an established method of power generation using "cold combustion" of a fuel (typically hydrogen, methane or methanol); the oxidiser is usually oxygen, but other substances such as sodium ferricyanide solution are often used for the biological type. They can generate electricity without significant temperature rise so they are not subject to the thermal cycling limits of the second law of thermodynamics. Anaerobic micro-organisms have been used to convert waste organic matter into flammable gases (principally hydrogen and methane), which could be burnt to produce power. The difference between conventional and biological fuel cells lies in the substances that catalyse the chemical reactions, which in the former type are usually transition metal alloys (such as platinum-rhodium or nickel-tin), whereas EFCs use enzymes. In contrast to the EFC described above, the catalyst for the oxidation of fuel at the anode of an MFC is a culture of living microorganisms – either bacteria or yeast. The oxidation of food transfer energy to the compound NAD (nicotinamide dinucleotide) to produce its reduced form, NADH. This NADH represents the reservoir of energy that can be tapped to produce electricity. In mediated EFCs, a synthetic chemical compound is added to the culture whose purpose is to ferry the stored energy of the NADH from inside the microbes to the anode. In “mediator-less” MFCs, micro-organism’s which use extracellular ferrous oxide (rust) as their electron acceptor instead of oxygen are employed which use graphite electrodes directly as an electron acceptor. Photosynthetic MFCs are very similar to mediated MFCs in construction, with a sodium ferricyanide cathode and a synthetic chemical mediator at the anode. The mediator transfers the chemical energy from the products of the photolysis of water to the electrode, and thus is subject to the same criteria as a mediator in a conventional mediated MFC. A typical power area density observed was 1.9W/m^2, with a light conversion efficiency of 3.3%. There are two examples of macroscopic robots powered by biological fuel cells have also been constructed – Gastronome powered by sugar and EcoBot powered by fly carcasses. Both require lengthy charging periods with comparatively short periods of activity.

Immunological approaches to self-defense: The immune system operates in two stages: first, recognition of substances that are foreign and therefore could cause harm (termed “non-self”, as opposed to the “self” substances that make up the organism), followed by their destruction. Computer programs have been written using software analogues of the recognition mechanism. The distribution of control between the ground and the space segment through the communications link could potentially yield viral infection of the spacecraft onboard computer. This is particularly the case when uploading software from the ground to the spacecraft. Viruses are usually introduced as Trojan horses which appear innocuous and can comprise of only a few lines of code. The ability of viruses to replicate and infiltrate programs as Trojan horses coupled with the widespread use of sharing and networking in computer systems makes them a considerable threat. A 20 byte signature often forms a family signature which is common to several distinct viruses of the same family. This pattern matching process is similar to biological viral receptors which bind to 8-15 amino acid sequences in viral protein coats. An antibody is a Y-shaped protein with two identical antigen-binding sites at the tip of the Y and complementary binding sites on its Fc region. If an antibody meets an antigen, the two bind together thereby neutralising the antigen. The strength of antigen-antibody interaction depends on the affinity and number of binding sites. Antibody immunoglobins are large molecules with a characteristic structure – the Ig fold – which serves to recognise non-self molecules. Immunoglobins are large proteins of two heavy and two light polypeptide chains, the different types being distinguished by their heavy chain structures. The immune system provides a fuzzy match to the virus. In artificial systems, the antigen-antibody matching is modelled by comparing test data with sequences of characters (strings) called “detectors”, the equivalent of antibodies, and determining their similarity. The symbols making up the detectors can represent anything from individual bits up to commands in computer programs – strings of bits are most popular. If the two strings “match” then a virus must be present and appropriate measures should be taken. The fuzzy match to a short sequence of bytes in the virus that comprise its signature must be efficient in time and retain the ability to recognise variants and mutations of the virus. Once detected, the viral code is deleted and the infected file repaired. The antibody program resides in the operating system until a virus triggers its activation. In particular, computational analogues of the natural immune system are based the B-lymphocyte. There are two classes of lymphocyte – B cells which make antibodies and T cells which are involved in cell-mediated immunity response. Each B cell clone makes antibody molecules with a unique antigen binding site. T cells undergo transformation and mitosis generating a population of specifically reactive anti-antigen cells. B cells differentiate into plasma cells which secrete antibodies in cooperation with T-helper cells. T suppressor cells interact with B cells to inhibit antibody synthesis. When the antigen binds to the antibody in the membrane, B cells are activated to multiply and synthesise quantities of the antibody with the same antigen-binding site which is secreted into the blood. Clonal selection operates by the production of a vast diversity of lymphocytes which are selected by fit to the antigen. Once selected, those lymphocytes are clonally reproduced and perpetuated as antibodies. The body can produce antibodies against a vast variety of potential antigens, i.e. the ability is genetically determined (clonal selection). The immune system analogue has also been proposed for fault tolerance in multi-computer system architectures - immunotronics.

Bio-inspired hardware research may be partitioned into three orthogonal directions – phylogeny (P – evolution analogies such as evolvable hardware), ontogeny (O – development analogous such as embryonics) and epigenesis (E – learning analogues such as immunotronics). There are research activities that combine these, eg. evolutionary neural networks (PE). There are areas confined to electronic hardware in using genetic algorithms to evolve field programmable gate array (FPGA) architectures to solve control logic problems. FPGAs with ASICs (application specific integration circuits) are becoming increasingly commonly used for dedicated computations. Evolved hardware controllers are not programmed to follow a sequence of instructions, rather they are configured and allowed to behave in real time according to semiconductor physics. Artificial hardware evolution has been to design hardware circuits as on-board controllers for two-wheeled autonomous mobile robots displaying simple wall-avoidance behaviour in an empty arena. Evolutionary hardware has also been used for developing self-repairing circuits.

Biological behaviour control and navigation: The robotic system provides the most demanding test of AI techniques since it requires the intelligent control system to interact with the real world via a physical robot body. Sensory input and motor output are not analogues of the standard peripheral read and write commands in a computer programming language. Bodies are a necessary pre-requisite for intelligence and that disembodied computing systems will never exhibit truly intelligent behaviour. The body provides a fundamental system of reference for cognition. All animals that are mobile require general-purpose perceptual and behavioural repertoires sufficient to deal with a variety of different environments in order to survive, eg. foraging strategies for food require complex navigational skills. The simplest form of behaviour is the reflex which is characterised by a rigid relationship between stimulus and response. A rapid stereotyped response is triggered by a particular type of environmental stimuli and it is usually protective against sudden environmental changes. Much animal behaviour is the result of innate responses to certain environmental stimuli which allow animals to respond to different situations. Fixed action patterns are more complex involving temporal sequences of actions and are extended responses to triggering stimuli. Animal behaviours as instinct centres are hierarchically organised with negative feedback connections between them. The three relevant basic centres are locomotion, food-seeking and pain aversion to generate feeding and fleeing behaviours under the appropriate conditions. Each instinct centre is decomposed into finer-grained behaviours at lower levels.

The animat approach has become dominant in robotics in which autonomy, survivability and robustness are the key characteristics inspired by the capabilities of insects despite their limited neural resources in which ecological niche is an important design consideration. It focusses on complete systems and on the tight coupling interaction between the robotic agent and the environment in which it is situated (situated robotics). Braitenburg vehicles indicate that simple mechanisms can yield complex behavioural manifestation in the appropriate environments. Behaviour control methods based on insect behaviours are an approach which have been highly successful. The behaviour approach seeks to build intelligent control in robotics by developing modules in which each is capable of task competences and which performs its own sensing, behaviour-generation and actuation. The "nouvelle" AI paradigm called the Physical Grounding Hypothesis suggests that autonomy requires representation to be grounded in the physical world – the incompleteness of simulated world models can miss important parameters, ie. the world is its own best model. Reflexive behaviour is characterised by simple stimulus-response behaviour which can generate sophisticated and robust behaviours. A basic set of behaviours comprises of: wandering, avoidance, following, aggregation, dispersion and homing for individual and groups of mobile robots. The behaviours avoidance (of obstacles) and following (targets) imply the ability to discriminate. Each behaviour is selected by the environmental conditions. The behaviours may be combined to generate higher level behaviours. The subsumption architecture is a hierarchical organisation of behaviour based units to provide robust and reflexive robot performance with a precedence organisation to avoid multiple activations of mutually exclusive behaviours. To eliminate conflicts, higher hierarchical level behaviours can suppress lower level behaviours so only one layer controls the robot at any one time. Different combinations of active behaviours yield more complex behaviours. Wandering and avoidance generate safe wandering. Wandering, avoidance and aggregation combined produce flocking. The addition of homing produces flocking towards a specific location. General foraging can be accomplished by combining different behaviours triggered by different conditions, e.g. temporal switching between avoidance, dispersion, following, homing and wandering. Hence, group behaviours can be generated by local environmental conditions.

An extension to the behaviour based approach is through the use of cognitive schemas which may be combined to yield a cognitive map. Schemas are collections of generic units of information with variable slots associated with generalised plans or behaviours. Scripts are schemas for frequently occurring sequences of events enabling temporal ordering of events to be represented. They specify general courses of action by incorporating a series of interrelated decision rules. The potential field provides a form of cognitive map which are necessary for animal foraging search strategies. Obstacles exert repulsive forces while the goal exerts an attractive force. A generalised potential field is a function of both position and velocity of the robot which gives a smoother path in following the potential minimum path. The most significant problem with the potential field is that local minima occur where the solution becomes trapped. It is then necessary to impose a global approach to overcome local minimum, or alternatively to add a probabilistic Monte Carlo element, or techniques such as simulated annealing. The lack of "memory" in the behaviour-based methodology limits the sophistication of such techniques but are sufficient to mimic many aspects of insect behaviour. Learning provides the means to escape this limitation. The hippocampus is the biological site for storage of cognitive maps in the mammalian brain.

Biological learning and memory: Learning is an essential pre-requisite for intelligent behaviour. To program a machine with 10^{14} bits at one line of code per hour (rate of production of software from conception to installation) with an average line of code some 500 bits long would require 100 million man-years. This may be compared with the information input rate of the human vision sensors of 250 Mbps which may input 10^{14} bits in 20y. Mammals have two different types of long-term memory which use two different brain structures for encoding information. Memory traces in the human brain are localised to the cerebellum, hippocampus, amygdala and cerebral cortex. Human memory comprises 3 major components: sensory memory, short term memory (STM) and long term memory (LTM). Short term memory (STM) is highly restricted in capacity being limited to 7 ± 2 separate items or chunks. A chunk is a re-ordering and re-integration of items into an organised groups to increase the capacity of STM. In SOAR chunking occurred in working memory whereas in ACT* it occurs in the production memory. Working (short-term) memory is associated with the prefontal cortex, particularly the principle sulcus, which accesses relevant stored information in the cerebral cortex – working memory provides moment-to-moment awareness and retrieves archived information (similar to a "blackboard"). Long-term memory storage occurs in the cerebral cortex which is also the region responsible for higher cognitive processes in mammals. The temporal lobes on the inner surface of the underside including the hippocampus are responsible for the storage of new long-term memories. Declarative knowledge comprises general and specific facts about the world while procedural knowledge defines how to perform actions in the world and mental strategies. Declarative knowledge involves association of simultaneous stimuli about single events whereas procedural knowledge involves the association of related sequential stimuli. Declarative memory is subdivided into episodic memory (experiential) and semantic memory (encyclopedic). Episodic memory is autobiographical and is self-referential while semantic memory is a mental lexicon stored without reference to personal events during which the facts were learned. The hippocampus is associated with semantic/episodic information storage. The hippocampus appears to be responsible for spatial memory of objects in the environment. It participates in the formation and retrieval of semantic memory and the cerebral cortex is associated with storage.

The addition of learning provides the ability to adaptively use the world model to predict and compare predicted and actual outcomes of courses of action to update the knowledge structures for further use. The world model allows planning of future behaviours through simulation – simulation is reversible unlike action in the real world. Reinforcement learning is based on punishment/reward signals to learn complex behaviours by favouring actions which generate the maximum numerical reward value produced the environment. Non-associative learning is associated with primitive response variation to single event stimuli such as habituation, dishabituation and sensitisation. Learning automata have the characteristic ability of improving their behaviour over time in a manner similar to neural networks and are modelled as stochastic systems with little or no a priori information. A random environment acts as a probabilistic teacher while learning essentially optimises an explicitly unknown functional. There are a number learning schemes. The linear reward-penalty scheme (L_{RP}) is optimal is all stationary environments. The linear reward-inaction (L_{RI}) scheme ignores penalty inputs from the environment so that action probabilities are unchanged in those cases (benevolent automaton). Genetic algorithms are similar to the L_{RI} learning procedures in that candidate solutions that survive into the next generation correspond to reinforcement and those that are discarded corresponded to zero reinforcement. L_{RI} procedure may also be used in artificial neural networks such that each connection is treated as a probabilistic switch. The symmetric L_{RP} scheme is ε-optimal in restricted random environments. The learning rates of non-linear models are faster than those for linear models and are ε-optimal in all stationary random environments. Associative learning is based on causal inference in deriving relationships between two or more phenomena – the basis of concept formation. There are two types of associative learning – Pavlovian (classical) conditioning and Thorndikian (operant) conditioning. The A_{RP} learning algorithm (associative reward-penalty) is an associative reinforcement (competitive) learning systems which treats each neuron of a neural network as a stochastic automaton. It is similar to the Robbin-Monroe gradient descent approximation procedure and is a generalisation of the perceptron algorithm and the Widrow-Hoff rule (adaline). It reduces to the non-associative L_{RP} algorithm under the appropriate conditions. These learning automata can function in stationary random environments only. A non-stationary environment will appear inconsistent because the automaton receiving the same input at two different times will generate the same actions which may generate different world states. The only way to overcome this is to make the algorithms efficient in time such that the automaton's time constant is shorter than that of its environment by using simplifying assumptions about the environment to maintain environmental regularity.

ANN's (artificial neural networks) are artificial analogues of neural architecture of the human brain. Connectionism is an holistic and nonsymbolic approach in that information representation is mapped directly into its architecture. Unlike the symbolic paradigm, the connectionist approach has an epistemological basis being founded on bottom-up processing through self-organisation. Neurons may be modelled as distributed RC networks with the cell cytoplasm providing a series resistance and the lipid cell membrane providing the shunt capacitance. An RC circuit sums the inputs according to the weighted synapses. However, this is probably an oversimplification as neurons have complex dynamics as multi-input devices with many non-linearities from in-out cross-couplings. The extended leaky integrate-and-fire neuron is a more sophisticated neuronal model. Output cells have graded analogue input with a digital all-or-nothing output. Interneurons however, have both graded analogue inputs and outputs – the graded potential output spreads out a short distance attenuating as it does so. Artificial neural networks (ANN) are based on electronic neurons arranged into regular layers connected by weighted links. Individual neurons react specifically to a given set of stimuli – this set represents the receptive field of the neuron. The information is effectively stored by the specificity of the neuron. The receptive field acts as a matched filter for those stimuli. A neural network thus learns to model input-output patterns. The commonest ANN architecture is the multilayer perceptron comprising of an input layer, a small number of hidden layers (usually one or two), and an output layer. Such ANN's can model continuous non-linear functions. The artificial neural network performs a weighted summation of its inputs and its output depends on that summation exceeding a threshold – it may be regarded as a form of Kalman filter. The error backpropagation algorithm is a generalisation of the least squares error procedure for single layer networks. Although there is no evidence that biological synapses in the brain use backpropagation of derivatives, reciprocated dendrite-dendrite synaptic connections do occur possibly for feedback. A variant of the backpropagation algorithm is the A_{R-P} procedure which although slower does not require a separate pass to backpropagate the error information.

It has long been known that conventional neural network models can exhibit chaotic behaviour by virtue of their non-linearity. Chaotic attractors in neural dynamics can store dynamic behaviours with several initial conditions allowing the system to randomly between solutions. Chaotic dynamics arises a result of non-linear feedback with time delays with neurons possessing more than one firing level to produce a range of possible behaviours, i.e. higher level cognitive processes such a generalisation and abstraction are not predictable from microscopic neural dynamics. This differs from the view of cognition as derivable from sets of finite state machines. Biological neural networks act as chaotic associative memories which provide robustness to noise. Many linear systems exhibit a fixed point attractor point to which the dynamic trajectory converges over time. In contrast, non-linear systems have no such fixed points but exhibit trajectories that never repeat. They usually comprise an infinite set of unstable periodic orbits which define the strange attractor. Memories are no longer associated with fixed points but rather with unstable periodic behaviours, in principle there are infinitely many potential memories. Chaotic dynamics arises as a result of non-linear feedback from re-entrant pathways which introduces time delays. The extreme sensitivity to initial neural input conditions displayed by chaotic neural networks makes them unstable and unpredictable but it also makes them susceptible to control by very small inputs. The OGY (Ott-Gregobi-Yorke) control method forces the chaotic system to stay near selected unstable periodic orbits. These techniques are not yet well developed but offer the potential for massive increase in neural memory storage without increasing neural network size.

Neural modelling has been particularly applied to motor and perceptual tasks rather than logical reasoning and knowledge organisation as they provide a better description of low level processing such as pattern recognition and of motor skills which are all highly parallel processes. This is particularly true of biological motor control which has particular application in robotics. Of particular interest are artificial neural network models of the cerebellum which store pre-learned of associative motor activation patterns, eg. CMAC (cerebellar model articulation computer). The cerebellum is primarily responsible for muscular coordination and regulation of motor skills – indeed, unlike the case of the cerebral cortex, each cerebellar hemisphere controls its own side of the body. The cerebellum essentially simulates muscular movement through feed-forward models of motor commands before feedback from the muscles is available. It provides the means for rapid movement through pure prediction – the biological analogue of the Smith Predictor. The forward predictor model simulates the effect of the commanded movement based on the body's current state (emulation). This model is supplemented by an inverse control model which represents the mapping of the motor commands to the muscular movements (control). The emulator "forward models" the system's behaviour to provide rapid error-correction (similar to model-referenced adaptive controllers or Kalman filter-based controllers). CMAC uses input-output measurements stored in a look-up memory to provide feed-forward control signals – the mapping generated through associative learning was used as feed-forward information to calculate the control signals. Parallel fibre inputs to the cerebellum code for the desired limb trajectory and proprioceptive feedback from the muscles. Climbing fibre inputs encode motor feedback commands which activate Purkinje cells which encode sensory error signals for the motor commands. The Purkinje cells utilise the long-term synaptic plasticity as a temporal averaging mechanism triggered by the stimulus from the climbing fibres.

Human reasoning is demonstrably associative in that it is content addressable and connectionist architectures are implicitly associative. Classical conditioning deals with unconditioned stimulus (UCS) that generates an unconditioned response (UCR), and a conditioned stimulus (CS) that is associated with a UCS. This CS-UCS relation creates a CR that involves the specific generation of the UCR by the CS. If some given CS precedes another UCS that creates a response, this CS will trigger the UCR. A neural model of classical conditioning has been implemented through connectionist strength adjustment. with the learning rule defining how the weights are updated using the Widrow-Hoff delta rule. Each time an action is selected by an agent, it changes the state of the environment which in turn emits an observable reinforcement signal which measures the agent's performance. This scheme corresponds to the Rescorla-Wagner model of classical conditioning. A variation modelled the prolonged trace latency effects of chemical concentration traces in neurons that persist after the CS ends and the US starts.

There are two main forms of symbolic learning: inductive learning and analytic learning with the latter including explanation based learning and analogical learning. Inductive learning involves generating a general concept representation from a set of positive and negative examples. Analytical techniques allow learning from a few examples and use deduction and background domain knowledge rather than induction. It involves extending existing knowledge to situations that have high similarity. Learning by analogy consists of tranforming existing knowledge for one domain for use in related domains. This method is most appropiate for increasing the efficiency of rule-based systems where reasoning and inference performs well in well structured environments. Explanation based learning (EBL) is an analytical and knowledge-intensive form of learning. It addresses the problem of generalising from single training examples (as opposed to inductive generalisation which requires many training examples). It can improve problem-solving performance from single examples in the context of background knowledge.

Genetic & neural approaches to learning: Algorithms based on biological analogies have become commonly adopted for solving difficult problems: genetic algorithms (and variants such as genetic programming and evolution strategies – collectively known as evolutionary algorithms) provide the basis for the evolutionary search of solutions to NP-hard problem spaces; neural networks (highly simplified idealisations of biological neural architectures) provide the basis for the generation of non-linear functions for robust pattern recognition and for control systems based on non-linear dynamics. Classifier systems are a machine learning architectures that employ a genetic algorithm where bit strings codify situation-response rules similar to an expert system and have been used to implement behaviour control through reinforcement learning – they are algorithmically equivalent to Q-learning. A neural net which requires precise connection schemes may be represented as a space of high fitness in the genetic search space. The **evolutionary robotics** approach is an automatic design procedure whereby a robot's controller, and possibly its overall body plan, is progressively adapted to the specific environment and the specific problems it is confronted with, through an artificial selection process that eliminates ill-behaving individuals in a population while favouring the reproduction of better-adapted competitors. It involves a population of genotypes (i.e. of information that evolves through successive generations) and a phenotype (i.e. the robot's control architecture, its body plan, and its behaviour) that is encoded in any one genotype. Evolved software controllers may be implemented as control programs (in a high level language or in machine code), as a variety of production-rule systems, or as neural networks. Behaviour based control methods may be evolved using genetic algorithm based classifier systems. The classifier system provided a mechanism for reinforcement learning based on the performance of robotic behaviour and generate feedback concerning the success or failure of actions. As hierarchical structures learn faster than single level distributed architectures, an hierarchical structure of behaviours is produced. The hierarchical structure has the lowest level classifier system learning basic behaviours from sensory data while those at a higher level learning to coordinate those behaviours with no direct access to the environment. This approach was extended and applied to a small robot AutonoMouse which had two eyes. Each classifier system of 100-500 rules was implemented on a distributed set of transputers as its brain, each with a specific behaviour arranged hierarchically (Alecsys).

Evolutionary neural networks apply biological analogies more strongly by hybridising the two approaches – the structure of the neural network is encapsulated into a genetic code and populations of these neural network structures evolve according to their success in providing appropriate control systems for robots. GA's can be used to search for optimal weights of an ANN to correctly classify training examples. Representation is characterised by four properties: (i) neuronal activation level; (ii) synaptic weights relating neurons in the network; (iii) dynamic trajectories of activation patterns across neurons; (iv) genetic representation of neural architecture. The ANN structure is responsible for generating functionally fit behaviour triggered by environmental input and it is determined by the adaptation of the synaptic weights. The development of the species adaptation algorithm (SAGA) provided the basis for using this technology in open-ended contexts by allowing variable length genomes. GA’s have been used to evolve neural network modules (GenNets) to construct hierarchical control systems. GenNet modules may be constructed into hierarchical structures to exhibit arbitrarily complex behaviour. Each GenNet with its specific behaviour may be evolved separately and then organised appropriately. Genetic Planner used GA-based artificial selection to breed computer programs that generate plans to purposively control a mobile robot. The genetic planner used crossover recombination and artificial selection through a fitness measure to breed computer programs. As the genetic planner did not reason about the world, it randomly generated plans and ran them until it came up with effective plans.

An extension of these approaches which offers great interest is incorporating the process of development of neural networks from genetic coding. This relies on encoding of the mapping between the genetic coding and the expression of that coding in the structure of the neural network. This emulates the process of embryonic development in providing information amplification from the genetic code to the connectivity matrix for a neural network. In the simplest case, the genetic algorithm codes a direct description of the circuit network wiring. In biological systems, the genome codes a more complex interpretative process through development into a phenotype. A number of possibilities for sensorimotor control are applicable, particularly using genomes to code for neural network topologies. The genotype may encode a graph generation grammar with each node of the tree being specific developmental genograms where the rewriting process is considered a developmental process. Alternatively, parameters to systems of ordinary differential equations, may be encoded into a genotype. Neurons are positional in a 2D plane and exert attraction forces on dendrites from other neurons. The dendrites grow in the force field generated by the distribution of neurons. If a dendrite moves within a given distance of a neuron, it establishes a connection to it. The most successful approach encoded the developmental process as a grammar tree controlling a cellular developmental process. In this approach the process starts with a single cell that divides, producing daughter cells that can in turn further divide. How much division occurs and how the cells connect together at each stage is controlled by the genotype. Cells inherit their connections from their ‘parents’ and no context sensitive development is possible. Alternatively, a developmental model based on Boolean networks can represent gene regulation networks to evolve autonomous agents through developmental processes. Recurrent dynamic neural networks have been developed using incremental evolutionary learning algorithms to provide the basis for visual navigation behaviours beyond simple behaviours. Recurrent neural network employ physical signal propagation delays and can model any class of dynamic system such as networks of augmented finite state machines. A range of control architectures and sensor layouts are represented as a population of genotypic chromosomes. These chromosomes interbreed with mutations according to a task-oriented fitness function which imposes a selection pressure on the population. In order to ensure generalized behaviour, a range of unstructured, dynamic environments should be used in the simulation. Cross-over is achieved through the SAGA (Species Adaptation Genetic Algorithm) principle which allows variable length cross-over between homologous genotype sections by minimising length changes. It is only through the use of increasing genotype lengths that the genotypic space can itself evolve beyond finite structures, i.e. a form of species evolution which allows local searches whilst retaining existing adaptations. Simulators are now widely used in real-world evolutionary robotics applications. Simulation has the potential advantage of speed as well as controllability and ease of access to data for analysis. Unless special generalized noisy simulations are used there will usually be a mismatch between simulated and real sensory data, so that real world noise must be included in the simulation. The approach is to evolve first in the imperfect simulation and then continue evolution in the real world. After initial evolution in the simulated environment, only a few additional generations are usually required to achieve successful behaviour in the real world – a kind of adaptive fine-tuning. A number of visually guided navigation behaviours have been successfully achieved including navigating around obstacles and discriminating between different objects. The use of minimal simulations which are ultra-fast by using crude models of carefully chosen sets of basic interactions between the robot and its environment have enabled speed up with vision-based systems. These employ multi-layered structured noise on all modelled aspects which means there are no stable simulation 'features' that an evolved robot could come to rely on (and therefore fail when transferred to the real world). Highly robust generalized behaviours are forced by this technique as they are the only ones that can score reliably. Neural networks for controlling locomotion in an artificial insect have been evolved in simulation and then successfully downloaded on a real 6-legged robot. A genetic algorithm has been used to derive the optimal gait parameters for an 8-legged robot. The controllers evolved have been capable of deriving walking gaits that are suitably adapted to a wide range of terrains, damage or system failures. Using a developmental approach called Cellular Encoding which genetically encodes a grammar-tree program that controls the division of cells growing into a discrete-time dynamical recurrent neural network, single-leg neural controllers have been evolved for a walking robot which generated a smooth and fast quadrapod locomotion gait.

GasNets incorporate virtual diffusing gaseous neuromodulators and are used as artificial nervous systems for mobile autonomous robots. The basic GasNet networks used in many recent experiments are discrete time step dynamical systems built from units connected together by links that can be excitatory or inhibitory. In addition to this underlying network in which positive and negative `signals' flow between units, an abstract process loosely analogous to the diffusion of gaseous modulators is at play. Some units can emit virtual 'gases' which diffuse and are capable of modulating the behaviour of other units by changing their transfer functions in ways described in detail later. For mathematical convenience there are two gases, one whose modulatory effect is to increase the transfer function gain parameter and one whose modulatory effect is to decrease it. This is loosely analogous to the length constant of the natural diffusion of NO, related to its rate of decay through chemical interaction. As well as network size and topology, and all the parameters controlling virtual gas diffusion and modulation, the robot visual morphology, i.e. the way in which the camera image was sampled, was also under unconstrained genetic control. Controllers based on networks with the virtual gas diffusion and modulation mechanisms turned on evolve significantly faster than those that are identical in every respect (including genotypes and all the evolutionary machinery) except that the gas mechanisms are rendered inoperative. Nearly all the successful GasNet controllers that were examined in detail exhibited surprisingly simple structures. A number of interesting sub-networks, such as oscillators making use of spatial aspects of the modulation and diffusion processes were independently evolved in several runs, suggesting that they are easily found building blocks.

The evolution of plastic neural networks provides the basis for networks that adapt during their operation according to rules that allow their connection strengths to changes depending on activity in the network. Evolved plastic controllers has some similarities with the GasNets approach but the form of plasticity is quite different. The conventional approach involves evolving fixed, genetically-determined properties of a robot controller, for example the synapse strengths and signs (excitatory or inhibitory) and neuron properties of an ANN. These limitations of this approach are overcome by evolving ANN robot controllers that can continuously modify their synaptic strengths using plasticity rules encoded in the genotype. The genotype represents Hebbian learning rules that determine how the synapses are changed at each time step. Evolution determines which variant of four Hebbian learning rules (plain, postsynaptic, presynaptic, covariance) governs each synapse in the ANN as well as its associated learning rate and sign. In each of these 4 variants of the Hebbian rules the change in synaptic change is completely determined by the activation of the pre- and postsynaptic neurons. The plain Hebb rule (1) only allows synaptic strengthening, with the greatest increase when both the pre- and post-synaptic neurons are at their maximum activation. The presynaptic rule (2) enables strengthening and also weakening of synaptic strength, when the presynaptic unit is active and the postsynaptic one is not. Conversely, the postsynaptic rule (3) causes a weakening of synaptic strength if the postsynaptic neuron is active and the presynaptic one is not, increasing the synaptic strength if both the pre- and postsynaptic neurons are activated The covariance rule (4) increases or decreases the strength of a synapse relative to the difference between the pre-and postsynaptic neurons. If the difference in activation is greater than half of the maximum activation level then the synaptic strength is reduced in proportion to this difference, otherwise it is increased in proportion to the difference. In the adaptive synapse approach, synaptic strengths are not encoded, rather they are initially set to small random values, so individuals have to learn how to solve tasks from scratch, changing their synaptic strengths through their interactions with the environment. It was applied to a simple navigation task where the goal was to avoid obstacles and move in as straight a line as possible within a looping maze. The adaptive synapse approach took about half the evolutionary time to solve the task compared to a genetically encoded network. Although the synaptic strengths continuously changed while the robot was moving, the overall pattern of change became stable. The robot controllers with adaptive synapses were able to adapt to new environments, in contrast to the hardwired controllers.

Biomimetic sensing: Sensors are a necessary component to any closed loop control system. They provide feedback data on the status of the environment. In particular, biological organisms are characterised by distributed sensing capabilities. There are five basic sensory modalities – vision, hearing, smell, taste, and touch. Hair-like mechanical sensors may be based on pressure sensing piezoelectric cantilevers – such sensors are ubiquitous in the biological world. Perception is fundamentally a categorisation process and relies on specific input features (cues) to perform such inferences of identification.

Vibration sensing through the use of hair-like mechanical sensors (which sense pressure through cantilever operation) are ubiquitous in the biological world. Vibration sensing, or more commonly, hearing is generally achieved through mechanical transduction receptors in the form of hair cells. Bending of the hair cell generates an action potential. The cochlear of the inner ear is a transducer which converts mechanical vibration into an electrical output. The vestibular system bears similarities with the acoustic system, particularly the transduction mechanism. Contact and bending sensors are to be mounted on the lobster robot's flexible antennae in order to provide it with a sense of touch. The MEMS flow sensors will allow the lobster to sense water currents and thus to adapt its behaviour accordingly.

The primary sensory modality for most animals which provides distance data is typically vision. Biological visual processing is based on feature detection involving three different regions of the brain: the retina, the lateral geniculate nucleus and the visual cortex. An object is recognised by the simultaneous appearance of a number of visual features in their correct relative spatial positions. Feature selection involves filtering an image to enhance relevant features and remove irrelevant ones. For example, convolving an image with a Difference of Gaussians (DoG) operator can enhance edges. In biological systems the receptive fields of cells respond only to some properties of the environment; for example, centre-surround antagonist neural organizations are found in the periphery of both vertebrate and invertebrate vision systems and they only respond to edges of particular orientation. Gabor filters in conjunction with centre-surround and corner-sensitive filters can provide such functions. Insects generally possess compound eyes with multiple receptors of constant, low resolution. They employ optic flow for vision with high frequency resolution which is computationally much simpler than mammalian vision. Their eyes are static and do not move relative to the body as they are too close for stereo-imaging. Optical flow field methods provide the means for motion-detection. Optical flow methods generate a vector field representing the motion of the pixels of an image. Associated with each pixel is a velocity vector and the vector field of all pixels of the image form the optical flowfield. Pre-smoothing the images with a Gaussian filter to reduce aliasing introiduces blurring. Active vision can exploit simple features to perform complex shape discrimination. active (foveated) vision which emulates foveated vision with automatic pan and tilt feedback control to provide gaze stabilisation, smooth pursuit, optokinetic nystagmus (OKN) and vestibular-ocular reflexes (VOR). Such foveated vision which reduces the image processing requirements of robotic vision systems. Control of foveated visual gaze through the vestibular-ocular reflex (VOR) reduces the requirements for image processing. The vestibular-ocular reflex (VOR) allows tracking of moving objects with smooth pursuit. Gaze shifting and gaze holding movements of the eye provide a means to rapidly aim this narrow FOV. Fast saccadic eye movements move the fovea between different targets in the visual field while active gaze control keeps the fovea on a specific target. The vestibulo-ocular reflex is a feedforward mechanism which uses the vestibular apparatus of the inner ear to detect movement of the head which responds much faster than the optokinetic feedback loop. The Artificial Retina consists of 9 square cells which move over and explore a visual image in order to perform shape discrimination. Initially the Artificial Retina starts with a low resolution covering most of the image which then shrinks when it moves around. The active vision systems develop sensitivity to position specific oriented edges. Co-evolving active vision and feature selection enables individuals to solve position and size invariant tasks using position and size variant mechanisms.

Touch is closely related to proprioception in some respects. Sensory feedback is provided by the proprioceptive sensors in the muscles, joints and skin while exteroceptive sensors provide direct observation feedback (vision, etc). Proprioceptive awareness of the body is central to controlling the agent in relation to the outside world. Touch is associated with mechanical stress or pressure but is a highly complex sense in its most general form as it is often associated with sensitivity to temperature, pain, moisture, etc. Touch is a distributed sensing modality which imposes the requirement for significant local processing. Touch is an active, exploratory process in which sensing and actuation are intimately linked such that the actuation process through movement of the hand and fingers define the sensory data directly. Tactile sensing is particularly important in grasping manipulation through modification of impedance of the hand. Grasping is a fundamental component of manipulation involving interaction with a high variation in impedances from zero (prior to contact) to high impedance if the payload is stiff. Internal models are likely to become complex in certain tasks which involve rapidly changing motion and/or forces. In such situations, low impedance is usually adopted so that the limb can "give" to unexpected disturbances so that large forces are not generated. Once a greater area of contact is obtained, greater forces can be exerted by stiffening impedance of the grip. Preflexes maintain stability while minimising the effects of disturbances. Preflexes are often implemented – preflexes are near zero-order responses due to the intrinsic, non-linear properties of the structure whereby disturbances are rejected within ~20-70ms without the time delays ~200-500ms inherent in feedback control from proprioception. Biological touch sensors are often hair cells, stress sensors (campaniform sensors in insects which reside at the base of wings) or vibration sensors (chordotonal sensors in insects analogous to hearing sensors). The primary sensitivity to touch is pressure for which there are several alternatives. Most force sensors are thin film strain gauges, piezoelectric sensors or capacitative array sensors. Carbon-impregnated polymers offer variable resistance depending on applied mechanical pressure and are used as tactile sensing “skins” in robotics.

Chemical sensing is the most fundamental of biological sensing modalities, common to almost all life forms, and is exhibited by single celled microbes for finding food (attractant) and avoiding poisons (repellent). Odour in particular is such a sensing modality. There is also the ubiquitous scientific sensor onboard planetary spacecraft which attempts to determine the elementary composition of planetary atmospheres and soil, eg. the gas chromatograph-mass spectrometer. Taste and smell are chemical senses – taste is in fact augmented by smell indicating the partial fusion of these senses in the brain. Chemiresistor is the simplest of chemical sensors in which the resistance of a thin film material may be measured. Metal oxides have limited utility in terms of their specificity, high temperature operation up to ~500°C, and requirement for ambient O_2. Conducting polymers may be used as chemical sensors – intrinsic conducting polymers are conductive without doping and extrinsic conductors are made conductive by impregnating with carbon black. Extrinsic conducting polymers are more versatile due to the wide range of polymers that can be used and they are almost entirely reversible. Chemical sensors made from different conductivity polymers embedded with carbon black particles provide an array of large bandwidth. Gas absorption causes swelling of the polymer which lowers the conductivity. They are particularly suited to deployment as polymer coatings for SAW sensors. Chemosensory cells of high sensitivity are the basis for the electronic nose. Given that most odours comprise an array of chemicals, the pattern of activation allows discrimination of a wide variety of smells. Electronic noses possess arrays of sensors, each sensor producing an electronic output as a function of sensor-specific chemical interactions. The combined electronic signals from the array give a characteristic pattern which can be correlated with a given odour or set of volatile compounds. The strength of the signal reflects the concentrations of each set of compounds. One difficulty is the degree of selectivity which suggests that a range of relatively unselective sensors should be employed with different responses to a range of gases – such an approach would require the use of neural networks trained to deconvolute the data. This makes such sensors suited to neural network processing and classification. The most traditional transducer in electronic noses is the gas-sensitive field effect transistor in which sensor arrays are used for pattern recognition. The selectivity and sensitivity are controlled by the composition and thickness of the gate metal, and the temperature of operation. The detected molecule causes a shift in the voltage characteristics of the transistor due to electric polarisation in the metal-semiconductor interface. Ion-selective FETs (ISFET) are commonly used for biosensors in which an insulating semiconductor oxide layer on Si substrate results in an inversion layer in the underlying substrate. Equivalent RC circuits formally describe the frequency behaviour of such devices. There are two types of acoustic wave sensor which are sensitive to mass – the quartz crystal; microbalance (QCM) and the surface acoustic wave (SAW) detector. Both devices are based on the generation of acoustic waves in piezoelectric materials by the application of metal electrodes either side of the crystal. A change in mass causes a shift in the frequency/velocity of the acoustic waves. The application of thin film coatings turn such devices into chemical sensors. Bulk acoustic wave (BAW) devices have been proposed but surface acoustic wave (SAW) devices offer much greater flexibility. SAW devices are constructed from a piezoelectric substrate – typically $LiTaO_3$ or $LiNbO_3$ - with metallised electronics including reflector gratings implanted on the surface. The reflector gratings act as mirrors for the SAW. Operating the SAW at high frequencies increases its sensitivity. Covering the wafer with imprinted polymer will ensure that the mass will change when target molecules are present. The detection of such mass changes also occurs in biological systems – this provides the basis for the recognition of antibodies by epitomes of antigens. Two types of SAW devices are used as chemical sensors – the delay line and the resonator. Polymer films are used primarily to detect organic gases but tend to be of limited specificity, so several sensors with different coatings are used to detect specific substances. An electronic nose based on a micromechanical array of Si cantilevers has been developed in which the cantilevers are sensitised with a number of different coatings such as metals, monolayers and polymers. An array of such devices with individually coated cantilevers provides the basis for a wide range of sensitivities. A cantilever coating of zeolite ZSM-5 which absorbs water selectively provides the basis for a water sensor.

A distributed and diverse sensor suite enables a system to cope with failures gracefully due to functional redundancy. They can also provide multiple information sources which operate in parallel. Each sensor output is degraded by noise so if multiple sensors are used to determine the same property, uncertainty is reduced. Furthermore one modality may suffer degradation under certain conditions but by using additional modalities such problems can be eliminated. Multisensor fusion is a natural extension of using multiple sensors for verification. For complex environments, it is desirable to simultaneously integrate such sensory data from diverse multiple sources. Multisensor data fusion improves the overall performance of interpretation and classification by providing a more reliable and robust estimation of the environment than each sensor considered alone. Sensor fusion can occur at different levels from the integration of raw data to the integration of object properties of a world model. Fusion essentially involves the integration of different sensory modalities into a single representational format such as a world model such as a blackboard. For similar modalities, redundant information may be fused at a lower level of representation, eg.through Bayesian or Kalman filters.

<u>High-level cognition</u>: Classical artificial intelligence (AI) methods are based on emulating higher human cognitive functions (physical symbolic systems hypothesis). There are three types of knowledge acquisition: rationalism – knowledge that is derived internally through reasoning; nativism – knowledge that is innate (e.g. Chomskian language organ); empiricism – knowledge learned through experience. The physical symbol hypothesis states that all aspects of human cognition are the product of information-encoding symbol manipulation within the brain. This paradigm is not open to empirical test in the Popperian sense, but is based on a doctrine of functionalism such that the physical medium be it brain or computer for such information processing is irrelevant to the processing mechanism itself. Cognition and thought therefore is synonymous with computation. A central issue to such a paradigm is the role of logic: programs are based on proof generation methods for symbolic logic theorems as a search for solutions to problems. A formal logical system is a set of abstract symbols together with rules for concatenating these symbols into strings. Formal logic is defined syntactically and semantics is defined by the interpretation of the symbols outside the formal system. A formal system is consistent if no contradictory statements can be proven. Truth is thus synonymous with provability. This approach is typified by symbol-manipulation architectures such as SOAR and ACT*. Both of these systems are universal computational architectures which attempt to model mental human intelligence. There is a standard form of first order predicate logic (Horn clause logic) that limits the conclusion of an implication to at most a single atom. Automated reasoning for the discovery of concepts employs theorem proving based on formal first order predicate logic. Horn clause logic is computationally equivalent to the universal Turing machine. Inferencing is achieved through Robinson’s resolution principle. Logic is an inefficient means for modelling intelligent behaviour as logic cannot tolerate contradiction (and much human behaviour is contradictory).

Plan generation is basically the process of specifying the ordering of actions to achieve desired goals. Plan generation is very much part of the control problem in that it generates the input to the control system as a sequence of commands to achieve the desired goals and monitors the plan execution for feedback. Plan generation has been traditionally viewed as the generalisation of the programming language process in that the plan represents a possible solution to a problem through the execution of an effective well-defined procedure, i.e. planning is an algorithmic problem solving process. Plan generation as generalised problem solving may be characterised as a state space search for actions which when applied to the outside world change it into a new goal state. STRIPS has provided the basis for many subsequent planners and it has a central role in planning systems. The environment is constantly changing and reasoning in the real world must reflect this. STRIPS and similar planning programs suffer from the "frame" problem which arises out of the dynamic nature of the world. The frame problem arises from the need to represent those aspects of the world which remain invariant during state changes. The difficulty is in keeping track of the consequences of performing actions and the required alteration to the world model representation in real time. As the world model becomes more complex so the computation time required to update it explodes exponentially. Since successive states of the world are logically independent, the difficulty is in deciding which statements continue to be true and persist after the performance of a given action. Indeed, generally almost all the assertions will continue to be true after the action has been applied as most of the world is unaffected by events as they tend to make only localised changes to the world. Newell's SOAR (Smalltalk-On-A-RISC) is an extension and generalisation of the GPS architecture system for general problem solving. It is an approach to derive a unified computational architecture for general intelligence particularly for goal oriented problem solving. SOAR is firmly rooted in the Physical Symbol Hypothesis. It is essentially a production rule based expert system which uses a heuristically guided A* search to find operators in the problem space that can transform the current state to a desired goal state through the generation of hierarchies of subgoals to resolve impasses that occur when knowledge is incomplete. The chief problem of STRIPS and all its variants is that they are capable of dealing with only simple planning problems. Problems such as the Tower of Hanoi problem which require overcoming obstacles to goal reduction are not soluble by these methods based on reducing searches.

Many cognitive functions such as search (such as the A* and D* algorithms) invoke NP hard type problems which cannot be readily solved in real-time. The Knowledge Principle states that search and reasoning alone is regarded as insufficient for intelligent behaviour – a great deal of knowledge of the world in which the agent operates is required. Much of this knowledge should be sufficiently general to enable its use for multiple but specific domains (the Breadth Principle). The knowledge base loosely represents stable long term memory (LTM) and they are limited to specific problem domains. Real world data in a dynamic global working memory database (loosely correlating to short term memory (STM)) must be matched against knowledge stored in the knowledge base. Events comprise the inputs to the working memory which is acted upon by an inference engine. Production rules represent relatively independent chunks of knowledge and are used by the inference engine to infer solutions to problems. Production rules define relationships between the rule's terms through implication. These rules consist of one or more antecedent precondition statements linked to one or more consequent action statements (condition-action rules). They essentially model plausible hypothesis-conclusion inferencing logic for problem solving through deduction. Production rules often have numerical weights assigned to them to quantify the level of belief in the consequent given that the antecedent is true. The certainty factor is a measure of the association between premises and actions to model uncertainty. The basis of determining uncertainty is probability theory either through Bayesian statistics (usually) or Dempster-Shafer theory of plausible inference (rarely as it is complex to compute) to give a truth value as a probability on the real unit interval [0,1]. Alternatively, fuzzy logic may be used based on possibility theory. There are two main forms of knowledge representation commonly used in expert systems, frames and production rules, of which the latter are the most popular. Frames are schematic, modular data structures for representing generic concepts concerning objects, classes of objects and situations. Frames provide a natural representation for declarative knowledge (knowledge about something) while production rules are usually used for procedural knowledge (knowledge about how to do something). Production rules can be attached to frame slots as values to be invoked as demons when the frame is invoked. This ability to attach rules or rule classes as demons to frame slots provides great flexibility in controlling reasoning particularly for reactive functions.

Semantic networks are one form of knowledge representation which comprise directed graphs whose nodes represent semantic entities and whose links represent relations between those entities. Semantic networks are directed labelled graphs of nodes are connected by binary relations with certain strength values measuring the relevance relations between the nodes. Each object is associated with other objects interconnected into a tree structure. They are similar in architecture to relational databases and form a dependency network of concepts. They define a relational logic through network links. Semantics is viewed as a contextual property of the concepts in the semantic network. Each node represents a concept and the concepts are represented as semantic attributes and stand in terms of their specific relationship to other concepts as represented by links. There are two main types of link between concept nodes adopted in such networks but both are "inheritance" links (superc) to provide a categorisation and classification capability: "is-a" and "is-a-kind-of" link. These two links provide the semantic inheritance network with the representation of both generic-generic relations of subtypes and generic-individual relations. General-purpose expert systems available include KEE (Knowledge Engineering Environment) and ART (Automated Reasoning Tool). Both KEE and ART support frame representations. KL-1 also offers a hybrid production rule/frame representation. The situation parts of all production rules are organised into a structured taxonomy of all the situations and objects in the knowledge base. A collection of concepts are linked together by "is a" links allowing the inheritance of properties between concepts. The action parts of the rules may be attached to the nodes in the structure as pieces of "advice" that apply to the situations described by the rules. NETL, a coarse coded distributed neural network/semantic network whereby each processing element is tuned to each concept or feature. The nodes represented symbolic noun-like concepts and links represented the relationships. The specific inferencing link adopted was the inheritance "is a" link handled by simultaneous marker propagation through the network and its parallel set intersection capability enabled recognition tasks to be implemented.

Logic in AI imposes a variety of problems, almost all of which involve the frame problem in one aspect or another. Logic deals with eternal truths, and all extensions to logic are attempts to cope with causality and change in the real world. Indeed, this is the fundamental requirement in robotic planning. Logic is descriptively universal and requires complete, certain and precise information. First order logic can handle only eternal truths. However, the world is dynamic and constantly undergoing change. Logic lacks the flexibility to cope with a changing world. Logics other than the first order type have been developed to attempt to cope with some of the limitations of first order logic: default logic, modal logic and temporal logic. Non-monotonic logics attempt to circumvent the rigidity of logic as an approximation to intelligent reasoning while retaining the symbol processing paradigm. Default logic augments the normal deductive axioms with default rules to infer conclusions that cannot be deduced from the axioms. Such conclusions are consistent with the real world but yet not verified and may be rescinded in the advent of further information. Modal logic is based on the notions of necessity and possibility through the use of modal operators to qualify truth through such expressions as “it is possible that…” and “it is necessary that…” Temporal logics enable reasoning about changes that occur in sequences of events. It explicitly represents the time dimension to enable reference to past events and rates of change. It reasons from past events to the truth of facts in the future. The past is a linear sequence of time points up to the present (now). The future is then a branching structure of time points from the present reference point to many possible worlds. The situation calculus may alleviate the problem of causality. For dealing with causality, situation variables may be introduced as extra arguments to logic variables. These arguments denote situations as time instants – time(s). The situation represents the complete state of the universe at an instant of time. Relationships between situations (including the consequences of actions) is represented by a result function independently of the problem. This leads to the concept of non-monotonicity which is closely related to belief, certainty and default reasoning. Human common sense relies heavily on the ability to make decisions in an uncertain world by using default reasoning to compensate for incomplete specifications. Non-monotonicity provides a way of expressing a fact that is true unless there is an abnormal circumstance or exception to the general rule, eg. birds can fly (unless it is a penguin). Axioms are required to state which features of the situation are unchanged by actions and this may be accomplished through circumscription, eg. objects do not normally change their locations. In these cases beliefs are held to be true in the absence of information to the contrary. Circumscription yields conjectures (jumping to conclusions) even to the point of allowing inconsistency (to which humans are prone) through over-generalisation of object classes. Default logics are a special case of generalised non-monotonic logics but remain in the framework of first order logic.

Truth maintenance is an important issue in large knowledge bases so that data which has changed should be updated so that the knowledge base remains consistent with the rest of the data in the knowledge base and with environmental data. Conflicting combinations of multiple assumptions can arise in Truth Maintenance Systems which depend on the order in which default rules are applied. This is fundamental to the frame problem. Truth maintenance deals with inconsistent knowledge by recording and maintaining reasons for beliefs, i.e.. truth maintenance which involves maintaining a database of beliefs that are consistent with propositions believed to be true unless there is evidence to the contrary. In TMS (truth maintenance system), each statement records an associated set of justifications (representing inferences), each of which represents a reason for holding the statement as a belief. These justifications function similarly to defaults and act as constraints. Justifications for a belief comprise of other beliefs. The Truth Maintenance System (TMS) provides the means for belief revision when discoveries contradict assumptions.

Few AI systems to date employ affect (emotion and motivation) in their design, though recent activities have suggested the inclusion of affect for value assignation (which is also important for learning). The James-Lange theory of emotion suggests that emotional experience is the product of cognitive processes indicated by body (visual) responses concerned with the maintenance of homeostasis. There is little doubt that the motivational portion of human cognition is an important complement to the rational component. Emotions are an essential component of higher level cognition to enable rapid decision-making. It provides greater flexibility than purely reflexive behaviours. Emotions are a major determinant of motivational behaviour: much behaviour may be characterised as pain aversion or pleasure-seeking. It is the motivational component which produces goals and desires to be achieved and generates emotions as an alerting device. The intensities of a motive determines the intentions of the agent. Emotions provide estimates of biologically valuable state variables and play a central role in behaviour to prolong pleasure (reward) and avoid pain (punishment). Intention maintains a rational balance between beliefs, goals, plans and commitments. The motivational state of an animal or human depends on its internal state, its perception of the external world, and its current behaviour and expected effects on the world. Motivational emotions allow specific activities to be selected in accordance with relevant goal stimuli. It is not quite clear how many primary biological emotions there are but suggestions include – fear, anger, depression/sadness, surprise, happiness/satisfaction and disgust. The CogAff architecture attempts to model emotions as the result of a control hierarchy – the first layer is reactive, the second deliberative and the third meta-management.

Hybrid Approaches to Intelligent Control: Reflexive behaviour-based robots lack explicit goals and goal manipulation. If a task requires knowledge about the world and requires reasoning and memory rather than just perception and survival, then reactive behaviour in insufficient. An autonomous system requires both goal-oriented planning and robust reaction. Plans may be viewed as resources for advice rather than as constraints on predetermined courses of action. Internalised plans were represented as gradient fields computed from graph searches to minimise the abstraction requirement. The map representation may be based on discovered landmarks distributed over a subsumption architecture of behaviours. The map may be encoded as a topological graph with each node representing a landmark and the links representing adjacency. Each landmark represents a behavioural/action module. Action modules with precondition list, action effect list and action effect delete list (similar to STRIPS) may linked into a network of successor, predessor and conflictor links which act as speading activation conduits. The interaction dynamics between action-oriented modules establish the sequence of actions in a distributed manner in response to environmental conditions and global goals.

There have been a number of approaches of attempting to combine the advantages of symbolic techniques with neural approaches. A rule-based execution monitor teaches a CMAC (cerebellar model architecure controller) network how to accomplish a task by observing the rule-based task execution. The knowledge-base represents declarative knowledge with inferencing abilities while the neural network represents associative procedural knowledge with pattern matching abilities. hybrid techniques have been proposed to augment fuzzy logic with other soft computing paradigms such as artificial neural networks (ANN), genetic algorithms (GA), and genetic programming (GP) to achieve the level of intelligence required of complex robotic systems. The fuzzy-ANN controller takes advantage of ANN for handling complex data patterns and learning from experience, the fuzzy-GA controller utilizes the ability of GA to optimize parameters of membership functions for improved system response, and the fuzzy-GP controllers utilizes the symbolic manipulation capability of GP to evolve fuzzy control rules. There is recognition of the shared characteristics between fuzzy and neural systems. To take advantage of the representational capacity of fuzzy knowledge-based systems, and the learning ability of ANN, the structure of a fuzzy-ANN system should take the form of a fuzzified connectionist network Under certain mild conditions, Radial Basis Function neural networks and fuzzy systems are equivalent. RBF network can be viewed as a mechanism for representing rule-based fuzzy knowledge by using its localized network structure, and performing associated fuzzy reasoning using feedforward computational algorithms. RBF network is essentially a network representation of if-then rules, with each hidden unit representing a rule and the basis functions being equivalent to the membership functions. The simplest and most straightforward attempt in the merging of fuzzy and neural controller techniques is to make the ANN learn input--output characteristics of a fuzzy controller. A conventional fuzzy controller is used to generate the data for training the ANN. The fuzzy neural networks (FNN) approach involves an ANN whose nodes have 'localized fields' which can be compared with fuzzy rules and whose connection weights represent input or output membership functions to generate fuzzy inferencing through the network. The membership functions and sets of rules are constructed from example data using multi-step procedures that involve learning the membership functions and weights, forming rule representations and constructing computational networks. Most learning of the membership function and the weights, is done using the back-propagation method and least squared methods. Integrating fuzzy-GA applications involves linguistic rules and their membership functions which are adjusted simultaneously. The membership functions are adjusted individually for each rule. A fuzzy classifier system learns by creating fuzzy rules which relate input variables to internal or output variables and, through performance assessment, credit assignment and reinforcement is able to achieve online improvement. The symbolic data processing done by the genetic program makes it particularly amenable to automatic evolution of fuzzy rules, which are comprised of symbols representing fuzzy sets and fuzzy inference. The use of GP allows direct manipulation of the actual linguistic rule representation of fuzzy rule based systems. GP has been used to evolve fuzzy rule bases that coordinate multiple fuzzy behaviours arranged in a hierarchical structure employed for autonomous control of mobile robot. The use of GA in fuzzy-ANN control has also been shown to enhance the performance of the system by overcoming some common bottlenecks, especially those requiring human intervention. GA is powerful for offline synthesis of optimal structures and parameters of FNN.

Group behaviour: Distributed information processing is fundamental and ubiquitous in natural and social systems as a form of complexity reduction through decomposition into modular chunks while providing redundancy for graceful degradation in the event of component failures. The vast majority of multi-robot research addresses the question of how best to design a group of robots that work together to perform some task or achieve some global objective. Multiple robots provides robustness through graceful degradation, parallel actions, distributed sensing and action, and cost effectiveness. The central problem is that of coordinating the behaviour of individual robots in order to achieve effective, coherent system-level behaviour. Wasps, ants and termites are social insects. The majority of insects in an insect colony are sterile workers which never have their own offspring. However, all insects in the colony descend from the same mother. Each unit is a simple processor, but through the interaction of many such units, complex behaviours can emerge, eg. the construction of termite mounds. The most appropriate biological analogy for distributed intelligence is the termite colony. Coordinated through pheromones, self-organising emerges as complex higher-order behaviour. Colonies of "virtual" ants (vants) have been simulated with global societal behaviour emerging from local interactions between individuals simulating pheromone trails. Each vant's behaviour is determined by its environment, but the environment is also the result of its past behaviour.

Problems requiring coordination arise when robots are engaged in activities that are interdependent. Typical group tasks include foraging, group movement, cooperative manipulation and team games. Foraging tasks are of interest because they present a scenario analogous to that which would be encountered in planetary exploration. Foraging tasks rarely require close coordination, as robots are primarily engaged in individually performing sub-tasks which relatively independent. Group movement tasks require robots to move together as a group whilst remaining within relative close proximity, eg. flocking. Alternatively, a specific formation movement task may be required. Cooperative manipulation tasks are tasks that require two or more robots to transport, reposition or otherwise manipulate an object. The commonest team game adopted for robots is robot soccer. Teams may be homogeneous (identical agents) or heterogeneous (different agents): heterogeneity affords division of labour while homogeneity affords robustness. Four main types of organisation are typically distinguished, these are: *centralised*, *distributed*, *hierarchical* and *hybrid* control systems. If control is centralised, a single agent is responsible for controlling all robots. Distributed control has been the dominant control paradigm in multi-robot systems research, but hierarchical control represent a hybrid approach. The main advantage of a distributed control system lies in its potential for robustness and fault tolerance. Global control which exerts its influence beyond the local agent is a necessary requirement of a centralised control architecture. Local control restricted to the local environment as emphasised in distributed approaches has greater potential for robustness and fault-tolerance. The most common task type addressed through co-evolutionary methods has been group movement (i.e., variations on 'flocking'). Generally, different (homogeneous) robots developed functionally distinct role contributions either to group formation, or to group movement. Non-cooperative group behaviours have generally been of the predator-prey type and their co-evolved dynamics. Certain parameters of each robot's control system are denoted as 'genetic material' and a set of such parameters comprise a genotype. An evolutionary population of genotypes is distributed across the physical robots. Genetic material is periodically exchanged between robots over communication channels.

Sociality is the basis of intelligence in humans - sociality complicates the environment, so that survival depends on predicting the behaviour of others in the social group. Language evolved as a form of social manipulation. In comparison to computer programming languages with its basis in logic, human language is highly rich. All human language according to the Chomskian paradigm comprises a three layered structure – syntax which determines the rules for the construction of language structure, semantics which maps the vocabulary to objects (nouns) and events (verbs) of the world, and pragmatics which determines the function of linguistic communication in altering the behaviour of the listener. The principle of social rationality maximises the joint benefit to its agent members and the society of agents – if an agent member of an effective society can perform an action whose joint benefit is greater than its joint loss, it may select that action. Joint benefit represents a balance between the individual agent and the social community of agents. Members indulge in activities for which their member benefit is less than their member loss if the member benefit if society gains more in total than it loses. Members indulge in activities which provide personal benefit but which are detrimental to society if the member benefit is greater than the social loss. This places a limit on the degree of cooperation – a member will not indulge in activity which brings a personal loss greater than the benefit accrued to society.

Distributed artificial intelligence (DAI) is based on the concept of organising a heterarchical team of interacting decision makers who indulge in coordinated activity operating under definite strategies to generate and develop solutions to problems which cannot be solved by a single individual. DAI systems are differentiated according to their control methods of organisation and their communications protocols. The amount of cooperation between agents can vary from fully cooperative (with high communication cost) to non-cooperative (no communication cost). Cooperation reduces control uncertainty and can coordinate multiple expert systems with different knowledge domains. Protocols keep communication between components to a minimum which may vary from a shared global memory (blackboard) in a hierarchical network structure to message passing whereby agents as individual active entities pass messages between each other to communicate. Synchronisation may be achieved through a protocol such as the contract net protocol. Small groups of agents may combine and abstract their activities as teams which in turn may combine at various levels of abstraction to form a functional hierarchy headed by a single entity. The paradigm for communication may be either through a shared global memory, message passing between agents or a combination. Using both overcomes the disadvantages of each and utilises the advantages of each, i.e.. using shared memory for local processes with message passing between groups of agents. The shared global memory is exemplified by the "blackboard" model which agents post messages, partial results and requests. The message-based contract net protocol (CNP) is particularly suited to a distributed set of heterogeneous specialist knowledge based agents. They cooperated by mutual sharing of information by exchange between loosely coupled agents, evaluation and mutual agreement. Negotiation provides the communications protocol mechanism to structure agent interactions so that relevant knowledge sources exchange information and come to agreement. In the contract net protocol, a hierarchical organisation is produced dynamically through a contracting and subcontracting process. The functionally accurate/cooperative (FA/C) approach involves asynchronous, opportunistic knowledge sources cooperatively exchanging and integrating partial and tentative high level results which may be incomplete or inconsistent being based on local information in pursuit of complete global solutions in bottom up fashion. The heart of the FA/C approach lies in the consistency checking of high level exchanged results to enhance global coherence. Inconsistent partial solutions are discarded and only consistent solutions are searched for. This bears resemblance to the scientific community metaphor with reference to paradigm-consistent theories and the generation of alternate concurrent partial plan bears resemblance to rival scientific theories. Both the FA/C and the CNP approaches are mutually consistent and may be combined. The blackboard can provide broadcast communications with directed message traffic between nodes provided by the CNP. In general complexity favours a heterarchical organisational structure while uncertainty favours a hierarchical organisational structure. The blackboard model is inspired by cognitive psychology in that it essentially models aspects of human cognition: knowledge sources represent permanent long term memory while the blackboard represents short term working memory. Blackboard systems use multiple knowledge sources to analyse different aspects of a complex problem. The blackboard control architecture provides a uniform system which integrates a number of diverse, specialised and independent knowledge sources which communicate through a common global database (blackboard). Cooperation through a central blackboard accessible by all agents alters the search problem by introducing massive speedup over using non-cooperative agents. The blackboard may be partitioned into a hierarchy of levels for different representations or abstraction levels of the problem. A separate domain independent goal blackboard with similar levels of abstraction as the domain dependent data blackboard improves cooperation control by allowing goal transmission between agents so enhancing coherence. Single blackboards are vulnerable to failure suggesting that a hierarchical network of blackboards which use a form of message passing may be advantageous. The blackboard architecture may implement a version of the contract net protocol.

Logic cannot accommodate inconsistency and so logical deductive inferencing is unsuitable for the foundation of intelligent machines due to the indeterminacy arising from the use of shared systems generating conflict through deductive indecision. Any form of local concurrent and interdependent activity will generate conflict and indeterminacy. Inconsistent beliefs are not meaningless since they are based on empirical knowledge which involves categorisation and interaction with the real world.

1. INTRODUCTION

1.1. OBJECTIVES

Biomimicry applied to space engineering promises the prospect of high miniaturisation, high integration and packaging efficiency, high energy efficiency, and high autonomy and robustness. We suggest that the primary *raison d'etre* for adapting biological solutions to the design of spacecraft and space missions is to enhance autonomy, i.e. relieve the ground-based (for robotic spaceflight) or spacecraft-based (for human spaceflight) human personnel from interacting with the space mission. This implies the need for extensive adaptability, robustness and reliability. The primary objectives (as stated in the Statement of Work GS 03/L27 [AD2]) of the project are:

- To provide an insight into recent, current and planned research activities in the field of biomimicry, in Europe and elsewhere (comments: much of the research in this area has been conducted in the USA)
- To provide an overview of the unique characteristics and properties of various life forms found in nature (e.g. animals, plants, microbes, etc) and to ascertain whether these characteristics could be an inspiration to create innovative space systems (comments: the Agency to perform the space systems application analysis)
- To conceptualise several innovative space systems and components which incorporate the design features and mechanisms of nature's life forms (comments: two case studies have been pre-selected by the Agency, and this proposal includes provisional suggestions for the other two)
- To assess what research and development might be required to bring life form properties and features to the point where they could be used in spacecraft (comments: wholesale replication of biological solutions is unlikely to be appropriate)

1.2. PREAMBLE

There is little doubt that biological systems are hugely successful in solving problems encountered in their environments. There is thus much that engineers can learn from biology to emulate in their design of engineering systems – this is biomimicry (or biomimetics, or bionics, or bio-inspiration). The success of biological organisms is attributed to the process of natural selection whose primary metric is such success – failure implies extinction. No wonder engineers have begun to examine biological systems to learn how organisms solve problems. It seems appropriate therefore to examine biological solutions to common biological and engineering problems. The reverse engineering of biological organisms has recently become a significant research effort to solve engineering problems by learning from biological solutions to similar problems imposed by the natural environment. Such reverse engineering of biology has met with varying degrees of success for a number of reasons, but renewed impetus has emerged following technological developments and increased scientific understanding. Such biological solutions offer insights into alternative strategies for designing engineering systems. Biological systems represent millions of years (billions of years in the case of microbes) of trial-and-error learning through natural selection according to the most stringent of metrics – survival. This project seeks to attempt to assess the current state-of-the-art in biomimicry and to apply this knowledge to the next generation of space systems. In this application, there is a curious mis-match in that biomimicry seeks to learn from organisms that evolved in environments on Earth, while space systems operate in a variety of hostile, non-terrestrial environments. However, by exploiting the robustness of biological organisms and the wide variety of environmental conditions in which they can survive, this mis-match should not present a difficulty. It is important however to ensure that we distinguish between functional aspects of biological evolution (which contribute to functionality) and incidental aspects (which do not). This is often difficult to do in complex systems as relationships and functionality are not necessarily well-defined – this problem may be illustrated by the question of whether consciousness is a necessary condition for intelligent behaviour, or is merely an incidental artefact resultant from the convoluted evolutionary history of the human brain. However, in many cases, abstraction of the central properties of biological systems which impart a given functionality is readily achieved. Evolution by natural selection does not provide optimal solutions to environmental problems but generates "satisficing" solutions due to the historical constraints of available genetic resources and the general requirement that an organism exists on the minimum amount of energy. However, although optimal solutions are characteristic of engineering, biological systems exhibit high robustness and adaptability to environmental variables implying a higher degree of redundancy in overall design. Certainly, the test for biological evolution is survival through functional effectiveness – the most stringent test within the hostile natural world.

Such robustness and adaptability is particularly critical in space exploration as the environments to be explored are typically unknown with unknown dynamics and variability. In particular, as space exploration missions demand greater sophistication and functionality, so the requirements for miniaturisation, autonomy, robustness, and survivability become paramount. According to Mjolsness & Tavormina (2000): "…how biological systems store and retrieve information, control development, fabricate structural components, build molecular machines, sense the external environment, reproduce and disperse themselves throughout the environment, engage in error detection, and carry out self-repair can pay big dividends to space exploration." This is particularly the case as space exploration missions give way from global reconnaissance to more focussed *in-situ* investigation requiring enhanced capabilities. A critical issue for space systems is that of autonomy – the distances involved, particularly once spacecraft venture beyond Earth orbit, preclude real-time control, necessitating high degrees of onboard autonomy. Indeed, although such missions are typically associated with exploration spacecraft, the issue of autonomy is also relevant to the Earth orbiting fleet, as ground station control is the dominant factor in operational costs. The greater the autonomy of a spacecraft, the greater its self-reliance as they operate at increasing distances from Earth deeper into outer space. The essential characteristic of an autonomous agent (a system with a particular function) is the capacity to allocate resources to perform its functional role and maintain its viability (Steels 1995). To be able to do this requires the ability to repair themselves and/or adapt to the environment. Such autonomy includes, but is distinct from, the property of automaticity – automatic systems are self-regulating in that they follow steering laws, but autonomous systems are self-governing in that they develop their steering laws. They must do this continuously by sensing, and acting upon its environment. This requirement imposes a consequent degree of loss of control but this loss of control is only partial – the relinquishing of only those aspects of control that ensure the survival and function of the spacecraft in response to uncertainty. The autonomous spacecraft is essentially a robot designed for function in outer space or within the influence of hostile planetary environments. However, engineered robots are highly deficient in comparison with the simplest organisms with regard to their sensory robustness, pattern recognition capabilities, adaptability to variability in the environment, and actuation flexibility. Biological systems on the other hand exhibit autonomy *par excellence*. They evolve, adapt and learn in a variable environment whilst maintaining their functionality. Even evolutionarily-primitive creatures such as insects exhibit high degrees of adaptive and flexible behaviour despite being endowed with limited neural tissue, a characteristic probably related to their skeletal design. Space systems more than many other engineering applications place emphasis on these qualities – a uniquely symbiotic relationship: as space systems demand more capabilities from biomimicry engineering, so biomimetic science will advance.

We are concerned with the application of biological solutions to the problem of spacecraft missions for exploration (and other space missions). Specifically, we are interested in replicating the capabilities of biological organisms in robotic spacecraft. Norman (1980) defined an animate system as one that has goals and purposes, maintains and protects itself (survival), regulates its own operation (homeostasis) and reproduces itself. Although biological life is characterised by self-replication as one of its basic properties, we exclude this aspect from further consideration due to its limited application to spacecraft of the near future. Fritz et al (1989) suggested that the primary goal of an intelligent system is survival and self-preservation (though more properly self-replication), and this is the guideline we use in applying biological lessons to the engineering of spacecraft. There are certain aspects that we exclude from this proposal as they are too immature technologically or of limited direct application to spacecraft design:

(i) Human hibernation and related closed loop life support systems as these relate to long-term technologies with applications limited to human spaceflight; furthermore, the skills mix of the proposal consortium has been built to contribute to the main issues of biomimicry rather than human and animal physiology.

(ii) Human-machine interfaces such as those being developed via neural-silicon interfaces are omitted as having limited applicability for space exploration missions as well as low maturity. This form of neural-silicon interfacing may be a long-term capability leading to "cyborgs" but is considered too long term for further consideration (and associated ethical issues).

(iii) Communication as exhibited in the animal kingdom has limited applicability to space communications and most lessons are currently incorporated (e.g. pulse code modulation and frequency modulation in animal signals rather than amplitude modulation); human communication through language is complex in structure (with syntactic, semantic and pragmatic aspects) and is not readily abstracted into applicable form beyond ready-existing programming languages based on syllogistic logic.

(iv) Generic artificial life, genetic networks and self-replication studies as these are undeveloped in terms of technological application (e.g. von Neumann's kinematic self-replicator).

(v) Nanotechnology including carbon nanotubes, RNA processing, etc are under-developed for technological application in the near future and would require considerable effort for treatment to do the subject equivalent justice.

(vi) Techniques of generic rather than specific biological origin, e.g. in-situ resource utilisation, which might be regarded as a form of "feeding".

We highlight five main broad biological principles that are particularly relevant to space systems:

6. The principle of autonomous behaviour to minimise reliance on ground systems
7. The principle of integration of mechanics, control, software and electronic systems
8. The principle of robustness for survival under widely variable conditions
9. The principle of compliant, multi-functional structures as part of a control system
10. The principle of neurally-inspired control systems with selectionism (e.g. neural Darwinism)

We believe that biomimicry has much to offer space systems design and the future of space missions. The planetary application aspect of space missions fits well with the role of biological organisms – the need to survive and function effectively in a variable, hostile and uncertain environment. Indeed, we submit that space applications provide a unique driver as a user application to biomimetic research in such a way that current piecemeal approaches to biomimicry will give way to structured and well-defined research programmes.

Biomimicry involves "reverse engineering" of the principles of evolutionary design of biological organisms in order to implement biological solutions to general engineering problems. The problems encountered by biological systems are in many respects similar to those encountered in engineered systems. Biological evolution has evolved robust solutions to many problems encountered in the environment through the most stringent of tests – survival (or more properly, inclusive reproductive fitness which may be approximately correlated with survival). In robotics, the commonest model for such biological inspiration is the insect. There are however, a number of important provisos:

- Biological solutions are not globally optimal, being dependent on evolutionary history and heritage, so application of bio-inspired solutions must be selective
- Biological organisms have evolved within a terrestrial environment rather than in space or other planets (though biological solutions may be sufficiently robust to accommodate the different physical and chemical properties encountered in space exploration)
- Biological evolution does not carve up the world into mechanical, structural, electrical, software control modules, rather that all aspects of biological solutions co-evolve together

Hence, the biomimicry employed for space systems is more properly bio-inspired rather than replicated wholesale from the biological organism. There are numerous examples of differences in approach to solving the same problems, e.g. fixed wing aircraft that separates lift from propulsion compared to flapping wings of flying animals that provide both lift and propulsion. However, given recent advances in materials and computing, bio-inspired solutions have become more amenable to emulation in engineered systems. Unlike conventional engineering approaches, biological solutions are integrative involving simultaneous co-evolution of mechanical structures with control systems to provide their functionality. The appropriate level of abstraction of such bio-inspiration is not a trivial issue, e.g. artificial neural networks (ANN) used for control systems are highly simplified and re-structured abstractions of biological neural architectures (neurons are typically represented as RC network-based processors) [Barto et al 1983]. It is not currently clear if such ANNs abstract sufficient properties from their biological analogues to replicate their function fully – biological neural networks exhibit complex dynamics including strange attractors characteristic of chaotic systems. Some aspects of such abstraction are enforced, e.g. abstraction of material properties in biomimetic materials. The chief motivation for pursuing biomimicry is that it promises a degree of miniaturisation, high integration and packaging efficiency, high energy efficiency, and high autonomy and robustness. As the robotic exploration of the solar system proceeds, so surveys will give way to more detailed and more challenging scientific investigation necessitating implementation of these characteristics. Similar characteristics will be required for robotic systems that support human missions. Such robotic systems will relieve astronauts from "cable repairman" tasks to those strategic tasks more suited to the human brain leading to greater cost-effectiveness of human deployment in space [Ellery 2002].

There are two major problems with the application of reverse engineering biological systems to spacecraft, both related to the notion that terrestrial biological systems have evolved in a terrestrial environment so the biological solutions may not have direct applicability:

1. The survivability of materials under space or planetary conditions - it is unlikely that many of the materials considered here (particularly those associated with biomimetic capabilities) are qualified for space operation. This will entail a programme of space qualification which we consider lies outside the scope of this study.

2. The problem of terrestrial robotic control research applied to planetary surfaces of rugged, rocky terrain – most research in mobile robotics has been based within artificial, indoor environments. Indeed, one of the commonest behaviours adopted in behaviour-based robotics is wall-following – walls do not exist on other planets! A related issue concerns vision processing, which is dominated by the extraction of straight-lines from images in order to delineate objects. Straight-lines however are not found in natural environments (though curiously, the mammalian visual cortex possesses neurons sensitive to straight-line orientations [Linsker 1986a,b]).

2. BIOMIMETICS EXPERTISE SURVEY

A separate database has been constructed for this section.

3. BIOMIMETICS RESEARCH SURVEY

This section reviews the field of biomimetics and its application to engineered systems. In particular, it provides the only integrative review of the field with such a wide coverage within a single document that we are aware of.

3.1. INTRODUCTION

We present a different, finer-grained division of the field of biomimetics in this section than that recommended in the Statement of Work [AD2]. This allows us to view the field in a more structured manner and to determine for TN3 a provisional mapping to spacecraft subsystems and technologies. We first consider the nature and difficulty of such division in terms of biological modularity. Our concern is to bear in mind the application of biomimetic technology to space missions. Our perceived mapping takes the current form, which reflects in part the structure proposed in this document (though not all aspects are covered:

Table 1 – Biological Analogues to Spacecraft Subsystems Matrix

Spacecraft Subsystem	**Relevant Biological System**
Space systems engineering	Evolution, embryonic development and cognitive psychology applied to modularity
Human element	Human-machine interfacing, closed loop ecology and hibernation strategies
Space environment	UV protection and oxidant tolerance
Propulsion system	Animal locomotion
Attitude control system	Animal navigation and vestibular system
Power system	Photosynthesis, ATP energy storage and in-situ resource utilisation (food)
Thermal control system	Thermoregulation, psychrophilic and hyperthermophilic extremophile strategies
Command and data system	Animal ethology, biological neural nets, central pattern generators, learning and behaviour control
Communications system	Animal communication and human language
Structural system	Biomimetic materials and structures
Payloads (sensors)	Special senses, active vision, optic flow, electronic noses and tactile sensing
Reliability	Autonomy, self-repair, immune system

A spacecraft is coupled to a hostile and highly variable physical world. The importance of physical bodies for intelligent robotics has been traditionally denied explicitly in the Good Old Fashioned Artificial Intelligence (GOFAI) approach. However, cognition is not abstract symbol manipulation and cannot be divorced from bodily constraints imposed by interaction with the real world – the biological world explicitly declares this as natural selection is imposed on organisms by their interaction with the physical world. This requirement for interaction with the real world is also imposed on engineering structures that must adapt to their surroundings such as robotic spacecraft. This interaction with the real world is the key to biological evolution – evolution effectively incorporates aspects of the environment into the structure and behaviour of the organism which enhance survival (or more properly, inclusive reproductive fitness).

This is how organisms are "designed" by natural selection. The artificial carving of the physical world into electronics, mechanics, structures and control systems, as is traditional in engineering, is one which is not adopted in biological evolution. A biological organism is an integrated machine including active sensory input, dextrous mechanical output and a complex integrated control system. The issue of control and that of structural compliance are tightly correlated and cannot be treated separately. Part of this capability is a direct product of the compliance of the organism's mechanical structure. The evolution of animal control systems occurs in parallel with body morphological evolution – they have co-evolved. Thus, biological control is one part of the functional component of biological structure.

Pfeifer (1996) suggested that there were a number of general design principles for situated autonomous agents (fungus-eaters), which are required to survive robustly in real-world environments:

1. fungus eaters principle – the situated agents must be complete systems with the properties of autonomy without human intervention, self-sufficient to sustain themselves over extended periods of time, embodied within a physical system, and situated within the environment with which it interacts.
2. ecological niche principle – the agent must be designed for a particular environmental niche.
3. parallel, loosely coupled processes principle – as exemplified by Braitenburg architectures, complex behaviours emerge from a large number of asynchronous independent processes without central control.
4. value principle – the agent must be embedded in a value system that is based on self-learning rather than pre-programmed categories.
5. sensory-motor coordination principle – sensory-motor coordination of the agent is dependent on its interaction with the environment, eg. movement often generates correlations in sensory patterns as the basis of grounded concept formation and object representation.
6. ecological balance principle – there must be a match between the complexity of the agent's sensors, control architecture, actuation capability, and its ecological niche.
7. cheap design principle – good designs are cheap in that they exploit the physics of interaction with the real world niche environment to minimise information processing requirements.

A simplified model of an autonomnous agent may be represented in Figure 1:

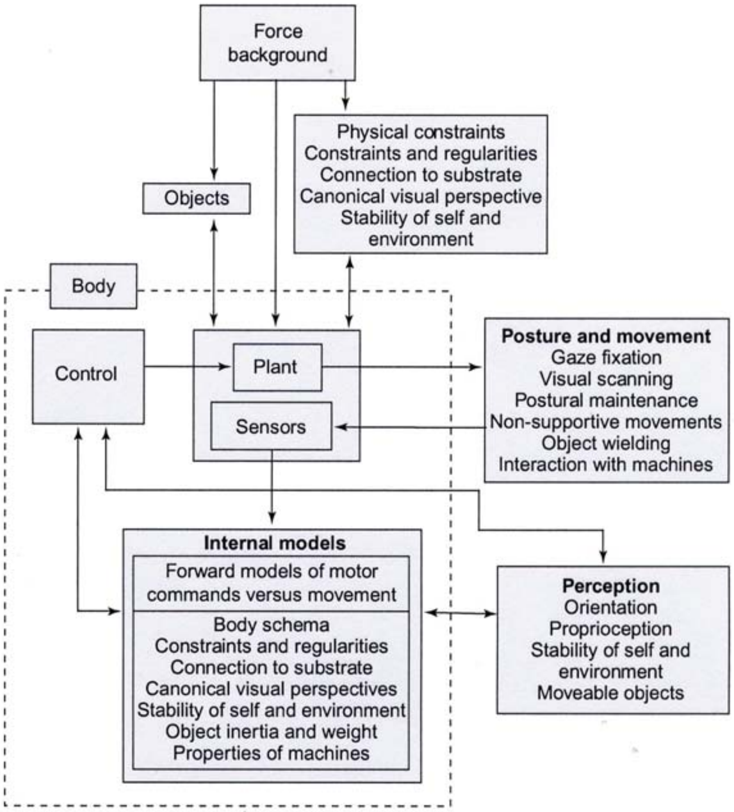

Figure 1 – Movement control requires complete agents (Lackner & Dizio 1998)

3.2. BIOLOGICAL MODULARITY & PLEIOTROPY

The first aspect of biomimetics in its potential application to spacecraft engineering that we consider is the issue of systems design. Our prime concern here is the issue of modularity (and hierarchical structuring). Modularity and hierarchical organisation is a fundamental approach in both science and engineering as the means for problem reduction, necessary for complex systems. Such modularity has been suggested as an approach to robotics design (eg. Figure 2).

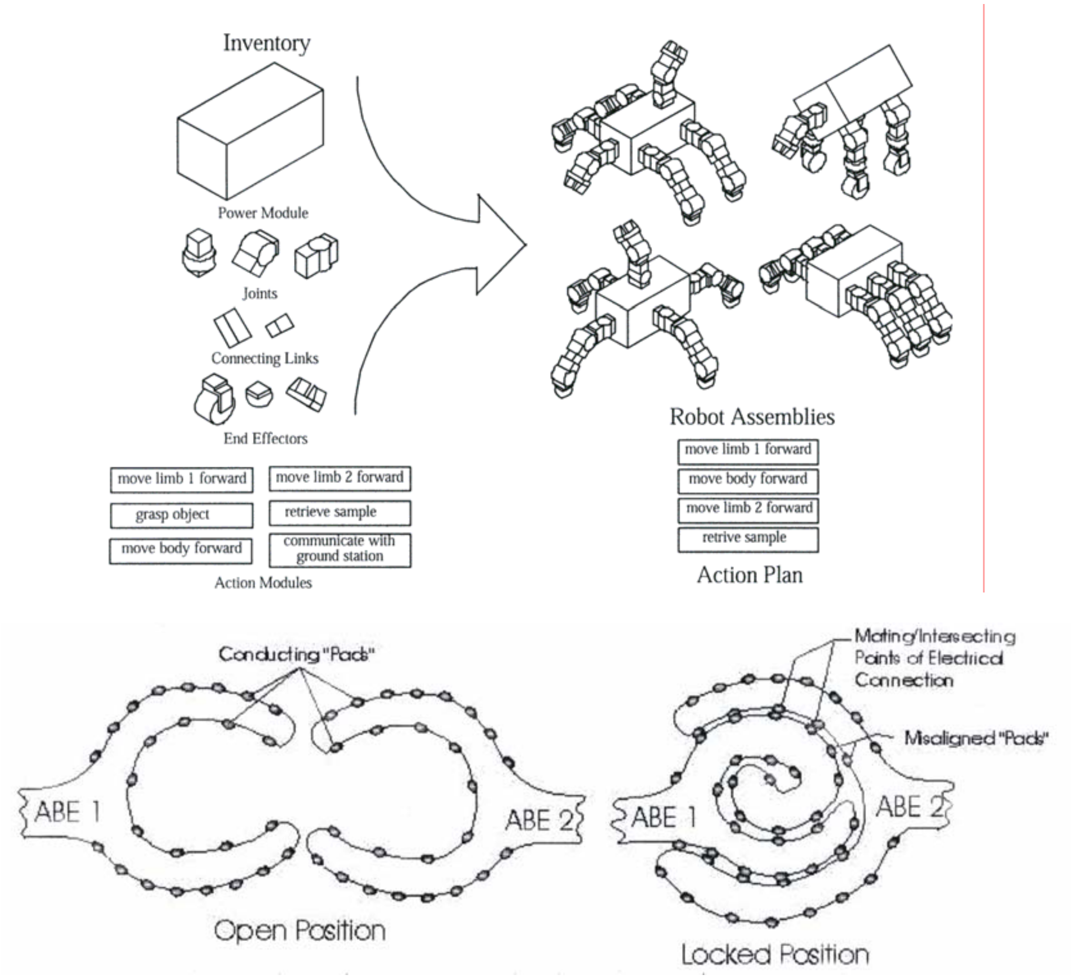

Figure 2 – a) Modular approach to robotics; b) Bi-stable compliant latch for connecting modules

3.2.1. Modularity

Modularity is central to cognitive psychology models of human information processing. Specialised modules are attributed to certain functions of the human mind whereby human information processes are characterised by a signal flow of transduction, modular input systems, and a generalised central system. Cognition represents the organisation of an organism's behaviour within the environment. Cognition comprises a number of independent neural modules each with separate functions and operating on different principles or but there are also cognitive processes which are different manifestations of the same general-purpose neural processing system (Fodor's thesis) [Fodor 1985]. Perceptual input systems are considered to be modular whilst central systems are unitary. Modularity implies special purpose, innately hardwired autonomous function. Such input systems are automatic, low level processes. Special purpose (domain-specific) modules exist for the processing of perceptual information and coordinating motor responses as interfaces to the real world. Perceptual input analyses are automatic and fast performed by genetically inherited hardware modules. Each module is functionally distinct and independent and do not communicate with each other (ie. distributed). These processes are not accessible either to conscious control or introspection and their outputs are highly specific. Other capabilities however are general manifestations of the same processes, eg. memory, problem solving, reasoning as these abilities overlap in their fundamental operation. Central systems are learned general purpose processes requiring attentional allocation of resources for high level cognition. Central processes provide higher level cognitive functions by integrating diverse forms of information from the different input modules with the knowledge structures of a world model stored in memory to derive an assessment of the present state of the external world. The exception to this of course is the Chomskian language module which is a general-purpose symbol processor but which is modular, innate and automatic like a perceptual processing module. Certain types of processes are performed in parallel and heterarchical in architecture such as linguistic analysis (debatably) and perceptual processing whilst other processes are inherently sequential and hierarchical in architecture. Some mental processes are automatic and outside conscious control and tend to be data driven, eg. perceptual processing and linguistic analysis while other mental processes are subject to conscious control and monitoring by the focussing of attention and tends to be goal directed, eg. problem solving and planning. Data driven processing is characteristic of low-level pattern recognition and runs on special purpose hardware while goal directed processing is characteristic of high level interpretation through expectation generation. Parallel processing is automatic, of high capacity and data driven and is characteristic of low level cognition. Serial processing requires conscious control, has severe capacity limitations and is goal directed and is characteristic of high level cognition. Inputs to the perceptual modules are in specialised data formats and are output in a common for format for the central processing system. It also provides integrative functions from the outputs of the specialised perceptual modules. The central processes are top-down hypothesis-driven processes characteristic of higher level cognition. Consciously controlled goal directed processes can become automatic and data driven with repetitive use to relieve the limited capacity of conscious control. This is an example of skills learning. It is generally argued that higher level processes are unitary with the possible exception of language processing, and that lower level processes involved with perception and motor control are modular and specialised [Slack 1989]. The whole processing architecture form a processing chain of sensory transduction, perceptual processing and central cognitive processing. Low level perceptual (and perhaps linguistic) processing is modular and each module is domain specific and self-contained while high level cognitive processing is non-modular and centralised employing higher levels of representation. The modules are genetically-determined, domain-specific, independent modules associated with perceptual input systems [Coltheart 1999]. Each module has its own dedicated, hardwired processes with fixed neural architectures to transform sensory data into a form suitable for the central, domain-general processing associated with general problem-solving processes. The modules are specialised, automated, rapid processors that perform fixed processing on their specialised stimuli in parallel. There is no interaction with more general, sequential, central processing. They are hardwired in the nervous system and operate outside of conscious awareness. Low level perceptual processing generally involves the extraction of invariants from sensory inputs (Gibson invariants). The central processes are associated with higher cognitive functions of problem solving and planning in providing top-down attention-focussing and expectations from long-term memory in a common. They are general-purpose, domain-general, sequential, slower and influenced by global cognitive goals. The high level cognitive processing involves inferential processing of the symbolic data integrated with the high level knowledge structures stored in long term memory for goal directed problem solving. It is this level that imposes order on sensory impressions through spatial and temporal structuring. Perceptual processes occur in the primary and secondary sensory areas of the brain while central processes occur in the association cortex. There are reckoned to be ~50 specialised modules for the 5 senses such as taste (sweet, bitter, salt, acid) and touch (temperature, texture) each with a dedicated pathway to the central processor. The blind spot is an indication of top-down processing interacting with bottom-up processing in human cognition. There are no receptor cells on the retina at the point where all nerve fibres coalesce into the optic nerve. This is no corresponding gap in the visual field as higher levels of processing fill in the blind spot. The knowledge structures of memory are unitary and global and cannot be organised as separate structures such as frames or schemas. The problem with this is that it imposes a computational explosion incurred by the nature of searching through large knowledge bases. However, it is inherent in the nature of neural systems that knowledge is stored in a distributed fashion throughout the network holistically. Hence knowledge representation is critical to any theory of human cognition. This also suggests that perhaps neural mechanisms at the hardware level may be critical for the functioning of human cognition. If cognition is studied as independent self -contained subsystems their behaviour may deviate from reality when they operate collectively. It is well established that specialised hardware exists in the brain concerned with sensory processing and language processing. Language processing is a curious capability. A compromise between modular automatic (for low level) and goal-driven conscious (for high level) processing is required for this uniquely human activity. There are several indications that language function is essentially hardwired and genetically inherited [Lenneberg 1964]: there are specific neural circuits in the cerebrum devoted to language analysis and motor speech coordination; the order of development of speech in children is always invariant across all cultures; language is not learned through generalisation; there are language universals of objects and relations which follow basic formal properties across all languages in all mediums. It is generally accepted that language ability arises through genetic pleiotropic effects such that it is coded through the interaction of many genes which code for other functions and that it is a separate phenomenon from intelligence and other behavioural traits as defective intelligence has no effect on basic language capability. Language, unique to humans, is fundamentally a symbol system which is generatively unbounded in that language is in principle capable of generating an infinite number of sentences. Syntactic analysis certainly must be automatic, yet semantic analysis must be at least partially conscious. Natural language understanding lies at the forefront of human cognition. Again, an information processing model has been advocated sequencing from syntax to semantics to pragmatics [White 1988]:

INPUT - SYNTAX - SEMANTICS – PRAGMATICS

The system involves analysing the input language to rules of set grammar (parsing). Semantic analysis comprises interpretation of the input's meaning. Pragmatics involves merging the information within all ready existing knowledge structures (eg. knowledge stored in a knowledge base of frames of context-related data). The knowledge base may have a parallel form to support all parts of internal language processing. One proposed structure for such knowledge bases are semantic networks which are linguistic structures with network links defined by relational logic which determine the inter-relation between concept nodes. A concept then is an aggregate of its relational links.

Karmiloff-Smith (1994) has attempted to synthesise Fodor’s modularity hypothesis with Piaget’s general cognitive development processes of assimilation, accommodation and equilibriation across all cognitive processes. Assimilation is based on previous behaviours towards objects while accommodation involves changes to those behaviours imposed by the environment. This views the modular/central processing dichotomy as too rigid – development involves a gradual process of modularisation. However, cognitive development does not proceed through well-defined stages as in the Piaget notion but is more gradual. However, there are innate biases to learning which focus attention on certain perceptual inputs to direct cognitive development along certain routes. By innately tuning attention to specific perceptions, this generates changes in cognitive representation over time, i.e. increasing modularisation. Physically, this involves progressive selection of brain circuits through the loss and gain of synaptic connections for different domain-specific functions similar to neural Darwinism. There is little doubt that some form of modularisation exists in infants – autism affects only social communication and reasoning about mental states leaving numerical and spatial cognition unaffected, while William’s syndrome affects numerical, spatial and problem-solving cognition leaving language, face recognition and reasoning about mental states unaffected. Similarly, in adults, brain damage often yields localised dysfunction. However, dysfunction in one sensory modality (e.g. deafness) often causes the brain to selectively adapt and reconfigure itself to receive other sensory inputs (e.g. visuo-tactile). This capacity for functional adaptation critically depends on environmental influences and so depends on the pattern of sensory activity. Furthermore, biological environments are structured and the interaction of innate biases and similar species-niches give rise to common developmental paths. The environment acts as more than a mere trigger but shapes the developmental process. Modularisation occurs when input-output processing become progressively less influenced by other neural processes as knowledge becomes chunked into procedural representations. Central conscious cognition is domain-independent and involves the development of cross-modal representational re-descriptions in the brain similar to natural language (which develops at around 18 months in the child). Only central cognitive processes are accessible to conscious introspection and verbal reporting.

This modularity is also dependent on experience during maturation in which competitive interactions between different regions of the brain, e.g. portions of the visual cortex are recruited for some auditory processing in the congenitally deaf [Jacobs 1999]. It is the process of competition between neural modules that generates functional specialisation to create a “mixtures of experts” architecture to the brain. However, it appears more likely that such modularisation arises through the process of development and maturation such that mental skills develop of increasing complexity which require increasing modularity and hierarchy formation as outlined by the Piaget framework – abstract representations are built from sensory-motor actions [Fischer 1980; Mallot 1997]. Nonetheless, the modularisation of human cognition is central to cognitive psychology. Human information processing follows this pattern and may be regarded as being based on two separate processes: automatic detection and controlled search [Schiffrin & Schneider 1997]. Controlled search is highly demanding in terms of attentional capacity while automatic detection does not require attention to select inputs. Controlled search is serial while automatic detection is parallel. Controlled search is under conscious control such that attention selectively filters the processing while automatic detection is hidden from conscious perception and difficult to modify. Controlled search is highly load dependent while automatic detection is unaffected by load. During motor skills learning controlled search gradually shifts to automatic detection through knowledge completion but requires extensive consistent training for routinisation. The implementation of skills learning requires consistent input/output mappings while controlled search deals with variable input/output mappings. This allows procedures to become implemented through automatic parallel detection which is unaffected by the load in bypassing the controlled search. Controlled processing is used in novel situations for slow but accurate performance. As situations become more familiar automatic processes become adapted to relieve attentional demands for other important and novel tasks. This allows efficient use of limited capacity processing. Only if stimuli are not appropriate for automatic response is the controlled mode activated. Automatic processes also provide the means for interrupting and overriding controlled search to draw attention to error.

The behaviour-based robotics paradigm suggests that modularity should follow evolutionary lines based on competence – this is the basis of the subsumption architecture for mobile robot control systems [Brooks 1986, 1989, 1991]. The subsumption architecture comprises hierarchical behaviour based units to provide robust and reflexive robot performance with a precedence organisation to avoid multiple activations of mutually exclusive behaviours:

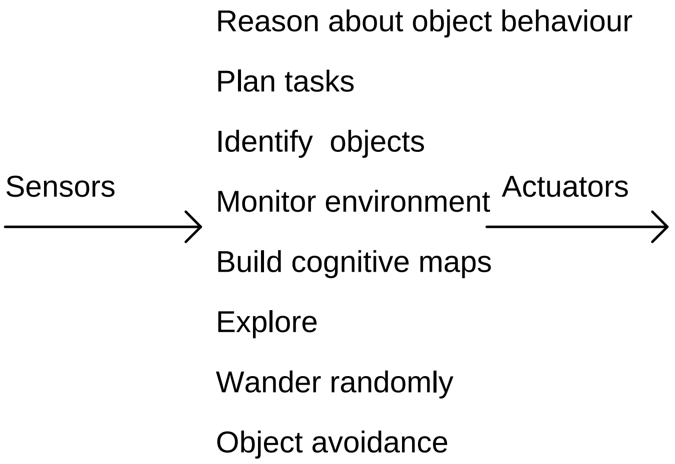

Figure 3 – Subsumption architecture task decomposition (Brooks 1986,1989)

The problem is decomposed vertically into a number of levels of task competence based on behavioural units rather than functional modules. Each behaviour was "wired in" between the sensors and the motors. The behaviours as functional units comprise of well-define control laws that are goal-oriented. Each task module (typically implemented as finite state machines) is connected directly to the outside world via sensors and actuators and operate in parallel. The control system is built bottom up incrementally and each level of competence includes the lower level as a subset. Simple tasks are solved first with new layers of competence added for greater complexity. This is an incremental approach but the addition of layers generates an exponential increase in the number of possible interactions between modules. Individual layers can work on individual goals concurrently whereby higher layers or modules inhibit (subsume) the behaviours of lower layers. Each level is implemented and debugged from the bottom up analogous to the evolutionary process of adding layers of greater behavioural complexity to animal brains to provide greater capabilities. This approaches bases its modularisation on behavioural lines rather than functional lines.

Koza (1993) applied the technique of genetic programming for the control of autonomous mobile robots. Genetic programs search for solutions to problems in the hyperspace of all possible compositions of functions that can be composed. Genetic programs operate on the level of symbolic expressions. A computer program comprises of a set of functions (e.g. logical operations). The computer program represents the solution to a problem. Toyically, S (symbolic) expressions in LISP are used as parse trees. S-expressions are thus rooted trees with ordered branches. A population of computer programs comprises a generation. The initial population of computer programs is random. Each program is executed and assigned a fitness value according to how well it solves the problem. The principal genetic operator is cross-over which swaps subtree portions of the parse tree from two parents to form two offspring programs. The point of swapping is randomly selected in both parents. As entire subtrees of such programs are swapped, the offspring are syntactically valid LISP S-expressions. Thus, new computer programs are created by swapping and recombining random segments of two programs. This process will produce increasingly fitter populations of programs. This process continues until an adequate computer program solution has been found or until a fixed number of generations have been run. Autonomous robots with 5 functional behaviours with inputs from sonar sensors were tested. The initial program population comprised of 4 primitive functional behaviours representing STROLL, AVOID, ALIGN, CORRECT comprising some 150 program functions such as move "back a fixed distance", "turn right/left a fixed angle", "move forward a fixed distance", etc. The purpose was to evolve the subsumption architecture. The subsumption architecture may be represented by the IF conditional to implement a decreasing priority of behaviours:

```
If $S_1$ then CORRECT
If $S_2$ then ALIGN
If $S_3$ then AVOID
If $S_4$ then STROLL
```

Fitness was defined as the maximum coverage of wall perimeter in the shortest time. The best 10% of parents surviving into the next generation intact and the crossover was applied to the other 90% of the programs to create new programs. A computer program was successfully evolved which exhibited the subsumption architecture via the composition of the conditional IF statement in performing wall following of an irregularly shaped wall.

Adaptation requires resolvability – the ability to produce improvement through random variations – which depends on the way in which genetic variation maps onto phenotypic variation through the process of development [Wagner & Altenberg 1995]. This representation aspect is central to the evolution of modularity in biological systems, and it requires a close to direct, one-to-one, relationship to constrain the effects of genetic variation on the phenotype. The primary reason for the evolution of modularity in biological organisms is to suppress pleiotropic effects between different modules [Wagner 1995, 1996]. Each module is an encapsulated, functional device, each being relatively independent of each other, i.e. there is little interference between modular functions. Only those developmental processes that can be dissociated can change their relative timing during development. Hence, modularity enhances the rate of evolution by minimising the detrimental effects of poor adaptations by limiting their effects. This is the basis of the building block hypothesis underlying the operation of genetic algorithms which ensures that sexual reproduction does not destabilise the gene string. Evolution by natural selection for the cumulative adaptations requires smooth fitness landscapes which is determined by the genetic-phenotypic mapping – independent functions can be coded independently with minimal interference with other functions so that pleiotropy is limited in scope. Pleiotropy is primarily confined to morphogenesis, which involves complex interactions of genes in the form of gene networks.

Arthropods, annelids and chordates are segmented structures but it is not clear if segments evolved independently or from a common ancestor (though it is clear that animal and plant evolved multicellularity independently) [Davis & Patel 1999]. Segmentation involves the serial repetition of structural units along the anterior-posterior axis of the animal’s body and appendages in bilaterians. All insect embryos start off a syncytium in which most of the nuclei later migrate to the periphery to form a blastoderm within which cellularisation occurs. Subsequently, embryogenesis in insects varies among different insects. In long-germ insects, segmentation of the different body segments occurs simultaneously. In short-germ insects, segments of the head form first, from which a growth zone progressively generates the thorax and abdomen. Intermediate-germ insects lie between these two extremes. Non-insect arthropods such as crustaceans, arachnids and myriapods appear to be similar to short-germ insects with thorax and abdomen developing from a posterior growth zone. In annelid worms, segmentation proceeds from the anterior to the posterior for each segment. In all cases, segmentation occurs in association with the expression of segment-polarity genes in a reiterated segmental pattern. In vertebrates, segmentation proceeds as progressive anterior to posterior growth zone similar to short-germ insects and crustaceans. It appears that there are some significant similarities in segmentation, the implication being common ancestry, but this is based on the assumption that pair-rule patterning common to all segmented animals is not a “logical” process – if it were “logical” then independent evolution becomes possible.

3.2.2. Hierarchical Structures

There is also the related issue of hierarchical structure – all multicellular creatures are hierarchical systems comprised of molecular assemblies, cells, organs, organism, and perhaps, groups of organisms. Hierarchical structures are common in natural, biological and artificial systems: language, organisations, cosmological structures, economic structures, biological organisms, ecosystems, etc.

Table 2 – Hierarchical structure of life (Bradley et al 2001)

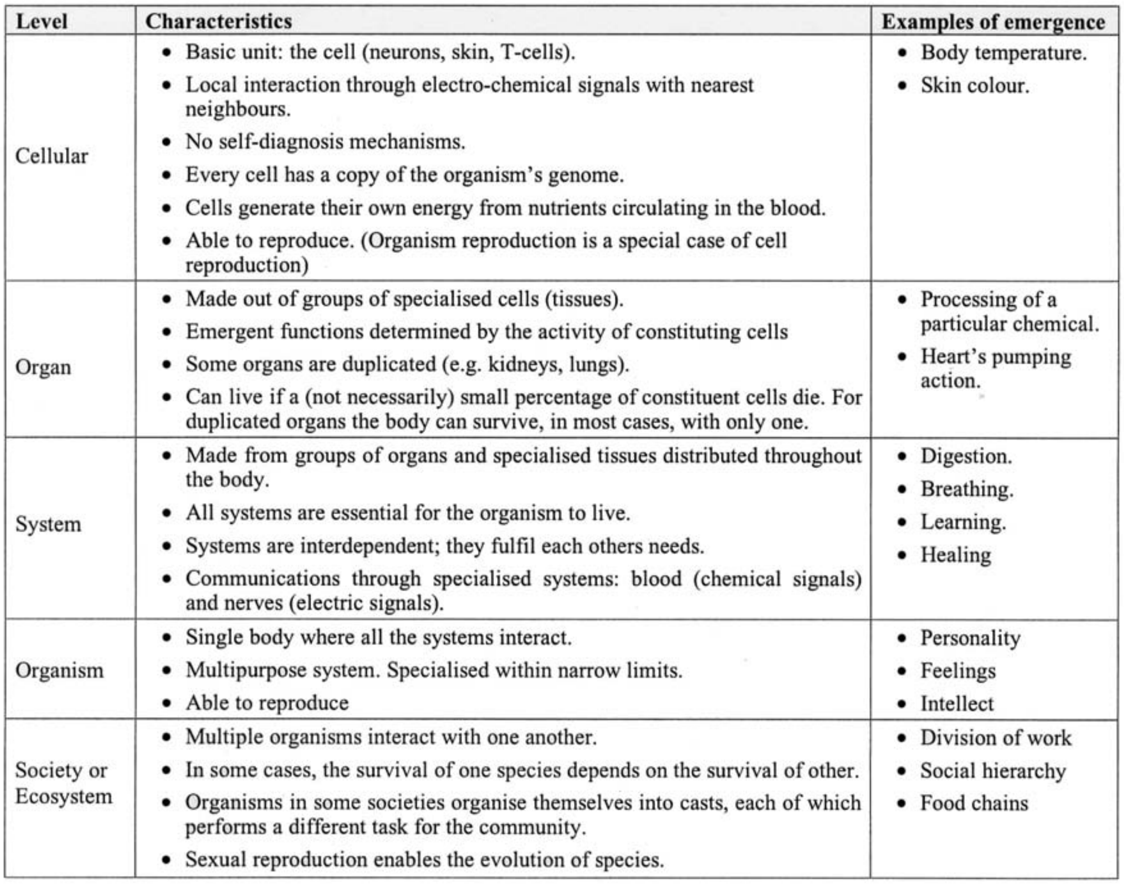

Level	Characteristics	Examples of emergence
Cellular	• Basic unit: the cell (neurons, skin, T-cells). • Local interaction through electro-chemical signals with nearest neighbours. • No self-diagnosis mechanisms. • Every cell has a copy of the organism's genome. • Cells generate their own energy from nutrients circulating in the blood. • Able to reproduce. (Organism reproduction is a special case of cell reproduction)	• Body temperature. • Skin colour.
Organ	• Made out of groups of specialised cells (tissues). • Emergent functions determined by the activity of constituting cells • Some organs are duplicated (e.g. kidneys, lungs). • Can live if a (not necessarily) small percentage of constituent cells die. For duplicated organs the body can survive, in most cases, with only one.	• Processing of a particular chemical. • Heart's pumping action.
System	• Made from groups of organs and specialised tissues distributed throughout the body. • All systems are essential for the organism to live. • Systems are interdependent; they fulfil each others needs. • Communications through specialised systems: blood (chemical signals) and nerves (electric signals).	• Digestion. • Breathing. • Learning. • Healing
Organism	• Single body where all the systems interact. • Multipurpose system. Specialised within narrow limits. • Able to reproduce	• Personality • Feelings • Intellect
Society or Ecosystem	• Multiple organisms interact with one another. • In some cases, the survival of one species depends on the survival of other. • Organisms in some societies organise themselves into casts, each of which performs a different task for the community. • Sexual reproduction enables the evolution of species.	• Division of work • Social hierarchy • Food chains

Marr's hierarchical structure for cognition with different levels of abstraction is one example. The top level is the computational information processing level; the middle level is the level of algorithmic analysis while the lowest level is that of the physical neural machinery. The hierarchical structure is universal. Pattee (1970) and Conrad (1970) suggest that biological systems possessing alternative modes of behaviour determined by events in the environment (i.e. adaptability) must evolve hierarchical levels of control to avoid the possibility of conflicting behaviours. The hierarchy is essential when the number of behaviours becomes large. The hierarchy is characterised by pyramid-like structure of decision and goal levels varying in complexity. Each level is coordinated by the level above and higher levels are distinguished by longer time constants [Bastin 1970]. The principle of division of labour is prevalent in biology by virtue of the competition for limited resources inherent in biological evolution – for instance, the evolution of multicellular structures illustrates the trend towards specialisation of different organs. One of the simplest multicellular creatures is the nematode worm, which comprises a simple entry cavity for the ingestion of food, a simple alimentary canal for the absorption of nutrients, and a simple exit cavity for the expulsion of waste. More advanced multicellular animals are variants on this body plan with more sophisticated and specialised organs – a mouth for mechanical mastication, a stomach for longer food storage and chemical processing, an intestine for further food processing and waste packaging and an anus for waste disposal. Such serial "assembly-line" division of labour prevents interference between physiological activities [Ghiselin 1978]. For instance, the male sperm duct in mammals also doubles as the male urine duct, both of which are activated intermittently and so without interference. If no interference occurs between functions, there is no need for such division of labour. A fundamental issue concerns the modularity of human cognition [Slack 1989], i.e. whether cognition consists of a number of independent neural modules each with separate functions and operating on different principles or whether all different cognitive processes are different manifestations of the same general purpose neural processing system. The brain appears to be constructed from a large number of distributed local but overlapping associative nets, some in parallel and some in series [Crick 1979]. The brain has a modular architecture of specialised components. The human brain has an approximately three level hardware hierarchy according to the simplistic MacLaren triune model – the reptilian part comprising the brain stem, mid-brain and basal ganglia; a palaeo-mammalian part comprising the rest excluding the neocortex and parts of the forebrain; and the anthropological portion comprising the neocortex and the forebrain. There exist $\sim 5\mathrm{x}10^7$ distributed, overlapping neural aggregates of $\sim 10^3$ neurons with well-defined functional and modular characteristics. Neurons are further subdivided into discrete groups of 50-150 neurons (ganglions). The connectivity within units is believed to be largely irrelevant with behaviour being primarily determined by their interactions. Hence, interactions between neuronal populations are not random but exhibit an ordered structure, so locally neurons may exhibit $\sim 10^5$ connectivity but modules exhibit low connectivity to maintain the synapse population at 10^{14} across the 10^{11} neurons.

The hierarchy is a method of decomposing complex systems into smaller interrelated subsystems nested into levels until some lowest level is reached. Hierarchies are based on successive delegation of functions from higher to lower levels, so that the number of distinct tasks increases down the hierarchy, each task at each level characterised by more limited contextual horizons. Each hierarchical level exhibits its own characteristic type of dynamics, and the dynamics at a given level is determined by the interaction of the components at the next lowest level. Each level relies on all the levels below it and imposes boundary conditions on lower operations. The decomposition generates simpler subsystems of approximately uniform information processing capacities at each level. There are, however, properties at each level that cannot be reduced to those at the level below as the whole possesses properties not exhibited by the sum of its parts [Polanyi 1968]. Dissipative structures are also characteristic of biological systems whereby macroscopic behaviour is an emergent property of the groups of independent behaviours of the structure's microscopic components. Weak interactions between these components determine the complete system's long-term dynamical behaviour. Decomposition into hierarchical modules is essential for the analysis of large and complex systems [Courtois 1975, 1985]. It assumes that the interactions within subsystems are much stronger than interactions among subsystems. Furthermore, it assumes that time and size scales are sufficiently far apart at different levels and that weak interactions have negligible influence on the system's overall behaviour due to their random fluctuation nature. Such decomposition relies on the assumption that any underlying processes at finer scales of time or space may be ignored. Hierarchical systems assume firstly that the complex system can be decomposed into inter-related subsystems until at the lowest level the elementary component is reached, and secondly that interactions within each subsystem dominate over interactions between subsystems. The basis for the decomposability of systems into a hierarchy of subsystems lies in the fact that complex structures grow and evolve more rapidly when they have stable intermediate forms characteristic of the hierarchy. Those stable forms are the subsystems with internal couplings that are stronger than the couplings between different subsystems. In the multilevel hierarachy the difference between weak and strong interactions exists at each level of the hierarchy: the weakest interactions occur between subsystems at the highest level and the strongest interactions inside the subsystems occur at the lowest level. Such nearly decomposable structures behave differently in short and long term periods due to this difference between weak and strong coupling. In short term periods as a result of stronger internal bonds subsystems tend to reach internal equilibrium "approximately" independent of one another. In long term periods when a whole structure evolves towards global equilibrium due to weak interactions among subsystems, the internal local equilibrium of subsystems reached at the end of the short term period are approximately maintained in relative value over the long term. Near decomposability depends on the partitioning to separate the short term dynamics of the subsystems from the long term dynamics of their interactions. However interactions between subsystems although weak are not negligible. Errors are caused primarily by these weak couplings between subsystems and are inversely proportional to the "undecomposability" of each subsystem - undecomposability of subsystems represents the shortness of the length of time to reach their internal equilibrium. However the constant of proportionality varies unpredictably. In the short term, the dynamic properties of the system are dominated by the strong interactions within each subsystem such that they quickly reach equilibrium independently. Long term dynamics is dictated by weak interactions between subsystems such that global equilibrium is reached after local independent subsystem equilibrium. Coupling between subsystems is small compared with coupling within subsystems. Because local subsystem equilibria are preserved in the long term, the macroscopic behaviour in the long term may be evaluated by a macroscopic model with aggregated equilibrated subsystems: each subsystem is replaced by a single state. We assume that in a nearly decomposable system, that the short-run behaviour of each component subsystem is approximately independent of the short-run behaviour of the other component subsystems. The hierarchical structure appears to be favoured by biological evolution under certain conditions. Biological systems extensively utilise the multilayer functional hierarchy for near autonomous subsystems. The control hierarchy constrains the behaviour of the elements such that they exhibit coherent behaviour.

Conant (1972, 1974, 1976) analysed the use of hierarchical structures in providing maximum efficiency of information processing systems. Every physical input/output device has a finite channel capacity which is the maximum rate at which information can flow through the device. Hierarchical systems come in two categories: classificational and organic systems. The classificational hierarchy are modular and requires logN bits to classify N labels – hence, such hierarchies are an efficient way of storing information. In computational systems, hierarchical structures are decision-making systems of computing elements with links representing causal communications media between elements. Vertical relations are authority relations that permit only unidirectional flow of information. Organic hierarchies are typified by societies, organisms and languages. Organic hierarchies permit horizontal information flow as well as vertical - they are characterised by two-way information flow with commands flowing downward and feedback flowing upward. The information flow cannot exceed the sum of the channel capacities of each element and the channel capacity of the entire hierarchy cannot exceed the sum of capacities of its elements. The number of subordinate elements directly linked to a superior element defines its span of control and its decision-making capability. Once information enters the system, it flows upwards to the element where a decision is made. This decision involves a loss of information such that only relevant features of the lower level information are retained. This allows convergence of information as it flows upwards in the hierarchy. It is the channel capacity limitation that forces this convergence. For efficient information usage, it is necessary for information to flow into the system from the environment at every level in a distributed fashion. Certain characteristics of the environment are relevant to different levels of the hierarchy and should be input to the relevant level to reduce the information flow burden between nodes. The biological organism must obtain information from its environment and make the appropriate responses in order to survive, i.e.

it acts as a communication channel. The cybernetic law of requisite variety states that it must obtain sufficient information from its environment and the cybernetic law of requisite transmission states that it must utilise this information efficiently in coordination with the environment. Hence, the biological organism's survival is bounded by its channel capacity. If the channel capacity of the organism required for survival is B bps while its senses are capacity limited to b bps such that b<<B, then it must structure its internal organisation hierarchically for efficient information processing. The number of elements and levels required to provide the requisite information to the decision-making elements is determined by the upward information flow. Hence, biological evolution will tend to yield hierarchical structures with weakly interacting nodes from the bottom-up. Hence, the organic hierarchy acts as an information channel.

Intelligence may be defined as the capacity to acquire and apply knowledge to reduce uncertainty. Furthermore entropy is also a measure of energy defined by the second law of thermodynamics - this is compatible with precision of execution at the lowest level of the hierarchy. Imprecision is a measure of uncertainty also. The entropy function should be minimised across all levels for efficient and optimal performance. A system S receives information y from the environment E and acts on it to produce an output response x. The information flow between elements provides coordination which allows solutions to problems beyond the capabilities of any constituent part. A measure of the total coordination between elements is given by the total transmission rate. The transmission is defined by the entropy of the system as a whole subtracted from the sum of the average entropies of each component. The information rate is bounded by the channel capacity of the input/output device: similarly, the transmission rate is bounded by the analogous constraint capacity which determines the maximum rate of internal information used for dynamic coordination within the system. In fact, the constraint capacity is bounded by the sum of channel capacities of the elements. The partition law of information rates may be expressed by [Conant 1976]:

Total information flow rate through all individual nodes of the system =

> ***input information flow rate which affects the output, i.e. the throughput rate***
> ***+ input information flow rate blocked by the system, i.e. blockage rate***
> ***+ internal information flow rate dedicated to internal coordination***
> ***+ noise rate***

$$F_o = F_t + F_b + F_c + F_n$$

The throughput rate gives the input/output flow rate through the communications channel and is limited by the channel capacity. The blockage rate gives the rate at which information about the environmental input is blocked within the system and measures the relation between the environment and the system, i.e. the effect of E on S. It is common to filter information from the environment and select only relevant information which affects the output. The blockage rate is then the rate at which irrelevant information is discarded. The coordination rate is a measure of the total internal coordination between all elements, i.e. the amount of relatedness. The noise rate is the amount of random uncertainty generated without deterministic cause. The noise term is decomposable into two terms – the external noise rate which is directly observable and reflects environmental uncertainty, and the internal noise rate which is internally generated but is not directly observable. Each of these information flows are bounded by the channel capacity. Since the total information flow is fixed by the sum of elementary channel capacities, then a high throughput is only possible for low environmental information inputs requiring little coordination and blockage. Complex information inputs however require significant internal coordination and high blockage at the cost of reduced throughput. Hierarchical systems are information-losing in the sense that their outputs carry reduced information rates than their inputs: $H(E) \geq H(S)$. The overall loss of information is due to internal blockage (which affects S) and the rejection rate at the system/environment interface (which does not affect S). Rejection merely requires insensitivity to the environmental information through passive filtering so it is desirable to lose as much information as possible in rejection to minimise the requirement for active blockage which is capacity limited. The objective of the internal blocking mechanism is to block information that cannot be so rejected by passive filtering, eg. habituation is a diminishment of response to repetitive stimuli. Similarly, internal coordination flows must be minimised. The system S may be partitioned in N subsystems S_i each of which receives input E_i:

$$H(S)=\sum_{j=1}^{N}H(S_j)-T(S_1...S_N)$$

The entropy rate of system S equates to the sum of entropies of each subsystem S_i less the transmission rate between subsystems. It is desirable that each of these information flows are operated at their maximum capacity, and a properly designed hierarchy enables this (preferably with minimum blockage and coordination to maximise the throughput with the overall channel capacity limits). This requires that coupling between subsystems is low, so that S is decomposable into independent subsystems to minimise internal coordination between subsystems. Each subsystem is thus modular. Furthermore, the hierarchy allows parallel information processing at each level [Singh & Singh 1978]. Traditional hierarchies are perceived to model human intelligence (and that of the machine): if knowledge flow measures the efficient use of knowledge, it provides a measure of intelligence. Intelligence involves the efficient use of knowledge defined as structured information. Increasing feedback knowledge reduces the amount of entropy which measures the uncertainty of the system. By equating machine intelligence as operating on a knowledge base to produce a flow of knowledge R the principle of increasing precision with decreasing intelligence can be derived such that the highest level of machine intelligence is associated with the smallest complexity of detail and the lowest level of machine intelligence is associated with highest complexity of detail [Saridis 1988, Valavanis & Saridis 1988]. Pollatschek (1977) suggested that hierarchical systems are fuzzy whereby the degree of fuzziness was proposed as a numerical characterisation associated with the level of the hierarchical system - more detailed lower levels of the hierarchy are consequently less fuzzy, i.e. the degree of vagueness increases up the hierarchy with increasing robustness to error. Fuzzy techniques may be applied to determine the maximum absolute errors permissible at each level that does not alter the description of the system. Both sensed data and stored models are approximate and incomplete and control must therefore accommodate this uncertainty pertaining to its hierarchical level. Courtois (1975) warned of the pitfalls of using too simplistic models of necessarily complex systems such as the hierarchical structure of intelligent control in cases where interactions between components or agents were significant. Dissipative structures are phenomena based on structures whose macroscopic behaviour emerge from individual independent behaviours of many microscopic entities. The macroscopic behaviour of these structures result from the amplification and cumulative effect of interactions on the microscopic level.

Autonomous robotics is based on the notion of distinct modules associated with basic behaviours. Emergent behaviours such as wandering and navigation arise through coupling between each more basic behaviour module. This coupling may be temporal or spatial. Spatial ordering involves concurrent activation of the component behaviours. Anderson & Donath (1990) held each of the primitive reflexive behaviours to be spatially ordered rather than hierarchically patterned. They used multiple components of reflexive behaviours which operated concurrently and independently. They found that high level global behavioural patterns emerge which are not ascribable to individual simpler components at lower levels. Control structures emerged from the distributed loosely coupled behavioural modules through parallel local interactions of a non-linear nature. This organisation gave rise to more complex and irreducible behaviours. The difficulty with spatially ordered behaviours is that they can cause conflict resulting in erratic behaviour. The total number of unique behaviours emergent from different primitive behaviours may be assumed to be a binomial function in that *i* possible combinations of *n* items taken *k* at a time is given by:

$$i = \sum_{k=0}^{n} \tfrac{n!}{k!(n-k)!}$$

Alternative orderings include hierarchical coupling of primitive behaviours by adding temporal ordering. Alternatively, behaviours may be weighted according to a learning algorithm. If temporal ordering is included through hierarchical coupling number of potential emergent behaviours will be increased. The behavioural modules are usually considered to be hierarchically organised such that each instinct is decomposed into finer grained behaviours. Instincts at the same hierarchical level compete with each other. Thresholding behaviour eliminates the possibility of chaotic behavioural responses. Spatially distributed ordering introduces the potential for cyclic behaviour due to the deterministic nature of mobile response and the inability to respond to events over several time periods. This is a consequence of a lack of memory which limits response flexibility when a reaction to an input may be dependent on previous inputs. Memory provides the capability of reacting to events over a number of time intervals and so alter behaviour on the basis of previous behaviour. Memory is based on the process of classification as its function is fundamental to recognition – concepts or situations to be recognised are organised into hierarchies which represent relationships between the concepts.

3.2.3. Human Physiology Model

We briefly look at human physiology as a biological system to gain some insights into biological approaches to systems design. Biological control is a fundamental component of animal physiology and is characterised by online adaptation. General patterns of metabolic and physiological activity are essentially constant (homeostasis). An animal's behaviour must maintain essential variables within prescribed limits in a changing environment, i.e. be stable. This overall adaptation must be persistent in the event of significant failures, i.e. be reliable. Animal physiology may be regarded as a system of interconnected feedback systems comprising the metabolism. Human physiology is based on multiple homeostatic feedback systems, the major ones being the cardiovascular system, the respiratory system, the digestive system, the excretory system, the nervous system and the endocrine system. The metabolic system provides MIMO (multi-input multi-output) integration through cellular activity. Many diseases are believed to be the result of loss of control of interacting elements rather than the breakdown of any single element in the body.

There are around 200 different cell types in higher multicellular animals but tissues are composed of four primary cell types – epithelial cells which form protective sheets of tissue lining, connective tissue cells based on the protein collagen, muscle cells which can contract, and nerve cells which can conduct electrical signals. There are 10 major organ systems in the human body:

(i) skeletal system provides structural support and attachment sites for organs;

(ii) muscular system provides movement;

(iii) circulatory system transports nutrients, dissolved gases and wastes through the body;

(iv) nervous system relays electrical signals through the body for the control of movement and integration;

(v) respiratory system provides gas exchange between the blood and the external environment;

(vi) digestive system breaks down and absorbs nutrients for metabolism;

(vii) excretory system filters out and expels wastes from the blood;

(viii) endocrine system relays chemical messages through the body for the integration of physiological processes;

(ix) reproductive system manufactures gametes to enable reproduction;

(x) the lymphathic/immune system provides defence against pathogenic invaders.

In mammals, these systems are interlinked in a number of ways, the primary linkages only described here. The circulatory system provides for the transport of materials for adequate energy generation while excretion expels waste products. Blood is the main transport system of the body which supplies the cells with nutrients and removes waste products. Blood is pumped around the body by the heart which comprises two pumps. The left side pumps oxygenated blood through the aorta and arteries to the tissues while deoxygenated blood returns through the veins to the right side which pumps blood through the pulmonary artery to the lungs. Oxygenated blood returns to the left side of the heart through the pulmonary vein from the lungs. Oxygen is absorbed from the air in the lungs by the blood as it passes through the lungs for transport to the tissues. Waste carbon dioxide is produced by the tissues and is offloaded at the lungs by blood for expulsion. Nutrients are ingested through the gastric system for the production of energy and growth of tissue. Enzymes in the digestive tract break down large molecules of food which is absorbed into the blood for transport via the portal vein and liver to the tissues. Undigested food is expelled via the rectum. Waste products of cells except carbon dioxide are excreted from the body in solution via the kidneys and urethra. There are two global control systems in the human body - the nervous system and the endocrine system. The electrical nerves and slower hormonal (endocrine) system act to provide online adaptation to the environment. The endocrine system controls the functions of the organs through the secretion of chemical hormones from the endocrine glands. This controls metabolism through enzymatic feedback through hormone secretion into the blood by specialised tissues. They play a major role in the reproductive system. The nervous system exerts control of the muscles through electrical impulses from the brain through the spinal cord. Feedback is provided back to the brain from the sense organs - skin (touch, pain, temperature) and the special senses (sight, smell, hearing and taste). The nervous system includes sensors to provide information about environmental conditions such as the involuntary human pupillary servomechanism and eye-head reflex control. Many physiological systems are under greater or lesser control by the brain in one form or other. The hypothalamus also controls the secretion of thyroid-stimulating hormones produced by the pituitary gland which in turn governs the secretion of thyroxine hormone, an iodine containing peptide which integrates growth. The hypothalamus also sends fibres to the sympathetic and parasympathetic systems of the brainstem and spinal cord. These efferent sympathetic and parasympathetic pathways pass onto the respiratory centre and vasomotor centre to control respiration and heart rate. From the spinal cord, pairs of spinal nerves are distributed from the cord – 8 cervical nerves, 12 thoracic nerves, 5 lumbar nerves, 5 sacral nerves and 1 coccygeal nerve to each side of the body. In addition, 12 pairs of cranial nerves project from the brain to innervate the head. Each spinal nerve has a ventral (anterior) efferent pathway and a dorsal (posterior) afferent pathway. Voluntary movement involves activation of striated muscle fibres by excitatory activity only. Involuntary movement is controlled by the autonomic nervous system, eg. gland secretion and smooth muscle innervation (eg. digestive tract, respiratory system, bladder, blood vessels). They are innervated by both excitatory and inhibitory nerves acting in coordination. Inhibitory/excitatory pathways are controlled by different systems – the parasympathetic nervous system based on acetylcholine transmitter and the sympathetic nervous system based on noradrenaline transmitter which depending on the function, eg. the eye pupil excitatory nerve is under the control of the parasympathetic nervous system. The sympathetic nervous system which emerges from the first thoracic and second lumbar segments is widely distributed throughout the body while the parasympathetic nervous system which emerges from the second, third and fourth sacral segments and cranial nerves is restricted to the trunk excluding the limbs. The sympathetic nervous system noradrenaline is secreted as hormones by the medulla of the adrenal gland to augment its function, while no such endocrine system exists for the parasympathetic system. The hypothalamus also controls multiple hormonal secretions from the pituitary gland (and body fluid volume through the kidneys). It also regulates temperature through sweating and shivering. It also governs primitive emotional responses.

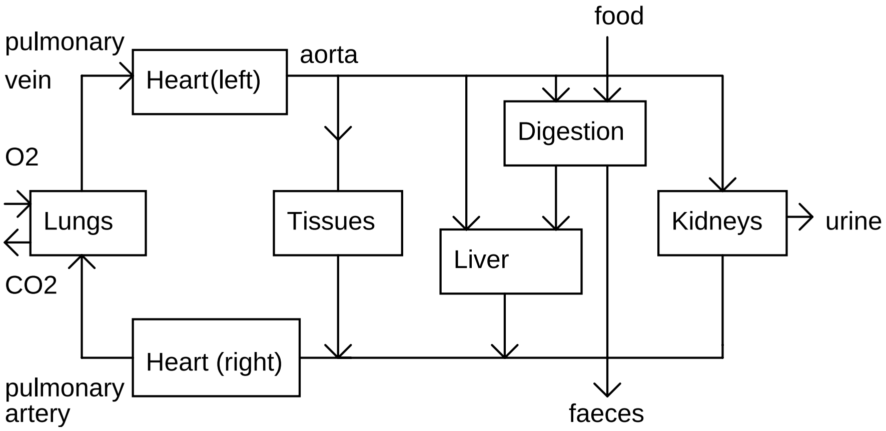

Figure 4 – Schematic of interleaved human physiological control loops

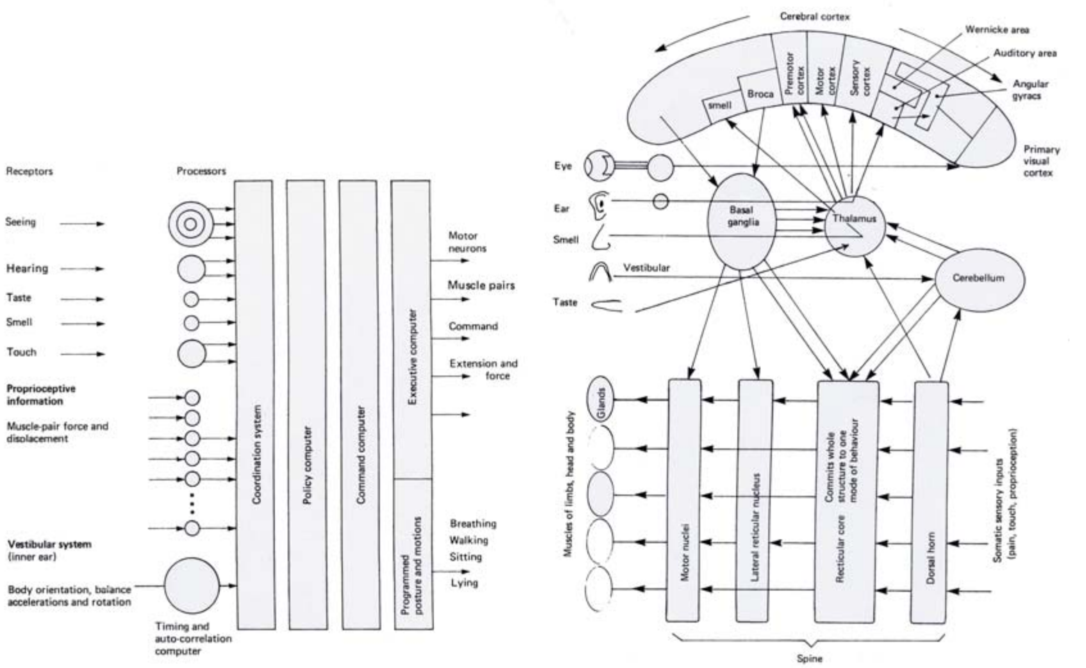

Figure 5 – a) schematic of human input-output system; b) schematic of human central nervous system

Homeostatic stability requires the ability to learn when the environment changes and reliability requires self-repair when the agent's state changes detrimentally. Adaptation is the process of generating and maintaining successful behaviour in a changing environment such that it maintains performance parameters within certain specified stability bounds. It also involves maintenance of overall functional reliability in the face of failure of machine parts, i.e. changes of specified stability bounds involving the self-repair of components [Slansky 1964]. The adaptive system automatically attempts to reduce uncertainties and the central feature is that this adaptation is a negative feedback mechanism. Systems involve multiple elements interacting with each other and can generate complex behaviour. This complexity is the product of feedback. Indeed, the comparison between animal and machine is striking; sensors (sense organs) provide information about the environment, the power transmission medium (blood supply) provides the energy to operate the system, while the control system (nervous system) provides the interface to process all the necessary information to perform the task. Organisms are open systems that maintain themselves in steady state through the expenditure of energy despite changes in the environment. Organisation is maintained despite changes in the environment tending to enforce its dispersal. Organisms receive information about the environment and this allows them to control their behaviour accordingly.

All smooth muscle, cardiac muscle and gland cells of the body receive motor innervation from the autonomic nerves including the entire cardiovascular, respiratory, digestive, endocrine, urinary and reproductive systems. The peripheral autonomic (visceral) motor nerves begin as axons from the brainstem and spinal cord via a number of ganglion relays. The autonomic nervous system comprises two systems – the parasympathetic and sympathetic systems. Terminals of post-ganglionic fibres of the parasympathetic system secrete acetylcholine (cholenergic) while the terminals of the post-ganglionic fibres of the sympathetic system secrete epinephrines (adrenergic). If the supply to the organ is dual, the two systems oppose each other. Both work antagonistic actions on the smooth muscle, cardiac muscle and glands, including the lungs, stomach and intestines. The parasympathetic nervous system supplies the secreting glands. Sympathetic nerves supply vaso-constrictor fibres to all blood vessels (except arteries supplying the myocardium). Pain, cold, and strong emotion evoke response from the sympathetic nervous system which mobilises the body for action. This is because the sympathetic nervous system has global effects on the body by virtue of its pervasion through the body and the release of epinephrine from the adrenal glands.

In carbohydrate metabolism, glucose is absorbed from the intestine and passes through the portal vein to the liver in which it is phosphorylated. Some is stored as glycogen, some is broken down to pyruvate, and some dephosphorylated for transport in the blood plasma to the cells. Inside the cell, it is phosphorylated again for storage as glycogen (3%), fatty acids (30%), or metabolised to pyruvate and then carbon dioxide and water (67%), releasing energy. Carbohydrate metabolism is regulated by the hormone insulin secreted by the pancreas. The kidneys are responsible for the maintenance of tissue fluid at constant composition (through resorption of water and Na), pH (by the excretion of H ions as many end products of metabolism are acidic) and volume. Renal function is essentially a filtration process followed by resorption. Buffer systems, such as the strong base Na bicarbonate, in the blood maintain normal acid-base balance in the blood by forming weak acids on the addition of H+. It is through the detection of Na concentration that the kidneys maintain blood volume. Carbonic acid, which is the primary product of the oxidation of food, is eliminated as anhydride carbon dioxide conveyed by haemoglobin to the lungs. Other acids such as phosphoric, sulphuric and organic (e.g. lactic) acids are excreted as anhydrides in solution through the kidneys. Blood contains three types of cells: red cells (erythrocytes for oxygen/carbon dioxide transport), white cells (leucocytes for immune defence) and platelets (thrombocytes for injury repair), all formed in bone marrow.

The endocrine system functions on a longer timescale to the nervous system and coordinates a wide variety of functions. It mediates metabolic activity through hormones which are discharged into the bloodstream. The target organs are self-selecting. Positive and negative feedback is exerted by the hormones and their metabolic effects. The main endocrine glands include the hypothalamus, pituitary, thyroid, pancreas and adrenal glands. Top level control is provided by the hypothalamus which integrates feedback and environmental signals. The hypothalamus acts on the pituitary gland primarily but also on the pancreas. For instance, the hypothalamus corticotropin releasing factor (CRF) causes the pituitary gland to release adrenocorticotropic hormone (ACTH). Cortisol provides negative feedback to restrain the secretion of CRF and so ACTH release. The pituitary hormones act directly on tissue receptors or other endocrine glands to release secondary hormones. Hormones secreted by target glands activate receptors which elicit response through a second messenger, eg. cyclic AMP. The adrenal gland secretes two types of hormone: glucocortisoids which regulate carbohydrate metabolism (cortisol); mineral cortisoids which control the transport of Na across cell membranes and excretion by the kidneys. Both the nervous and endocrine systems share evolutionary origins in primitive chemical mediators and the borderline between the two is often diffuse, eg. the pituitary hormone vasopressin (ADH) is synthesised in the hypothalamus and secreted down the axons of the nerve fibres. Adrenaline released at the nerve endings is a neurotransmitter while adrenalin released from the adrenal medulla is a hormone. Indeed, the adrenal medulla has both nerve and endocrine cells, differentiated only by the staining of its endocrine cells by chromic acid. Furthermore, the hormones of the gut and pancreas are strongly influenced by the autonomic nervous system.

3.3. BIOMIMETIC MATERIALS

3.3.1. Materials of Nature

Biological materials and structures are characterised by multiple levels of organisation which impart toughness and multi-functionality in comparison with more traditional engineering materials. Biological materials possess their unique characteristics by virtue of their organisation rather than their material substrate. Biological materials are adaptable with the capacity to sense aspects of the environment and react accordingly – they incorporate "smartness" to their structure by integrating aspects of control. The use of natural materials for engineered structures is not completely unknown – the British Mosquito bomber of World War II was constructed primarily with wood making it very lightweight and fast while affording resistance to impacting gunfire. Wood possesses resistance to crack propagation initiated by impact by virtue of its cellular microstructure [Gordon & Jeronimidis 1980]. Cellular compartmentalism within structures is the approach exploited in honeycomb panels for spacecraft primary structures.

Most biological materials are soft with stiffness imparted through mechanical rigidisation by fluid pressure or through viscosity changes in gel materials – the biological analogue to electrorheological fluids. They are generally compliant with low Young's modulus and provide strength for resistance to breakage rather than stiffness for resistance to deformation. Composite materials are ubiquitous in biological systems in which the structure incorporates sensors, processors and actuators [Srinivasan 1996]. Most biological materials are composites comprising polymer fibres embedded within matrix materials – like engineered composites, their physical properties are highly anisotropic. Engineered composites may incorporate fibre optic sensors to detect a range of physical parameters through distortion of the propagation path of the optical signal: temperature, strain, pressure, electric and magnetic fields by the interaction of the parameter to produce a modulation of amplitude, phase or polarisation of a transmitted signal. Interferometric phase sensors detect the influence of physical parameters on the phase of coherent laser light in a single mode fibre with respect to reference arm fibres. Biological polymer fibres include cellulose in plants, collagen in animal tendons, chitin in arthropods and silk in spider webs. Cellulose is a polysaccharide that forms fibrous structures in plants similar to collagen in animals. Such fibres have analogues in artificial materials such as Kevlar, graphite, S-fibreglass, and polyaramid. Biological matrix materials may be polysaccharides or proteins with their equivalent artificial analogues of epoxy or phenolic resin. Structures with high strength but low density, many of which are composite resembling fibre-reinforced composites, are used in aerospace applications. Some biological composites are reinforced with minerals such as calcium carbonate in shelled marine animals and hydroxyapatite in vertebrate bone. Shark skin is similarly a composite material of layers of collagen fibres arranged in alternating right and left hand helices around the shark's body. The shark's body represents an internally pressurised shell structure ~20-200 kN/m^2 whose muscles contract along the circumference to allow length changes along the longitudinal axis whilst maintaining constant volume. Mollusc shells are also comprised of composites of ceramic phases such as $CaCO_3$ "bricks" embedded in a protein matrix. The abalone shell comprises an outer layer of calcite and an inner layer of aragonite within an organic matrix with a strength of ~100-150 MPa. Shells also limit crack propagation by the boundaries between shell layers. Furthermore, the shell microstructure which orients cells at different angles increases toughness by making the crack path increasingly branched.

The outer covering of arthropods (insects, spiders, crabs, shrimps, etc.) is known as cuticle, which is a composite material consisting of arrangements of highly crystalline chitin nanofibres embedded in a matrix of protein, polyphenols and water, with small amounts of lipid. It is secreted by a single layer of epidermal cells that covers the entire surface of the animal. It can be as thin as 1 μm in the hindgut and over gills (e.g. may-fly larvae) and as thick as 200+ μm (wing case of large beetles). Chitin is a fairly completely acetylated polysaccharide akin to cellulose. It is assembled into nanofibres about 3 nm in diameter and about 0.3 μm long, each containing 19 molecular chains of stiffness at least 150 GPa. In suspension in water, chitin nanofibres are highly thixotropic and liquid crystalline. Stiffness and thixotropy are crucial properties for the self-assembly of the components of the cuticle. The number of chitin chains in the nanofibre is probably close to a minimum for internal stability of the crystallite; hence the nanofibres present the maximum surface area for interfacial interactions within the cuticle. The protein has to produce a matrix of varying mechanical properties, which will also interact with and stabilise the chitin. In soft cuticles such as the highly extensible soft intersegmental membrane cuticle of the female locust there may be only 20 different proteins whereas in some of the stiffer cuticles in the same insect there may be 200 or more. This complexity is probably more apparent than real, since the conditions of separation of the proteins on gels picks up slight differences in charge, which could be given by a single amino acid substitution in the sequence, which may not be so important in a structural protein. Also the proteins are probably not different in the different areas, as can be shown by immunological comparisons.

The interaction of the remaining proteins with the chitin seems to be fairly consistent in that even in the softest of cuticles (in which the interactions are least developed) a strong solvent (e.g. 5% NaOH at 100C) is required to remove the protein from the chitin. X-ray diffraction of the ovipositor of the wood wasp *Megarhyssa* suggested that the proteins surround the chitin in a regular manner, but while we cannot calculate the strength of the interfacial interaction between the chitin and the protein (fibre-matrix), we can at least put a lower bound on it, guessing that the shear strength of a single H-bond is of the order of 30 pN, and calculating that the area around each bond is about 10-18 square metres, then the shear strength will be about 30 MPa, or about half that measured for carbon fibres in a resin matrix. The addition of matrix material to chitin is at least a 2-stage process; the chitin nanofibres are first coated with a specific protein, which adheres through a silk-like conformation (beta-sheet) then the rest of the protein is added, interacting with the proteins already attached to the chitin.

Conformations other than beta-sheets exist in cuticular proteins, although they have not been much explored. They have been mapped in some proteins. In particular the highly rubbery cuticle protein(s) called resilin now appear to contain beta-turn structures similar to those found in elastin, a rubbery protein typical of vertebrate animals. Comparing the surface area available for bonding between matrix and chitin; if the minimum practical size for a carbon fibre is a diameter of 5 μm, but only half of the surface area of the chitin nanofibril is available for matrix bonding, then the surface area of chitin per unit volume of cuticle is 106 greater than the carbon fibre material. But with the huge surface area difference it is fairly obvious that the total fibre-matrix interaction per unit volume in insect cuticle is still 5 x 105 greater than carbon fibre composites. 'Soft' (compliant) cuticle commonly contains equal weight fractions of chitin and protein, and 40-75% water. 'Hard' (stiff) cuticles contain 15-30% weight fraction (dry) of chitin and only 12% water. Stiffness can vary from 1 kPa in the highly hydrated extensible intersegmental membrane of the locust, via about 1 MPa for the hydrated rubbery protein, resilin, to 60 MPa in the abdominal cuticle of *Rhodnius* to 10 or more GPa in well-tanned isotropic cuticles and 20 GPa measured parallel to the chitin orientation in the tibial flexor apodeme of the locust.

The cuticle of the cutting and grinding surfaces of the mandibles of plant-eating insects is reinforced by impregnation with heavy metals (form unknown) such as Zn, Mn, or occasionally Fe. These metals are present up to 16% of dry mass of the and treble their hardness from about 25 to 80 kg mm^{-2}. Elastic energy storage is important in the locomotion of insects for saving energy from one locomotory cycle to the next (e.g. resilin in the hinges of flapping wings) or as a power amplifier releasing energy more quickly than a muscle can (e.g. the flea). Some insects such as locusts and flea beetles use sclerotised cuticular tendons to store elastic energy.

An alternative method is to reinforce the fibres with minerals such as in bone to impart high stiffness. Mineralisation by precipitation of minerals can be achieved in a number of ways [Calvert 1994]. Seashells are constructed from close packed layers of calcium carbonate multi-sided tiles interspersed with a compliant polymer glue between the layers – crack propagation is diffused by the complex pathways imposed by the seashell structure [Sellinger et al 1998]. Nacre (mother of pearl) comprises mineral plates interleaved with protein sheets [Clegg 1990]. Bone offers structural adaptability to imposed loads by virtue of their dynamic equilibrium between bone loss and growth. Most biological materials are self-repairing by virtue of their continuous dynamics, e.g. bone is maintained by a continuous destruction and generation of bone material. Bone comprises connective tissue stiffened by the deposition of minerals. The connective tissue comprises collagen fibres within a polysaccharide matrix into which are deposited crystalline salts of calcium (phosphate and carbonate), the commonest combination being dahlite, $CaCO_3.2Ca_3(PO_4)_2$. Calcification depends on the concentration of Ca and PO_4 at the bone face – blood plasma contains ~2.1-2.6 mmol Ca/l of which 50% is ionised and 50% is bound to protein (such as albumen). Calcitonin is a hormone secreted by the thyroid, parathyroid and thymus which causes the deposition of Ca and PO_4 in bone. The function of bone is to support the body. It is composed of an extracellular matrix of glycoproteins and collagen fibres with a mineral component of Ca hydroxapatite crystals. Bone is continuously formed and resorbed by metabolic activity controlled by bone cells. Osteoblasts are cells which secrete the organic components of bone for bone growth while osteoclasts are cells which secrete digestive lysosomal enzymes for the dissolution of bone. Homeostasis maintains bone mass constant but as the body ages, bone resorption gradually exceeds bone growth (osteoporosis). Plasma Ca concentration controls bone cell activity as well as being influenced by indirect control by hormones especially parathyroid and calcitonin. When plasma Ca levels fall, parathyroid hormone (PTH) is secreted from parathyroid gland to promote Ca and phosphate absorption by the osteocytic membrane from bone fluid. If plasma Ca levels rise, calcitonin is secreted by the thyroid gland to reduce bone resorption by osteoclasts. These nanocomposites indicate a general approach to shatter-resistant, tough, lightweight structures. Such approaches are adopted in impact-resistant structures such as Chobham armour which comprises layers of ceramic shields interspersed with layers of Kevlar fibres. A variant on this technology is a third generation molecular composite comprised of polymeric macromolecules [Palsula 1993]. The production of such fibres may emulate spider silk production [Vollrath & Knight

1999]. This allows a hybrid approach of utilising artificial materials manufactured using chemical engineering techniques with biological components [Ahmad & Mark 1988].

The chief defining characteristic of biological materials is their integrative multi-functionality by virtue of microstructure. Many biological materials exhibit holes which can have well-defined advantages. If the holes have angular outlines or are too close together, cracks can start from them which can spread throughout the material. This is because they can concentrate loads into a small zone, over-loading the structure of the material in a localised area. In a bending beam, however, the main resistance to the load is from the material arranged around the outside. The material in the middle carries relatively little of the load. Therefore there is less need of material and some of it can be safely omitted. We either put a hole down the middle or replace it with a foam or cellular structure. In doing, so we make the structure lighter and cheaper. The stems of plants, particularly the non-woody ones, are designed like this, as are fungi, the spines of the hedgehog, the quills of the porcupine, most of the bones in the body, antler bone, the shell and spines of the sea urchin, the light filling of bird's feathers and, of course, trees. Most natural materials are made of fibres in a matrix. The fibres can then be so arranged that loads are guided around the edges of the holes, eliminating any stress concentrating effects. The foam-like structures down the centre of hedgehog spines support the thin outer walls against local buckling, allowing the structure to bend further without failing. In order to raise bending stiffness, one can either increase the thickness of the wall of the tube (using more material), or one can put more of the material on the outside, increasing its radius. The further away from the plane of bending the material is placed, the more it will contribute to the bending stiffness. Any tube then runs into the problem of ovalisation at the point of highest force and of local buckling. The foam core alone can resist the ovalisation, but local buckling demands radially orientated material for its resistance. This reinforcement is best provided as longitudinal, circumferential or orthogonally arranged stiffeners. If these stiffeners are sufficiently massive they can also do the job of the foam core, which can be removed with no reduction in mechanical performance. The production of cylindrical metal shells with an integral foam-like or honeycomb-like cellular core means that the excellence of this design, previously confined to nature, can be extended to lightweight tubular struts. Evans and Lakes have independently investigated re-entrant foams which are auxetic – they show a negative Poisson's ratio (Lakes 1993). There are several ways of making such materials, and one way is to take a conventional plastic foam, squash it from every side, then heat the compressed foam to slightly above its softening point and allow it to cool. Such foams are more resilient than normal foams and can be extended further than normal ones when stretched. Skin seems to be an auxetic material (Lees, Vincent and Hillerton 1991) whereby the holes in skin serve to transport blood to the cells which make the skin to provide hydraulic stiffening.

Wood is a natural composite of hollow cells of 0.03 mm diameter constructed from multi-layered walls. The middle layer (S2) is thickest and provides the major load-bearing structure. It consists of cellular fibres ~10-20 nm diameter wound into a helix at an angle of 5-20° from the direction of the grain. The cells are arranged parallel to the grain direction and bonded together in a phenolic matrix. The fracture properties are direction-dependent, being brittle along the grain with fracture energies of 0.1-0.2 kJ/m^2 which results in buckling and debonding, but strong across the grain with fracture energies of 12-30 kJ/m^2 comparable to low-grade steel. Wood is able to store elastic strain energy and absorb shocks due to the presence of holes which provide the means to store elastic strain energy and absorb shocks. Hardwoods (from broad-leaved trees, mostly deciduous in temperate climates) are immediately distinguishable from softwoods (from needle-bearing trees, mostly evergreen) by the complex morphology of their wood. Softwoods are composed of tubes about 0.1 mm diameter. In addition, hardwoods possess large vessels, about 0.5 mm diameter, which provide the main route for the transport of water. In a section taken across the tubes, these vessels appear as large holes in a field of smaller ones. They provide localised low density weak areas around which the wood collapses preferentially. If the vessels are arranged in layers (which happens in some oaks and willow) then the wood will crack fairly easily since the holes represent a weak layer within the material. But if the vessels are arranged uniformly throughout the wood then the damage is spread evenly and the full toughness of the wood can be realised [Hepworth Jeronimidis & Vincent 2001]. Each of the large holes in the wood allows greater local deformation, thus causing a concentration of strain and thus of stress. As each crack grows, it removes strain energy from the surrounding region and so inhibits the further propagation of other starter cracks. So the distributed holes in wood absorb strain energy before it can be transferred to the tip of a growing major crack. However, if there are too many holes in the structure it will be weakened since the cracks only have to join up the holes for the material to break.

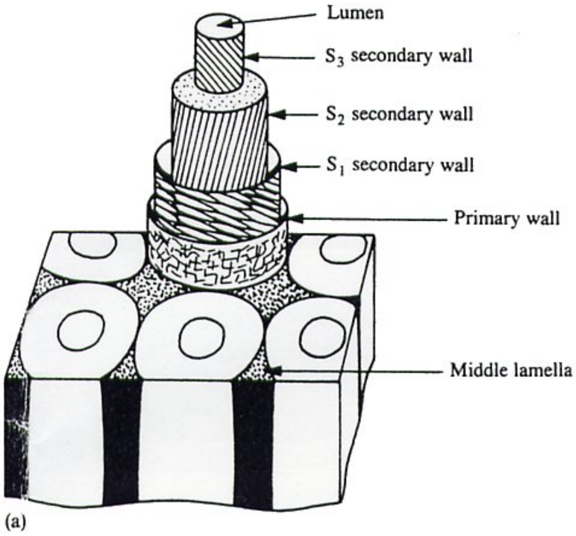

Figure 6 – Hierarchical structure of wood (Jeronimidis & Atkins 1995)

Holes can be used as strain gauges. Arthropods have an external skeleton made of cuticle whose structure is very akin to aerospace fibre/resin composite. The fibre is chitin and the resin is protein, to which chalk and other fillers may be added. They have strain sensors which are basically holes through this skeletal sheath. Spiders have slits for strain sense organs which, assembled into groups, are also known as lyriform (harp-shaped) organs. Insects such as cockroaches and flies have similar, simpler, organs which are basically holes with a domed bell-shaped cap, hence called campaniform sensilla. Crabs and other crustaceans have their own versions, probably all independently evolved. These organs allow the animal to measure displacements of molecular size, and they do so by introducing compliance. The hole which is the organ actually weakens the surrounding cuticle so that it deforms more under a given load. A circular hole will become oval or a slit (the spiders' equivalent) will get narrower or wider depending on the direction of the force. This has the effect of pushing the dome up and down. The main action is therefore to rotate the applied deformation through 90° so that deformation in the plane of the cuticle becomes out-of-plane deformation which can be measured remotely by a cell sitting in the epidermis which produced the cuticle. If the fibre orientations and stiffness of the various components are carefully arranged, the movement of the dome will amplify the deformations as well for easier detection. These organs are always placed in areas where the load is likely to be the greatest - the wing or the leg – which is most likely to break. The holes exist in a sheet of fibrous composite such that the fibres are placed to one side or the other of the hole. As the sheet of material is stretched, the fibres tend to straighten and squeeze the hole, thus increasing its deformation from round to elliptical. But the fibres also tend to carry the load around the hole, so that although the hole concentrates strain, the fibres dissipate stress. These same holes can also serve as crack stoppers, since upon reaching a hole the tip of a crack will be blunted, the stress concentration at its tip will disappear, and the crack will no longer be able to grow. In nature, many of the holes contain, or have contained, living matter which can repair the surrounding material such as bone which can repair itself using the activity of the cells (osteoclasts) which formed it. This concept can be used in many man-made materials, where holes which can stop cracks as well as start them can contain materials which can creep down a crack by capillarity and repair it. Self-healing fibres for composites comprise of hollow porous glass fibres filled with adhesive – the shearing of such fibres by crack propagation releases the resin to fill the cracks and so repair the breaches. Holes are important in insulating animals from the effect of heat and cold. Some desert beetles have a large space or hole beneath their wing covers which serves as a condensation chamber for their exhaled air. Thus they can conserve water more readily. Animals which have fur or feathers rely on the air trapped in the holes in their pelage to insulate them. The air in the holes insulates, but only if it cannot move between the fibres of fur or feather. The trick is therefore to immobilise the air using as little material as possible. The guard hairs in fur are hollow, especially in winter, so that the hairs can both be stiffer and keep the fur layer from being compacted by external forces, and the air inside the hairs provides the insulation. Penguin feathers are particularly good insulators, maintaining a temperature gradient of some 80^o over their thickness of 2 cm. They do this by dividing the air space around the base of the feathers into small holes only 50 μm across, using the thin (5 μm) fibres of the down layer [Dawson et al 1999].

Structures comprising living organisms are hierarchical in nature, e.g. molecular configuration of collagen proteins give way to tertiary structure such as fibres which build up lattices, etc. Most biological materials are self-repairing by virtue of their continuous dynamics, e.g. bone is maintained by a continuous destruction and generation of bone material. The function of bone is to support the body and is composed of an extracellular matrix of glycoproteins and collagen fibres with a mineral component of Ca hydroxapatite crystals. Skeleton is 80% cortical bone and 20% trabecular bone. Cortical bone is predominant in long bones and on the exterior of all bones - it is solid and hard. Trabecular bone is spongy and occurs in the inner regions of bones to allow passage of marrow and red blood cells. Bone is continuously formed and resorbed by metabolic activity controlled by bone cells. Osteoblasts are miniature bone cells which secrete the organic components of bone for bone growth while osteoclasts are cells which secrete digestive lysosome enzymes for dissolution of bone. Homeostasis maintains bone mass constant but as the body ages, bone resorption gradually exceeds bone growth (osteoporosis). Bone fluid within this osteocytic membrane over all bone surfaces contains some 500 mg of Ca in equilibrium, i.e. 0.05% of total bone Ca. Plasma Ca concentration controls bone cell activity as well as being influenced by indirect control by hormones especially parathyroid and calcitonin. When plasma Ca levels fall, parathyroid hormone (PTH) is secreted from parathyroid gland to promote Ca and phosphate absorption by the osteocytic membrane from bone fluid. If plasma Ca levels rise, calcitonin is secreted by the thyroid gland to reduce bone resorption by osteoclasts.

Liquid crystallinity is an example of the mechanisms available for processing and post-processing of materials. For instance, silk is partially aligned before being spun by forming liquid crystal structures. Thus, the energy required for spinning the silk is reduced, since a large part of the molecular orientation associated with spinning has been produced already. This is also an example of nanofabrication, and emphasises the general principle that because biological materials are made "from the molecule up", they are necessarily designed from this level, and have to be assembled into a number of hierarchies. This in itself is probably an advantage, partly since hierarchical structures tend to be more efficient (Lakes 1993), and partly because it is then possible to produce a greater variety of properties by varying the degree of interaction at the interfaces between the various levels of the hierarchy. However, although a hierarchical structure is adaptable, it necessarily introduces more scope for uncontrolled variability. So although the underlying genetic system is extremely conservative, yielding very uniform materials at the molecular level, biological materials in bulk can be rather variable in their properties.

The manufacturing technique most applicable to such structures is shape deposition manufacturing (SDM), a rapid prototyping manufacturing technology in which 3D parts manufacture and assembly occur simultaneously. Material is deposited and shaped in layers with components embedded into the workpiece when required. Different materials (including plastics, metals and ceramics) can be deposited/removed incrementally as required, especially to provide different structural properties such as stiffness or compliance within combinations of stiff and compliant materials. Structures may be fabricated with spatial variation in material properties such as stiffness and damping characteristics. Furthermore, sacrificial support structures may be used as scaffolding to support the manufacturing process allowing the construction of channels and cavities. Sacrificial layers can also encase critical parts during manufacture. The SDM cycle comprises of alternate deposition and machine or cast shaping (for hard and soft materials respectively) of part and sacrificial material layers. Workpieces are built up incrementally. SDM offers the capacity for complex 3D geometry, and engineered compliance and damping within these 3D structures. One example is a 3D plastic exoskeleton housing a distributed shape memory alloy actuator material and embedded spring sensors. Embedding pre-fabricated components adds to the great flexibility of this manufacturing process. Electrical connections and insulation such as aluminium oxide and thermal dissipation paths such as copper may be embedded as well as devices. The embedding of components allows the introduction of distributed sensors such as strain gauges and actuators as well as structural components for passive impedance during manufacture - biological sensing and actuation is typically highly distributed and local. Polyurethane can be manufactured with properties that vary from high stiffness to highly elastic. Actuators can integrate directly into the structure and connected directly to the limb that they are to move, and similarly sensors can be integrated directly into the structure of the limb. This eliminates fasteners and connectors which can comprise up to 40% of the mass of the device. Much of the structure can be reduced by using multi-functional structures for structural support with embedded electronics, thermal paths, sensors, etc similar to physiological tissue. SDM offers the capability of creating multiple materials, sensors, actuators and control within a single monolithic structure. The development of multi-functional structures is a step in this direction which incorporate energy distribution – electrical wiring, thermal piping, optical fibres, etc.

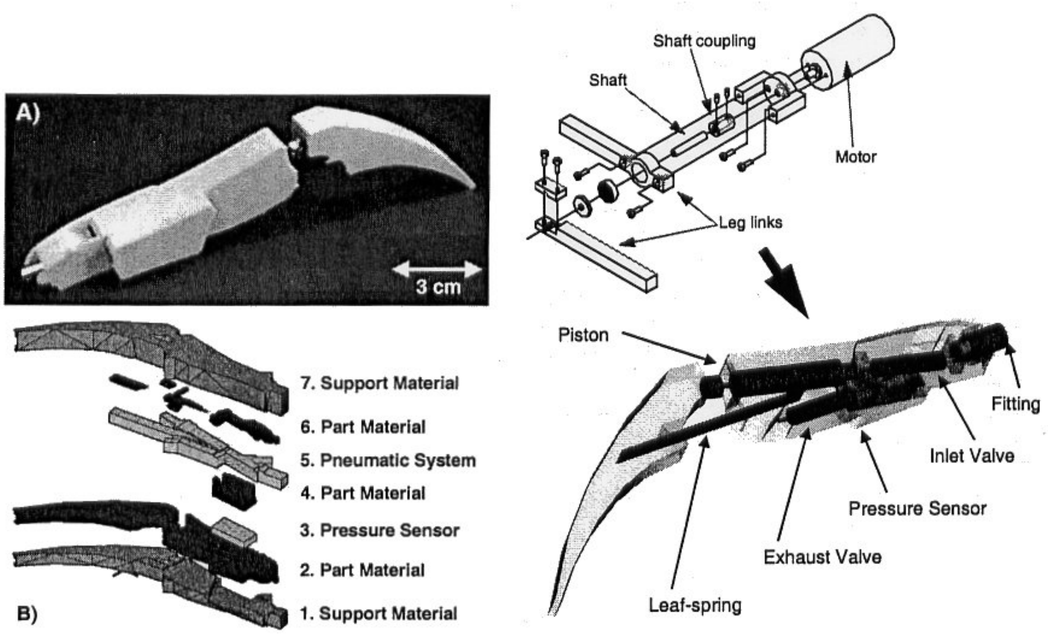

Figure 7 – SDM-manufactured linkage with embedded pneumatic piston, solenoid valves and pressure sensor (Bailey et al 2000)

3.3.1.1. Gecko-Hair Adhesives

Geckos and several other animals such as species of beetles and spiders have the remarkable ability to climb on smooth surfaces, even supporting their own weight on the underside of such surfaces. Such animals are noted for their possession of pads of small hairs (setae) on their feet which are the source of their climbing ability; the hairs range from simple clumps of filaments on beetles' legs to geckos' sophisticated arrangements of flexible flaps (lamellae) on their toes supporting branched hairs about 100μm long and 20μm in diameter, which in turn bear up to 1000 tips (spatulae) about 200nm in diameter [Arzt et al., 2003]. Gecko setae have an adhesive ability of about 200μN per seta [Autumn et al., 2000].

Figure 8 – Tokay Gecko Showing Lamellae on Toes

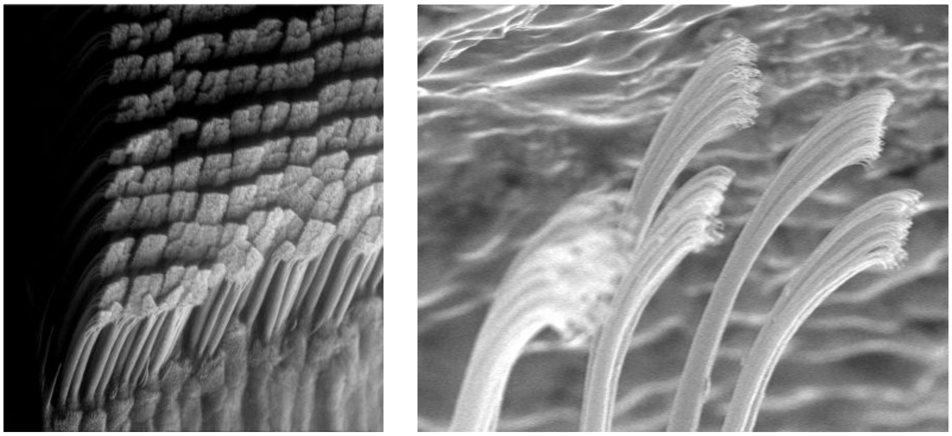

Figure 9 – Close-up of a) Gecko Feet Showing Setae on Lamella and b) Spatulae on Setae

Recent research has shown that the main mechanism of setal adhesion in geckos is van der Waals attraction between the spatulae and the surface with which they are in contact [Autumn et al., 2002]. The force is electrostatic in nature: the arrangement of electrons within one surface can produce electrical dipoles on a molecular level that attract corresponding dipoles in the other surface. Different materials vary in their susceptibility to such dipoles, but in principle the forces are due to the close contact between the surfaces rather than any specific chemical properties. Geckos and similar animals may also take advantage of capillary forces, relying on the surface tension of a thin film of water coating the two surfaces – however, the load that can be supported by this mechanism depends on the hydrophobicity of the material which can vary considerably.

Normally such forces are too weak to be noticeable on a macroscopic level since the area over which two rigid surfaces are in sufficiently close contact is, counter-intuitively, much smaller than the apparent (or "nominal") contact area. However, since a gecko's spatulae are mounted on springy, compliant hairs, a very large fraction of its 2 million setae will be in contact with a surface at one time and so its feet are able to support a significant load; following the best principles of overengineering of safety-critical systems, a gecko is able to support its body weight (~100g) with only one toe in contact with a surface. This strong adhesion could present a problem for locomotion; to get around it, geckos detach their feet by peeling their toes off the substrate, and then uncurl them back onto the surface to reattach.

Gecko-hair type adhesives have a number of claimed advantages over conventional glues:

- They are removable and repositionable.
- They are "dry" adhesives, in that no liquid substances are involved and no residue is left on the surface when the adhesive is removed.
- They stick to almost every type of surface, albeit to differing degrees.
- They work in any environment, including in a vacuum or underwater.
- They are self-cleaning, though it is not clear at present how this is accomplished.

This has led to several efforts to try to produce synthetic gecko setae and spatulae. The group of American researchers that first investigated the mechanism of gecko adhesion has attempted to fabricate synthetic gecko setae by several methods [Sitti & Fearing, 2002]: the first involved indenting a wax surface with a conical atomic force microscope (AFM) probe, and then casting polyester resin or silicone rubber into the resulting moulds. When the layer of polymer was peeled off the wax, 2-3µm tall nano-hairs with a base diameter of around 10µm were produced on the surface of the polymer. An individual hair could adhere with a force of about 180nN, but investigation of the hairs' macroscopic properties was not carried out. The second method involved casting a polymer (silicone rubber, polyamide or polyester) into the holes in a porous membrane (porous alumina made by a self-assembly process or polycarbonate that had been perforated by nuclear bombardment) and etching away or peeling off the mould to leave the hairs. The extreme aspect ratio of the pores in the alumina (60µm long by 200nm diameter) meant that the hairs clumped together through the van der Waals forces the researchers were trying to exploit for adhesion to a substrate; the silicone hairs produced with the polycarbonate membrane were about 6µm in length and diameter, and were able to exert an adhesive force of about 2.8mN/cm² on a glass slide.

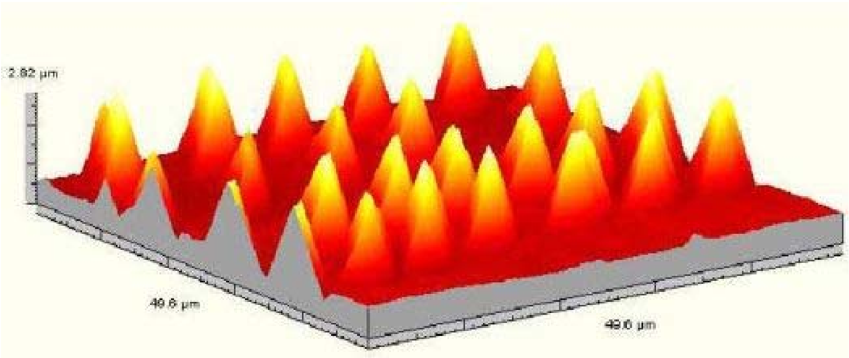

Figure 10 – An AFM image of silicone nanohairs moulded with a wax surface indented by an AFM probe

This work was extended by Sitti (2003) at Carnegie Mellon University. He was able to use the polycarbonate membrane/silicone rubber nanomoulding technique outlined above to produce nano-hairs 5µm in diameter by 10µm long and 0.6µm in diameter by 5µm long. A new technique was also used: the surface of a thin film of liquid polymer was subjected to a vertical electric field between two parallel metal plates, causing filaments to grow upwards and attach to the top plate. By moving the top plate vertically and horizontally after partially curing the polymer, the hairs were lengthened and aligned before final curing. However, macroscopic adhesive properties are not reported for either of these methods.

The most successful attempt was made by Geim et al. (2003) at the University of Manchester, who fabricated a 1cm² array of polyimide pillars 0.5 µm in diameter by 2 µm long. They prepared a number of aluminium discs on the surface of a 5 µm thick polyimide film; the discs then acted as a mask to protect columns of underlying polyimide while the surrounding plastic was etched by a dry oxygen plasma. The resulting film, when attached to a flexible backing (sellotape), was initially able to support 3N/cm²; however it gradually lost its adhesive ability due to the pillars breaking off the backing and tangling together.

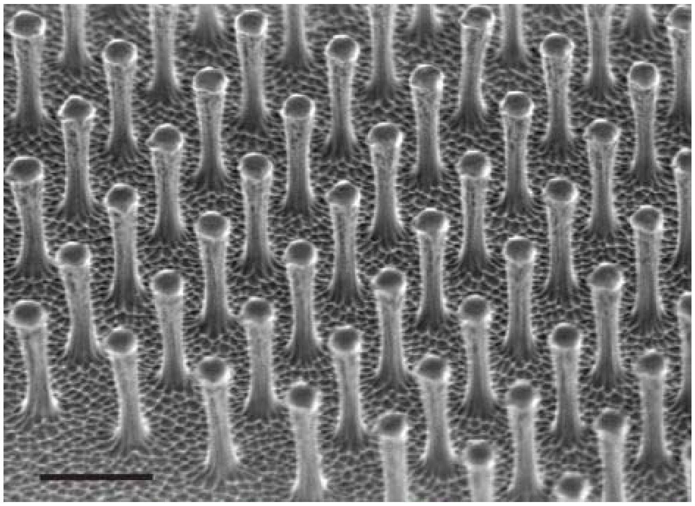

Figure 11 – Polyimide Gecko Hair Mimicking Pillars showing a 2 Micron Scale Bar

3.3.2. Adaptive Materials

Smart materials are modelled on biological systems with sensors acting as a perceptual processing system, actuators acting like muscles and real-time microprocessor based control system acting like a nervous system [Vincent 1992]. A definition of a smart or adaptive structure is: "a structural system whose geometric and inherent characteristics can be changed beneficially to meet mission requirements either through remote commands or automatically in response to external stimulation" [Wada 1990]. Adaptive structures generally utilise strain as the variable parameter – as acceleration and strain are the only sensors which do not require external reference (both can be realised with silicon cantilever structure sensors), this makes these sensing modalities ideal for intelligent structures [Crawley 1994].

Table 3 – Comparison of Strain Sensors (Crawley 1994)

	Foil	**Semiconductor**	**Fibre Optic**	**Piezo-polymer**	**Piezo-ceramic**
Sensitivity	30 V/ε	1000 V/ε	10^6 V/ε	10^4 °/ε	2x10^4 V/ε
Localisation/cm	0.008	0.03	~0.04	<0.04	<0.04
Bandwidth (Hz)	0	0	0	0.1-10^9	0.1-10^9

Smart/adaptive structures/materials come in three main types: passive smart structures, reactive smart structures and intelligent structures [Davidson 1992]. Passive smart structures often employ structurally integrated optical micro-sensors to determine the state of the structure. Reactive smart structures use embedded distributed optical microsensors and microactuators to effect a change in stiffness, shape or other aspect of the structure. Intelligent structures will be capable of adaptive learning. The key is that smart materials offer rapid response capability to a dynamic environment. Smart materials are in fact an extension of composite materials whereby the sensors and actuators are embedded in the material. Autonomous self-adaptation enables optimal performance over a range of operational conditions in an unstructured environment. Future developments will concern the implementation of control system for flexible behaviour.

The use of materials as sensors only provides the first level of smart materials. Sensing provides feedback information about the health of the agent through internal sensors and about the state of the environment through external sensors. Sensors and actuators comprise the input/output of the control system as a whole, and provide the interface for the digital electronic control system to the analogue non-electronic world. They provide the means for coupling to the environment. The speed of sensory processing is critical as it essentially limits the rate at which the control system can operate in response to sensory data and actuate to alter that environment. The three most important sensory measurements are displacement, temperature, and radiation. Displacement measurement is of fundamental importance as it provides the basis for velocity, acceleration, strain, force and pressure measurements sometimes in conjunction with time measurements by piezoelectric quartz oscillators and/or Schmidt trigger. There exist materials that react to external stimuli such as electric current/voltage, temperature, mechanical stress, electromagnetic flux, etc. Such smart materials perform functions dependent on environmental changes. Sensor properties are determined by their linearity, sensitivity, accuracy, resolution and frequency response. As sensors vary in terms of range, directionality, sensitivity, resolution, specificity and accuracy, a suite of sensors is usually required.

The second level of smart structures includes simple binary actuation states that switch according to a threshold level of stimulus, eg. shape memory alloys. Sensing of the environment for perceptual information is central to effective control and actuation. Sensors provide monitoring of task execution and the state of the external world in which the task takes place. An actuator may be defined as a device that transforms electrical or mechanical energy into controllable motion. Traditional approaches to actuation are based on either the electric motor for rotary movement or linear actuator for translational motion. Almost all engines are in fact variants on the slider-crank four bar linkage which gives the engine its reciprocating mechanism.

The third level of smart materials is characterised by continuous actuation in response to continuous sensory signals. Such intelligent structures will require microprocessor-based control of large numbers of instruments to generate flexible behaviour on the basis of the ambient conditions such as for counteracting destabilising effects. This effectively offers the integration of control avionics into the structures of aerospace systems. Critical issues for the integration of sensors, actuators and circuits include the maintenance of the structural integrity of the composite with the impregnation of the sensor-actuator package. They often require different dimensions to reinforcement fibres particularly in terms of diameter.

Actuation mechanisms vary but they generally involve changes in stiffness in response to environmental conditions without the use of moving parts: shape memory alloys, electrorheological fluids and piezoelectric solids. Desirable characteristics of adaptive materials include vibration suppression (shape memory alloys/magnetic shape memory alloys), self-healing materials (ionomers), tactile surfaces (piezolelectric/magnetostrictive/electrostrictive materials), multifunctional structures with embedded molecular wires and optic fibres, structures that can alter their viscosity (electrorheological/magnetorheological fluids), etc. We consider electrically actuated materials over their magnetically actuated counterparts due to their greater utility and controllability by the application of dc – the application of magnetic fields can corrupt devices such as magnetometers and magnetorquers.

Table 4 – Characteristics of Smart Actuators (Dario et al 1992)

Actuation Mechanism	Energy density (J/cm³)	Drive conditions	Scalability	Power/weight ratio
Electrostatic	0.4	E=300 V/μm	Good	High

Electromagnetic	0.9	B=1.5 T	Poor	Low
Shape memory alloy	10.4	P=1.4 W/mm^3	Good	High
Piezoelectric	5x10^{-4}	E=30 V/μm	Good	High

3.3.2.1. Composite MEMS Materials

Composites are laminated structures that are uniquely suited to the incorporation of sensor and actuation components [Crawley 1994]. Smart materials have been exploited within MEMS technology as they are a type of micro-machine (which perform mechanical work), particularly piezoelectric materials and, more unconventionally, shape memory alloys. Miniaturisation of sensors, actuators and microelectronics to micron-size offer reduced mass, volume (and so launch costs) and power requirements, and the potential for batch processing and mass production of common off-the-shelf integrated microsystems with increased reliability. MEMS are devices with physical dimensions of less than 1 mm which operate on mechanical principles [Angell et 1983]. MEMS technology encapsulates the notion of adding mechanical structures such as beams, diaphragms, gears, springs, etc to electronic integrated circuitry to add interfacing to the world through sensing and actuation. MEMS technology is premised on:

- micro-scale geometries (miniaturisation)
- precise dimensional control of etching techniques
- direct interfacing to control electronics (microelectronics)
- high repeatability and reliability (multiplicity)

The commonest MEMS devices have been based on sensors which exploit a flexible diaphragm which deforms due to a pressure differential across it. Such a device may be integrated with on-chip, bipolar op-amp and thin film resistance circuitry for signal conditioning. Other common sensors include accelerometers based on inertial mass suspended by compliant springs – the deflection of the inertial mass changes the capacitance between a cantilever beam and reference beam. Both acceleration and pressure sensing in MEMS devices can operate by either by capacitance sensing (most commonly used in accelerometers) or by piezoresistive detection (most commonly used in pressure sensors).

The development of Si microactuators $\sim \mu m^2 - mm^2$ are a logical extension of the more established microsensor technology. Such microactuators are mechanical devices that can allow motion control over very small distances ~0.01-10μ. The addition of actuation introduces a number of problems which are dominated by the need for microactuators with high power output to mass and high torque to speed ratios [Dario et al 1992]. There are ten actuation methods of transforming energy into motion [Gilbertson & Bisch 1996]. Electromagnetic devices include electric motors, solenoids and relays and micro-electromagnetic motors are capable of generating magnetic flux densities of 1.0 T. They are capable of smooth motion due to the continuous nature of electromagnetic fields (as opposed to electrostatic fields which are discontinuous). Electromagnetic micro-motors offer higher output torques but require large currents. They exhibit fast operational speeds with high efficiency, but require perpendicularity between current conductor windings and the moving magnetic elements making it difficult to manufacture in planar silicon. Electrostatic devices exert great forces over very short distances with low current consumption. They are readily fabricated on silicon substrates, large electric fields are readily generated through capacitance, and high rotation speeds can be achieved with ~10^3 Hz cycle rates. They exploit coulombic force generated when a voltage is applied across a pair of conductive plates separated by an insulator, which scales down to micron scales. The chief difficulty at such small scales is the friction in sliding surfaces scaling as L^2 while inertial forces scale as L^3 such that frictional forces are dominant over inertial forces at small scales. Furthermore, electrostatic micro-motors above 0.1 mm in size does not offer good performance. Electrostatic linear and rotary motors and gear trains have been successfully fabricated. Electrostatic actuators are suited to on-chip actuation (such as positioning of devices, heads, mirrors, sensors, etc), eg. the Ni electrostatic micro-motor with a 120 μm diameter, 7 μm thick disk rotor which can rotate at 10,000 rpm. Micro-grippers become feasible with micro-motors, particularly based on electromagnetic solenoids. A variation on this is the use of magnetostrictive alloys as the actuating element in response to applied magnetic fields. The piezoelectric ultrasonic motor offers high torque output with great potential but have yet to be miniaturised to micrometer scales.

Table 5 – Different Micro-motor Principles (Johansson 1995)

Actuation mechanism	**Torque (μNm)**	**Power (μW)**	**Comments**
Electrostatic	0.001	0.5	Suffers from friction problems
Electromagnetic	0.04	30	High current requirements
Piezoelectric	0.0004	0.03	Good performance for larger motors
Inchworm	1.6	36	Offers high torques

Johansson (1995) suggests that the linear inchworm represents a universal microactuator for micro-robotics offering high force and stroke. These micro-motors are based on piezoelectric devices to control position rather than generating resonant vibrations as in the ultrasonic motor. The inchworm comprises both gripping and moving actuators. There are two sets of opposing linear piezoelectric actuators around a central rotor. Considering one pair of piezoelectric actuators: two opposing piezoelectric elements are powered to grip the rotor, and the rotor is turned by contracting the piezoelectric elements longitudinally. The two piezoelectric elements release the rotor returning to their initial positions.

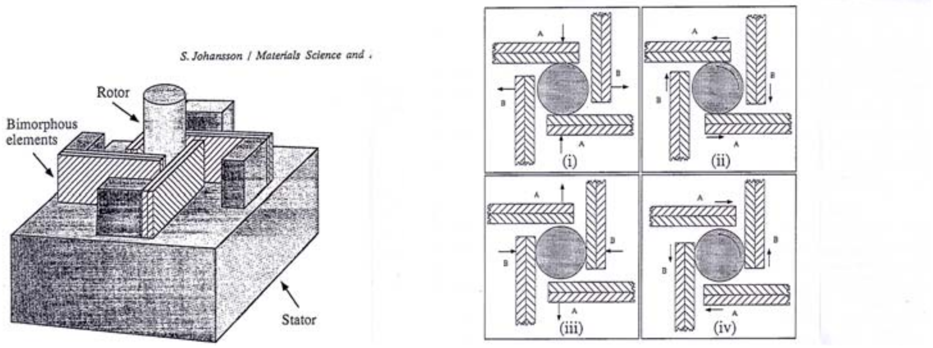

Figure 12 – Inchworm actuator mechanism (Johansson 1995)

There are a number of processing techniques used in IC manufacture, which tailor materials by doping with impurities through thermal diffusion or ion implantation of silicon-based materials. Photolithography is used to transfer a pattern from a mask to a film with a photosensitive chemical (photoresist). The mask comprises a glass plate coated with patterned Cr film designed by CAD software. The photoresist is coated onto the wafer and exposed through the mask. The film is then selectively etched, leaving the desired pattern on the film. Wet etching uses liquid reactants but is becoming superseded by dry etching. The commonest dry etching process is reactive ion etching, which employs a plasma of reactant gases at low pressure. To deposit SiO_2, metal, polysilicon, and Si_3N_4 films onto substrates other than Si, chemical vapour deposition (CVD) in which a heated gaseous mixture is used to deposit the film. Epitaxial growth is a special type of CVD in which a film (such as GaAs and SiC) is grown on a single crystal substrate. For the formation of 3D structures, bulk micromachining employs anistropic etchants of Si such as KOH and N_2H_4 (hydrazine) with SiO_2, Si_3N_4 and metals as etch masks. KOH etches Si <100> and <110> crystal planes much faster than the <111> crystal plane. Deep anisotropic dry etching using reactive gas plasmas can etch Si vertically. Surface micromachining involves using sacrificial layers which are generally subsequently dissolved using HF. In this way, 3D structures of high complexity can be created. For instance, NiTi shape memory alloy film has been fabricated within MEMS devices to provide thermally actuated microvalves.

Sensors that may be embedded into the composite material during manufacture may be optical waveguides or microelectronic devices of which fibre optic systems appear to be the most promising due to their capability to withstand the strains of the materials manufacturing process. They offer the possibility of distributed continuous monitoring of the interior of the composite at all stages of the structure's life as they are stretched or compressed, i.e. structures with optical fibres. They can also serve to reinforce the composite. Coatings must be temperature resistant, eg. polyimide is stable up to 200^oC. Diamond fibre yields a 50% increased stiffness than SiC fibres with a lightweight material. Polyimide coatings may minimise the degradation effects of integration on mechanical integrity. It is unlikely that such integration could extend from polymeric composite to metal matrix composites due to their high temperature processing requirements unless sapphire fibres can be developed. Fibre optic sensors can detect a range of physical parameters through distortion of the propagation path of the optical signal: temperature, strain, pressure, electric and magnetic fields by the interaction of the parameter to produce a modulation of amplitude, phase or polarisation of a transmitted signal. Interferometric phase sensors detect the influence of physical parameters on the phase of coherent laser light in a single mode fibre with respect to reference arm fibres. The sensed information is integrated strain and temperature information. Phase shifts are caused by a change in the strain and temperature state in the fibre and may be detected as fringe shifts in an interferometric pattern. If a fibre is operated at a wavelength less than the cut off wavelength the fibre may be used as a modal interferometer to detect intensity changes due to strain or temperature variations. Polarimetric sensors in a fibre can also be used to detect strain/temperature changes. By incorporating elliptical cores orthogonal polarisation modes will propagate at different velocities generating a phase difference as a function of strain and/or temperature. Dual core optical fibres operating at two different wavelengths offer the possibility of separating strain and temperature information. To make measurements as a function of position along the fibre length the fibre is sensitised at discrete points along the length to create larger signals from each sensing element by for example Bragg reflection or partially reflective splices. One molecule thick Langmuir-Blodgett films such as organic dyes may be doped with rare metal ions to create light emitting layers. By changing the film thickness or molecular orientations the optical properties such as wavelength of light transmission can be altered. Optical neural networks of simple non-linear processing elements capable of adaptive learning offer great promise for direct integration into composite materials [Coghlan 1992]. Such techniques may be implemented with new, 3rd generation molecular composites that are constructed from flexible coil polymer macromolecules, reinforced by molecularly dispersed rigid rod macromolecules [Palsula 1993]. The molecular composite is homogeneous in that the molecules are distributed isotropically. The molecular composite retains the thermomechanical properties of the matrix while attaining the mechanical strength of the rigid polymer in all directions. It eliminates the fibre/matrix interface, which can suffer from poor adhesion. Applications include vibroacoustically stable structures, which utilise their passive damping characteristics of the high strength rigid molecule.

Adaptive materials are likely to contribute substantially to these approaches to reliability specifically with regard to structure for in-situ self-diagnosis and repair. Failures may be prevented by the ability to monitor strain, fatigue, damage and delaminations. Optical fibres can continuously monitor composite structures at all stages of their life from manufacture, testing and operational lifetime (health monitoring and biometrics). Sensors must be able to withstand the manufacturing environment without degradation. Self-healing fibres have been developed for composites consisting of hollow porous glass fibres filled with adhesive. Shearing of such fibres by crack propagation releases the resin to fill the cracks. Corrosion may be offset by wrapping plastic based fibres around metal reinforcement fibres. Changes in acidity dissolve the fibre coatings releasing chemicals to halt the decay.

3.3.2.1.1. Polymer-based MEMS Actuator

The development work on polymer-based electrostatic actuator arrays that have many desirable muscle-like properties and potentially can meet the requirements of many biomimetic robots. For a typical robot, locomotion actuators require the largest forces and displacements, as they are required to move the entire vehicle with each step/flap/etc. However, power consumption during actuation must be held to a minimum, as batteries can become a significant fraction of the total robot weight. During actuation, though not when moving, the power consumption should be close to zero, to prevent additional power loss for holding a position or stance. Manipulator actuators may require much less force, though more finesse and likely a minimum of space and power.

Some actuator examples used in biomimetic robotics are electromagnetic servos, pneumatic actuators, and shape memory alloys. SMAs may prove to be the best option for biomimetics, as the others provide low power-to-weight ratios and require high input power or a high pressure system. SMAs are small and can easily be controlled with a minimum of voltage. Although not as much as the others, SMAs still use a lot of power to work, as electricity is used to create heat to expand and contrast the SMA, even when actuated but not moving.

Polymer-based actuator has the potential to meet most of the above requirements for a biomimetic robot. The polyMEMS actuator (PMA) is composed of micro- to milliscale unit cells configured in a three-dimensional array. A unit cell is fabricated from two thin, flexible, polymer sheets, bonded together at localized points. When an external force is applied to the cell, it opens like a spring. When a voltage is then applied to metal and dielectric films deposited on the inner surfaces of the cell, the electrostatic force causes the unit cell to completely collapse (see Figure 13). Additional cells added in the same plane increase the force of the system, while additional cells added vertically could increase displacement. Therefore, a large 3D array could be built to generate large forces and large displacements simultaneously. Customized layouts delivering specific forces/displacements at different points throughout the array could also be created, as could different shapes (cylindrical, prismatic, helical, toroidal, etc.)

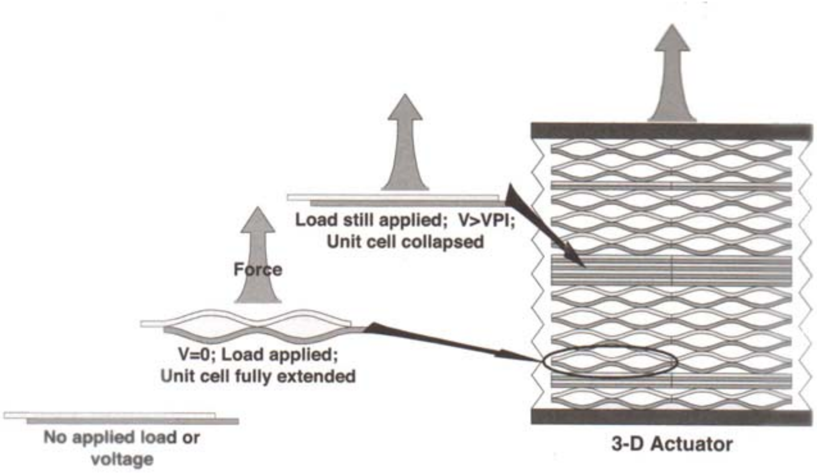

Figure 13 – PolyMEMS Actuator (PMA) Concept

Comparison with muscle

Muscle is composed of unit cells of myosin and actin protein filaments, arranged in a 3-D array. Although structurally different from the unit cells in the polyMEMS actuator, the multiplicative concept is the same. Similar to PMAs, the muscle is a collection of small units working together and the failure of a single or small group has minimal effect on the system as a whole. Both also work in tension only, as electrostatic force is always attractive.

PMA Unit Cell Design

A more detailed diagram of a unit cell is shown in Figure 14. Initially, the sheets are flat together with no forces applied to it. When a force is applied, they are pulled apart as shown below in the figure. Then, when voltage is applied, they are brought back together due to the force of the electromagnetic field.

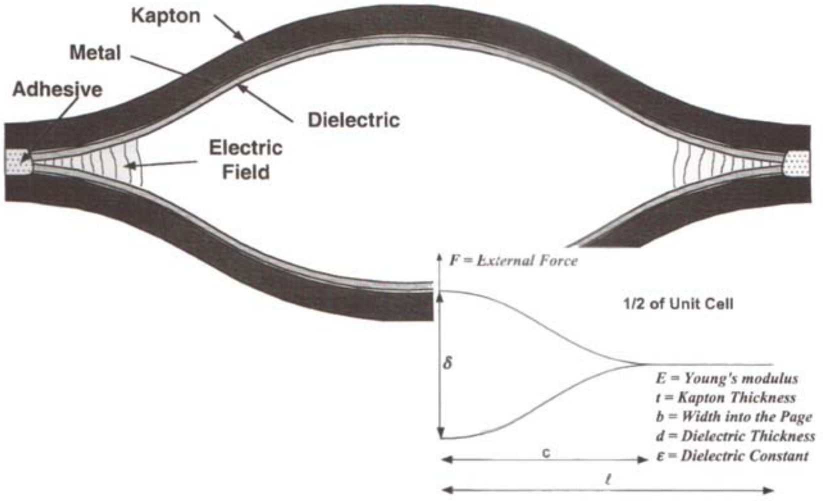

Figure 14 – Close-up look at a PMA unit cell

Both analytically and through finite element analysis, the unit cell can be modelled. The bending energy in the curved region that must be overcome by voltage is:

$$U_B = \frac{1}{2} kb\left(\frac{l}{c}\right)^3 \delta^2$$

where:

$$k = \frac{Et^3}{2l^3}$$

The parameters l, c, b, δ, E, and t are defined for half unit cell in the inset of Figure 14. Note that this is simply the energy of a spring with spring constant $kb(l/c)^3$.

The electrostatic energy is nonzero only in the flat, clamped portion of the profile and is therefore just the energy stored in a parallel plate capacitor at voltage V:

$$U_E = \frac{\varepsilon_0 \varepsilon b(l-c)}{4d} V^2$$

where:

d = dielectric thickness
ε = relative dielectric constant

The pull-in voltage is the voltage needed to start the compression process and is expressed as a function of the material properties, unit cell dimensions, and the constant external force, F:

$$V_{PI} = \left(\frac{6d}{kl\varepsilon_0 \varepsilon} \right)^{1/2} \frac{F}{b}$$

The strain is the ratio of the displacement, δ, to the total thickness of the open cell, δ + 2t. Therefore:

$$Strain = \frac{\delta}{\delta + 2t} = \frac{F}{F + 2kbt}$$

In the absence of voltage, $c = l$ and the unit cell is a spring with the spring constant kb. An external force, F, will extend the actuator a distance $\delta = F / kb$. When a small voltage is applied, almost nothing happens, however, when the pull-in voltage is reached, the unit cell suddenly collapses.

This model can also be used to determine the best choices with respect to unit cell length, polymer thickness, and dielectric thickness and constant when designing a PMA for a biomimetic robot.

The polyMEMS actuator is fabricated using a simple process adapted from flexible printed circuit manufacturing. Polyimide was chosen as the base polymer because of its excellent thermal and mechanical stability (and it is widely used in flexible circuit material), though other options can be substituted. Aluminium is coated with a dielectric film to prevent shorting between adjacent metal electrode films.

Results

Development and testing has begun and arrays have been created with unit cell lengths between 3 and 12.5 mm. The goal is to reduce the size to 1 mm unit cells. Problems have arisen with fracturing of some unit cells under high electric fields, a problem that can be reduced substantially using thicker dielectric films. Thicker films require a higher pull-in voltage, which imposes a very minimal penalty on the system drive electronics as long as the voltages remain below ~250V. The final issue is moisture absorption in the polyimide, which can cause swelling and film delamination. Barrier films and hermetic packaging should be able to control this in most terrestrial or aerial applications, though special packaging may be required for robots that are continually exposed to water or high moisture environments (Horning & Johnson 2002).

3.3.2.2. Shape Memory Alloys (SMA)

Shape memory alloys exhibit thermoelastic behaviour changes such that they can exist in two temperature-dependent crystal phases separated by phase change transitions. The commonest shape memory alloy is the NiTi alloy known as Nitinol (others include Cu-Zn-Al and Cu-Al-Ni alloys). The first such metal alloy discovered was Cu-Zn in 1938, but subsequent emphasis has been on Nitinol (Nickel-Titanium Naval Ordnance Laboratory). When plastically deformed in the low temperature thermoelastic martensitic phase, the alloy will return to its original shape reversibly by heating above a characteristic temperature to the austenitic phase (which is tailorable between 0-100°C) and if restrained from regaining its memory shape will induce stresses of ~700 MPa. Shape change can occur over very small temperature ranges through the transition temperature. The transformation temperature range can be varied by adjusting the amounts of the constituent alloys.

The martensite phase exists at the lower temperature and the austenite phase at the higher temperature. On heating from the martensite phase, the alloy is transformed into the austenite phase which starts at the martensite start temperature A_s and is completely transformed at the martensite finish temperature A_f. If the austenite phase is cooled, it is transformed back into the martensite phase, which starts at the martensite start temperature M_s and is completed at the martensite finish temperature M_f. The temperature range A_s to A_f is higher than that for M_s to M_f. Typically, these temperatures are in the region of 40-50^{o}C. Thus, SMA exhibits hysteresis with a 30-40^o temperature differential typically, though this can be tailored. NiTi has essentially three forms: martensite, superelastic and martensite phases. In the martensite form, the material is soft and ductile due to its complex rhombic crystal structure; in the superelastic phase, it is highly elastic; and in the austenite phase, it is strong and hard due to its simple cubic crystal structure. The superelastic phase is a property of the material from just above A_s to M_d where $M_d > A_f$ which defines the maximum temperature of superelasticity.

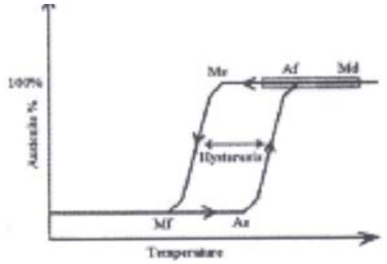

Figure 15 – Shape memory alloy hysteresis (Mihalcz 2001)

Shape memory alloys are limited to ~2 Hz frequency oscillations and ~4-8% extensions but are capable of generating high forces (but are limited to 2% recovery strains for deformation/recovery cycling). The heating and cooling rates of SMA vary as 1/L and this thermal inertia limits their speed of response. They can generate the highest control forces than other actuating materials but have high time constants, $\tau_{\text{heating}} = 1.2$s and $\tau_{\text{cooling}} = 3.5$s being typical:

$$\tau_h \frac{dF(t)}{dt} + F(t) = k_a i(t)$$

$$\tau_c \frac{dF(t)}{dt} + F(t) = k_a i(t)$$

where:

i(t) = input current
F(t) = effective actuating force
k_a = input influence coefficient

If we assume an average time constant of $\tau_0 = \frac{\tau_h + \tau_c}{2}$, we define SMA dynamics by transfer function:

$$\frac{F(s)}{i(s)} = \frac{k_a}{s\tau_0 + 1}$$

However, for small micro-type devices, heating times can be as law as 10-100 ms with low heating power inputs ~10-100 mW.

The shape memory alloy offers large actuator forces and displacements suite to Joule heating as the control input to embedded SMA wires in structures. A generalised Hooke's law applies:

$$\sigma = E\varepsilon' = E\left(\varepsilon - \alpha(T - T_0)\right)$$

where:

E = *Young's modulus*
ε = *stress*
α = *thermal expansion coefficient*

Heat conduction in the SMA is most generally modelled as [Shu et al 1997]:

$$c_v(T)\tfrac{\partial T}{\partial t} = k\tfrac{\partial^2 T}{\partial x^2} + \rho J^2$$

where:

k = *thermal conductivity*
ρ = *electrical resistivity*
J = *current density*

As SMA is based on heating effects, heat transfer within the SMA device must be modelled by:

$$m(c_p + h(T))\dot{T} = \dot{Q} - \lambda \Delta T$$

where:

h(T) = *heat of phase transformation*
$\dot{Q}$ = *electric heating power*
λ = *thermal conductivity*

Generally, Joule effect heating by the application of voltage is exploited for controlling SMA devices with control implemented through the measurement of electrical resistance (rather than temperature). The SMA wire temperature may be modelled by:

$$T = Fo^{-r/Th} e^{-Fo(Bi+\varepsilon)}$$

where:

Fo = *Fourier number*
Th = *Thring number*
Bi = *Biot number*
r = *emission constant*
ε = *strain*

Similarly, the time constant may be estimated by:

$$\tau = \frac{c\rho A^2}{\lambda(Bi+\varepsilon)}$$

where:

c = *specific heat capacity*
A = *geometric cross section*
λ =*gain factor*

The SMA is moulded into its required shape just above the temperature M_s then rapidly cooled below M_f and subsequently heated to above A_f to take up its austenite shape. This procedure is repeated 20-30 times. Superelasticity is characterised by NiTi's ability to return to its original shape after deformation – it is capable of being strained much more than ordinary metal alloys without plastic deformation up to 8% strain but thus behaviour is limited to a specific temperature range A_s to M_d. The optimal superelasticity occurs close to A_f while below A_s NiTi exhibits low elastic modulus. Critical transition temperatures such as A_f can be lowered by substituting Co for NI in 50:50 NiTi. Increases in the hysteresis range can be achieved by adding Nb to decrease the martensite transformation temperatures – however, there may be a temperature dependence on stress capacities, which drop at 4 MPa/°C. NiTi alloy may be manufactured with the required elastic properties at the required operational temperature by variations in the alloy composition, the amount of cold working and the heat treatment process.

There are two companies, namely, Furukawa NT and TiNi Aerospace Ltd which make the Nitinol alloys tailored to the needs of a specific application. Furukawa NT have a range of available Nitinol products manufactured by them. The TiNi Aerospace is the one which designed the NiTi alloy SMA Frangibolts (the bolts used for launch locking and detaching the lander from the orbiter for the EDL) for the Beagle 2. Such materials can be used for precise six degree of freedom positioning and pointing of payloads using the hexapod-like Stewart platform. The US Middeck Active Control Experiment (MACE) flew on STS-67 in 1995 to demonstrate the use of adaptive materials and neural networks for structural control and active vibration suppression [Denoyer et al 2000]. MACE II was a 1.5m flexible structure of Nitinol shape memory alloy with embedded piezoelectric materials in which rigid body motion is controlled by three reaction wheels at the centre of the beam, and two-axis gimbals at each end.

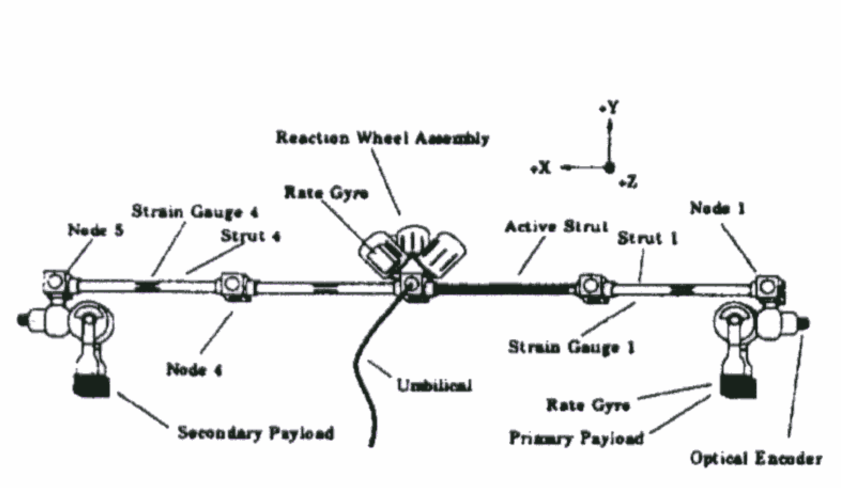

Figure 11- MACE II Test Boom (Denoyer et al 2000)

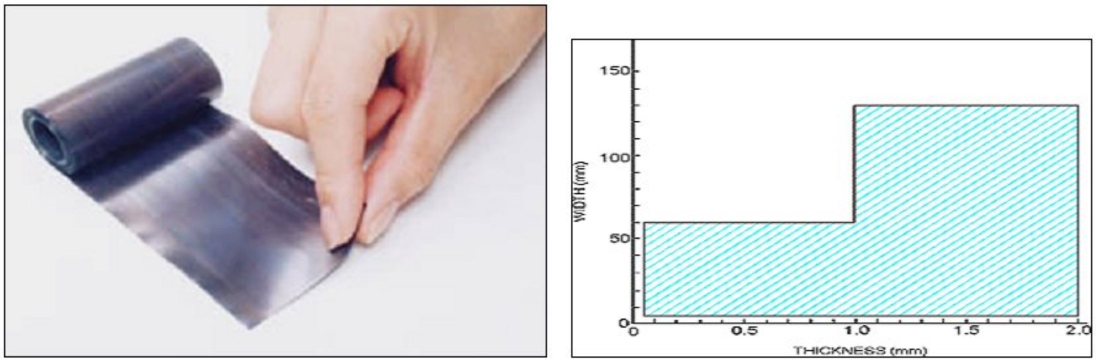

Figure 16 – Shape Memory Alloy Nitinol (a) SMA foil; (b) Width-Thickness Manufacture

Thermocouples may be embedded in composites to alter the temperature of the material but as they are metal conductors, temperature change can be induced by resistance heating. Shape memory alloys may be embedded as 200-400μ plastically elongated wires or springs constrained from recovering their normal memorised length during fabrication. If the fibres are heated for instance by a current, they will generate a uniformly distributed shear load along the fibre to alter the material's modal response and elastic stiffness or if offset from the structure's central axis the structure will deform. Such systems could be used to control vibrations in large composite structures such as space platforms. The fibres may be arranged such that they induce a state of strain and the structure thereby altering the structure's stiffness, natural frequency and modal responses. Other possibilities include latch release mechanisms, fibre optic switching, micro-valves of thin film NiTi on etched silicon (to act as diaphragm of valve for micro-gas chromatography). Shape memory alloys require cooling to undergo reverse transformation and have low efficiency.

Shape memory polymers (SMP), eg. Nippon Zeon polynorborene (NZPN) are similar to SMAs and exhibit the elastic memory effect such that their elastic modulus changes reversibly across the glass transition temperature T_g. Above T_g, the material is flexible but below this temperature, it is stiffer. Currently, SMP have high T_g limiting their applications, though NZPN has T_g = 35-40^{o}C. Thermomechanical devices utilise the expansion/contraction of materials during temperature changes such as bimetallic thermostats. Phase change, thermomechanical and shape memory alloys require heat as the primary driving force and are inefficient with low cycle rates due to difficulties in heat transport.

3.3.2.2.1. SMA Actuators Applied to Biomimetic Underwater Robots

There is the need for an accurate-to-nature biomimetic linear actuator as a muscle replacement. One solution to this need is the use of shape memory alloys (SMAs), specifically nickel-titanium (Ni-Ti, more commonly known as Nitinol).

Nitinol is an alloy material that exhibits the shape memory effect (SME). This property enables materials to recover residual strains of up to 8% upon deformation. SMA materials show two important phenomena:

1. Shape Memory Effect: the specimen exhibits large residual strains and fully recovers during a phase transformation initiated by a temperature rise

2. Pseudo-elastic Effect: the specimen exhibits large residual strains, which are fully recovered upon unloading in a hysteresis loop

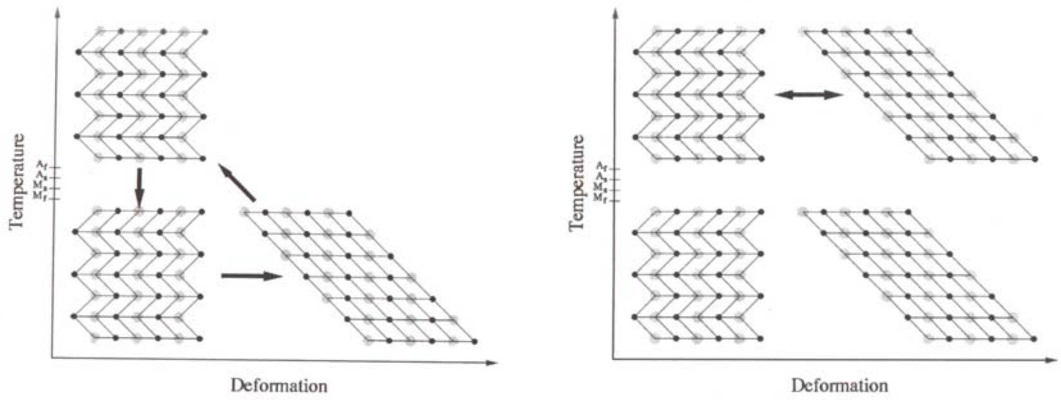

Figure 17 – a) Shape Memory Effect and b) Pseudo-elastic Effect

The materials ability to recover large amounts of plastic strain is due to the martensitic phase transformation initiated in certain conditions of stress and temperature.

SMA materials at zero stress and high temperature exist in the austentite crystalline phase, a body-centred cubic (BCC) structure. When the temperature is decreased, the material undergoes a crystal transformation into the martensite phase, usually a face-centred cubic (FCC) structure. This transformation is usually in the order of: $M_f < M_s < A_s < A_f$ (martensite-finish, martensite-start, austenite-start, austenite-finish).

Nitinol can be used as a artificial muscle through simply heating Ni-Ti wire with electrical current. This electricity/heat brings about the martensitic transformation, thereby producing strain in the wire. This strain is applied to a mechanical system of levers (formed into a robot's leg, for example). Really, this is quite simple.

Construction of this can be achieved with a Nitinol wire (ex: 250 μm Flexinol brand SMA wire by Dynalloy, Inc.). First, the wire needs to be annealized (500^{o}C for 30 minutes) in an argon atmosphere, then acid etched in a mixture of HNO_3 and HF to remove an oxide coating present on the wire's surface (this provides more uniform electrical resistance and creates better electrical connections). Then the wires are run through a PTFE (Teflon) insulating sleeve, providing thermal and electrical insulation from the surrounding water. Stainless steel crimps (302 SS) are then crimped to each end to form a terminal loop of Nitinol on each side of the "muscle". Kevlar stings are attached to the loops to create a means of a mechanical attachment to the robot legs, as shown in Figure 18. Electrical connectors are soldered to the stainless steel crimps then covered with shrink-tubing and potted in urethane sealant to insulate the terminals and keep the Teflon free from water.

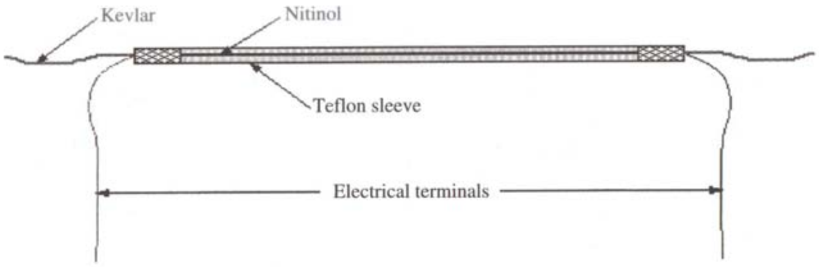

Figure 18 – Nitinol Muscle Schematic

Several constitutive models for SMAs exist, however, due to the complicated pathway for heat dissipation created by the insulations in the water of the lobster robot built in the testing of this muscle, they proved to be a poor predictor of Nitinol behaviour. Therefore a complicated empirical approach was taken involving testing of the Nitinol under various stress/strain and temperature conditions that are not discussed here.

Biomimetic underwater robot posture control with Nitinol actuation

Mimicking the posture states of a *Homarus americanus* (lobster) is important to accurately negotiate the rough surface through high water currents commonly found on the sea floor. Achieving the postural states in the main body pitch, roll, and height planes is accomplished through 3-DOF linkages, each actuated by Nitinol muscle modules.

The onboard controllers are able to provide 3 distinct levels of actuation power, which is analogous to the muscular recruitment levels of the real animal. The full power level corresponds to the power required to produce the maximum contraction of the wire, while the other levels are 1/3 and 2/3 exactly.

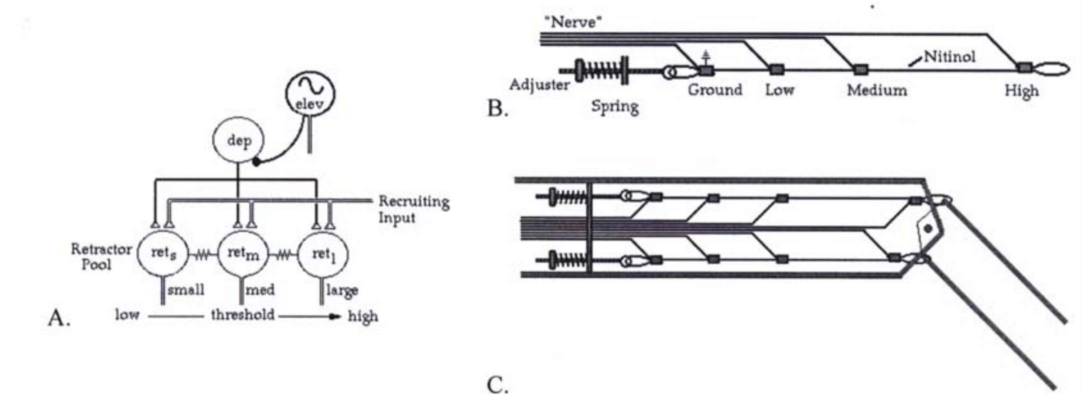

Figure 19 – Neuromuscular recruitment network for NiTi muscles

Although not the ideal material for an artificial muscle, this study shows that Nitinol SMA can be used in the construction of an underwater ambulatory robot. The benefits to this material are very high energy density, compact size, and low weight. However, the major drawbacks include relatively high power consumption and difficulties in controllability. A predictive relationship between the actuation power and muscle displacement must be established to accurately determine the power required to achieve a full compression under varying thermal conditions. The addition of a thermal insulation layer (PTFE tube) has reduced the power consumption to less than 4W per module in warm ambient water.

For any desired posture of the lobster, the best possible set of muscle activations (based on the minimum power consumption criterion) is decided based on the collection of posture possibilities collected through the study. Power consumptions as high as 99 W has been determined for some pitch and roll manoeuvres.

Using an artificial muscle in a biomimetic robot requires the careful matching of the response characteristics of the muscle to the control signals. In this study, Nitinol has shown to fulfil the requirements of the biological model: the *Homarus americanus*. However the high power consumption and heat dissipation limitations essentially exclude the use of Nitinol in terrestrial applications (Witting, Safak, Adams 2002).

3.3.2.3. Piezoelectric (PZE) Materials

Piezoelectric materials are insulators that produce charges on their surface when mechanically deformed. Piezoelectric ceramics and polymers become electrically polarised due to changes in their resistivity when subjected to mechanical stress. Maxwell's differential equation in tensor form describes the free energy function of the piezoelectric crystal which couples the elastic field and the electric field:

$$D = e^T \varepsilon + \varepsilon_r E$$

$$\sigma = c\varepsilon - eE$$

This gives:

$$\varepsilon V = -tD\sigma$$

where:

t = *crystal thickness*
V = *electrode voltage*

and:

$$dD_i = d_{ijk}^T d\sigma_{jk} + \varepsilon_{ij}^{\sigma,T} dE_j + p_i^{\sigma} dT$$

where:

d_{ijk}=*piezoelectric strain coefficient tensor (with 18 components)*
ε_{ij}=*electric permitivity matrix*
p_j=*pyroelectric coefficient vector*
σ = *elastic stress*
D = *elastic displacement*
c = *elasticity*
E = *electric field strength*
T = *temperature*

Similarly, when electric current is passed through them, they undergo mechanical deformation.

The Lagrangian of a piezoelectric body is given by:

$$\Im = \int_V \left(\tfrac{1}{2}\rho\dot{q}^T\dot{q}\right) - \left(\tfrac{1}{2}\varepsilon^T\sigma - E^T D\right).dV$$

The work done due to an external force over the surface S_M and an applied charge over the surface S_E to a piezoelectric body is given by:

$$dW = \int_{S_M} dq^T F.dS - \int_{S_E} d\phi\mu.dS$$

where:

F=surface mechanical load
ϕ*=electric field*

Hamilton's principle allows a description of the change of behaviour over time:

$$\int_{t_1}^{t_2} d(\Im + W).dt = 0$$

Assuming that the piezoelectric is a thin layer, this simplifies to:

$$M\ddot{q} + C\dot{q} + K_{qq}q - K_{q\phi}\phi = F \text{ = mechanical force}$$

$$K_{\phi q}q + K_{\phi\phi} = G \text{ = electric charge excitation}$$

Piezoelectrics are 10-20 times more sensitive than metal film strain gauges and are regarded as part of MEMS technology in which they are exploited in the form of pressure transducers. Pressure sensing is the best developed of such sensors, most of which are based on the piezoelectric effect [Angell et al 1983]. Such silicon-based microsensors are able to utilise piezoresistive and piezoelectric effects in relation to applied stress, capacitance changes from pressure and thermal resistance which varies as a function of temperature, all of which may be used to measure pressure and pressure-derived parameters such as force, acceleration and fluid flow. They change dimensions on the application of an electric field, eg. PZT ceramic (lead zirconium titanate), PVDF polymer (polyvinylidene flouride) and lead niobate, a property determined by the piezoelectric coefficient $\left(\Delta l/l\right)\left(\Delta E/E\right)$, eg. quartz has a piezoelectric coefficient of 2x10^{-12} m/V. PVDF has a Young's modulus of 1-10 GPa with high mechanical energy density. Lead titanate and lead zirconate have much higher values. Generally piezoelectric ceramics such as PZT are more efficient than piezoelectric polymers such as PVDF (which is also limited to low temperature and pressure environments). A voltage applied across the material generates longitudinal and transverse strains such that the material can act as an actuator. The crystal deformation produces an EMF which can be measured in conjunction with a high gain DC amplifier with high input impedance. Piezoelectric polymers are subject to electric charge generation in proportion to the displacement resulting from an applied force such that integration of the current value gives the force value. Since the Curie temperature must be higher than the processing temperature of the composite to retain its piezoelectric properties ceramics such as PZT with T_{curie}=360 K offer a wider operating temperature range than polymeric materials for use as embedded fibres in a composite. They can give precise motion with repeatable rapid oscillations ~10^3 Hz when wed with ac electricity, but they generate very low actuation levels. Piezoelectric materials can also act as sensors since an applied load will generate a measurable current. Stacked piezoelectric actuators of PZT can suppress high frequency vibrations of 50-300 Hz. PZT may be embedded within a structure as a component of a composite material such as graphite epoxy. This allows the distribution and transmission of significant loads over a wide area. PZT generally exhibits degraded performance at low temperatures making its deployment by spacecraft unlikely. Piezoelectric materials also provide the basis for many MEMS-based sensors including the highly versatile SAW (surface acoustic wave) sensors. Polymeric piezoelectric film transduction has great potential with high signal-to-noise ratios and good flexibility but are again fragile. An example of this technology is the use of polyvinylidene flouride (PVF_2). When stretched to 4 times its natural length it achieves a piezoelectric state. Piezoresistive transduction which also undergoes decreased electrical resistance under pressure also has great potential but the sensors tend to be rigid, stiff and slippery. However they are robust to high temperatures ~1100^{o}F and offer low hysteresis. Magnetoresistive materials (eg. Ni-based metals like Rumalloy 81/19 Ni-Fe) respond to external forces by altered resistance due to altered magnetic fields. They are used in STM (scanning tunnelling microscope) tips and in ultrasonic transducers, and extensively in MEMS devices. Electrostrictive and magnetostrictive materials appear to give enhanced performance for active damping over piezoelectric materials. Electrostrictive materials include lead magnesium niobate (PMN). The hydrocarbon molecules in electrostrictive polymers are arranged in semicrystalline arrays which give them their piezoelectric properties. Magnetostrictive materials exhibit shape changes when subjected to magnetic fields generating large forces over small distances to create moving mechanisms activated by an external magnetic field from a coil wound core, eg. rare-earth elements combined with flourine such as $TbCl_2$ (terfenol). Many are temperature limited and are brittle. They offer fast operation at low voltages but require large applied magnetic fields.

Examples of piezo-based actuators include miniature linear inchworm motors and low power sonar transducers. Piezoceramics have also been demonstrated on micro-robots with piezoceramic legs – a femur provides vertical motion while the tibia provides horizontal motion. The piezoelectric stack, in which displacement amplification is achieved, has been exploited as a mechanical impact drill which generates an ultrasonic pulse that excites a proof mass into high frequency vibration [Bar-Cohen et al 1999]. This is a variation on the travelling wave ultrasonic motor in which the elliptic motion of the stator induced by travelling wave propagation is converted into motion of the rotor by friction coupling. The typical ultrasonic motor exploits mechanical oscillations generated by a metallic resonator induced by piezoelectric component. Periodic frictional contact between the resonator and the rotor generates the driving action. They have efficiencies of 35% and can generate torques of up to 25 Ncm at 150 rpm. The ultrasonic motor has several advantages over the electrostatic motor – they do not require gear reduction as high frequency vibrations of the resonator are transformed to slower macroscopic motion, frictional coupling between rotor and stator exploits the ever present friction problem, and they have higher energy storage density. For precise positioning applications, the impact drive mechanism is superior to ultrasonic motors – the motor comprises two parts, one large mass with high inertial force, and another smaller mass connected to the piezoelectric element. The piezoelectric impact drill use low power, compact ultrasonic vibration ~20 kHz impact devices based on stacked piezoelectric motors with few moving parts - they require ~25-35% the power and generate 20-30 times push force of rotary drills. They operate like jackhammers so do not rotate, "drill-walk" or exhibit low frequency chatter and the drill bit does not require sharpening [Bar-Cohen et al 1999, Das et al 1999]. A typical rotary drill requires 20-30W to produce a 10mm core in hard rock. Generally, piezoelectric actuators have small operating strains making them more useful as sensors rather than actuators.

3.3.2.4. Electrorheological (ER) Fluids

An electrorheological fluid comprises a dielectric carrier fluid in which semiconducting particles are dispersed. Electrorheological and magnetorheological fluids alter their viscosity in response to applied electric or magnetic fields due to polarisation of the particles causing them to form chains with a given orientation to the applied field (Winslow effect). The particle dimensions generally range from 1 –100 μm. The ER fluid retains the viscosity of the carrier fluid until the electric field strength exceeds 1 kV/mm, becoming solid at around 2 kV/mm. Increasing the particle concentration increases the solidity of the material (with a power requirement increase). Electrorheological (ER) fluids may be used as actuators due to their reversible phase change in material characteristics especially bulk viscosity and flow rate when subjected to electrostatic potentials. Many ER fluids are based on micron-sized hydrophilic particles (such as starch) suspended in a hydrophobic non-conducting liquid (such as silicone-oil) in random orientations, eg. particles of crystalline Al silicate zeolites suspended in a non-conducting dry fluid such as silicone oil which can function up to 250^{o}C. On the application of an electric field ~2-4 kV/mm across the ER layer, its molecular structure changes to a solid as the particles orient themselves into a regular chain of columns in milliseconds imparting solidity. The energy dissipation characteristics are thereby altered and the natural frequencies and resonance responses can be changed to damp out vibrations in composite structures by controlling the voltage imposed on the material. ER fluid voids in advanced composites offer this capability. Other possibilities include gear transmission without moving parts, shock absorbers and hydraulic actuators. The chief difficulties are fluid/particle separation due to sedimentation (not a problem in space), electrophoresis and evaporation. Particle abrasion can cause wear and ER fluids tend to have high sensitivity to temperature.

Electrohydrodynamic motion occurs when particles of a polar fluid are subjected to an electric field to generate fluid flow pressure. EHD provides direct conversion of electricity into hydraulic fluid flow without moving parts. They require high operating voltages with low currents and can generate high volume flows. A microscale solid state ethanol pump constructed from charged grids of etched Si ~3mmx3mmx30μ in size has been demonstrated to generate pressures up to 2480 Pa at 700 V and a flow rate of up to 14 ml/min.

Polymeric gels comprise a tangled network of cross-linked polymer chains immersed in a liquid – when stimulated by acids and alkalis, they can swell or contract as "pH muscles" [Osada & Gong 1999, Rossi et al 1988]. The macromolecules form non-uniformly cross-linked 3D networks in ionic fluids such as water to form a gel. They are effectively chemomechanical systems which undergo shape change and generate contractile forces in response to chemical stimuli (typically involving a 1000-fold volume change depending on its cross-link density). They transform chemical energy into mechanical energy and are closely analogous to biological muscles. The most successful of such gels is polyvinyl alcohol (80%) - polyacrylic acid (20%) which contracts under the action of acetone and dilates under the action of NaCl solution [Caldwell et al 2000]. They have limited performance as they are dependent on ion diffusion time and are not well suited to space application due to their requirement for aqueous solution and high temperature sensitivity.

Table 6 – A Summary of PZE, ER, and SMA Materials

	PVDF	**ERF**	**SMA**	**PZT**
Young's modulus	2	0	42	64
Density	1780	1145	2340	7700

3.3.2.5. Electroactive Polymers (EAP)

Electro-active polymers (EAP) are electromechanical systems and are used to convert mechanical strain to electric current thereby acting as sensors. EAPs have also been used as high efficiency, flexible LEDs by virtue of their ready fabrication into films for conformal displays [Wax & Sands 1999]. Electroactive polymers (EAP) which are lightweight strips of flexible plastic that bend or stretch in the presence of an applied voltage [Ashley 2004] - these offer the potential basis for the reciprocating contractile linear actuator properties of artificial muscles. The simplest contractile muscles in the biological world are spasmonemes which occur in some protozoa. They contract on contact with solutions of Ca salts but they offer lower stresses than muscle. Animal muscle is based on the contractile protein, myosin, which slides along another contractile protein, actin, to provide mechanical movement – contraction is caused by the sliding motion of the head section of myosin along the actin molecule. Myosin acts as a molecular motor whereby ATP-generated energy provides for the power stroke followed by a recovery stroke.

Electro-active polymers and ceramics are characterised by an electromechanical deformation response when a voltage is applied. Electro-active polymers offer greater deformation response than electro-active ceramics (which include piezoelectric ceramics) and can provide high strains [Bar-Cohen et al 1998, 1999]. They provide fast response and good controllability. Electroactive polymers have been suggested for use as muscle-like actuators which contract under electrical excitation [Bar-Cohen et al 1998]. There are two types of electroactive polymer – ionic polymers and electronic polymers. Ionic EAPs include dielectric EAP, electro-viscoelastic elastomers and liquid crystal elastomers. Electronic EAPs include ionic polymer gels, electrorheological fluids, conductive polymers and hybrid inorganic/metallic polymers. Elastomers are one type of electronic EAP. Dielectric elastomers such as polyurethane, soft silicones and acrylics which, when exposed to a high voltage electric field, contract in the direction of the electric field and expand perpendicular to them due to Maxwell stress. Soft silicones, which can survive temperatures down to –100^{o}C, offer strains of up to 10-20%. Opposing charge accumulates on each electrode which attract each other. This squeezes the insulating polymer causing it to contract in that direction and expand in the perpendicular direction. At high strains, electrodes must similarly be expandable – carbon particle impregnating an elastomeric matrix offer such compliance. Pre-stretching the elastomers increases the strain capacity of these materials by up to 100 times – up to 350% linear strain, significantly larger than piezoelectric materials. Thin ~30-60 μ films of dielectric elastomers sandwiched by carbon particle impregnated compliant polymer layer electrodes allow expansion of the electrodes with the sandwiched elastomer. Such elastomeric materials act as parallel plate capacitors, which can act as actuators, sensors or energy transformation devices. The two metrics for the performance of artificial muscles include: strain (displacement/unit length) and force (pressure/unit density). Performance can be enhanced by the construction of spring rolls in which elastomeric sheets with expandable electrodes are wrapped cylindrically around a compressed helical spring. The spring pre-stresses the sheets circumferentially. Such devices can be used as linear or angular robotic actuators, valves and pumps with up to 30 N force capability, 2 cm displacement and cyclicity of 50 Hz (angular movement is generate by using dual parallel configuration of such spring rolls, similar to a bimetallic strip). Force amplification may be achieved through series and parallel configurations of such devices. They require comparatively high voltages ~1-5 kV but low current. Artificial muscles are well suited to micro-fluidic applications as fluid pumps, one configuration being a film diaphragm powered by low frequency voltage signals to act as valves to open and close pressurised fluid reservoirs. Furthermore, as capacitative devices, elastomeric polymers can act as sensors. They are particularly suited to application in tactile sensing in "artificial skins". As capacitors, elastomeric polymers can also be used as power storage devices which store applied mechanical energy such as from wind disturbances as electrical energy. Alternatively, polymer engines may be possible in which propane fuel is burned in a combustion chamber, which distorts an elastomeric diaphragm, generating electricity. Electronic electro-active polymers deform by virtue of electronic carrier motion but this requires high voltages ~100 V/μm close to breakdown level. Electronic EAPs include ferroelectric polymers, electrets, dielectric elastomers and electrostrictive elastomers. Electronic EAPs are driven by electric fields but require high voltages, particularly electrostrictive polymers (which require ~10^2 V). However, they have high rapidity of response, can generate high mechanical forces and require negligible current to maintain their configuration.

Ionic EAPs include ionic polymer gels, ionomeric polymer-metal composites, conductive polymers and carbon nanotubes. They are current driven rather than voltage driven allowing the use of low voltages. Ionic polymers, conducting polymers and polymeric gels perform well at low frequencies but cannot operate at high frequencies due to the inherent slowness of ionic diffusion. Fluorinated ion-exchange membrane platinum electroactive polymers, such as the perfluorosulphonic acid polymer Nafion plated with Pt electrodes, deform by virtue of ionic carrier motion and require only low voltages ~1-3V and currents ~mA. However, ionic polymers generally require moisture to operate as they are hydrophilic but may be sealed within flexible coatings. Such ion-exchange-polymer-metal composites (IEPMC) offer a potential use as biomimetic sensors/actuators which can undergo large bending displacement but offer low force actuation capability when an electric field is imposed across its thickness (greater than polymeric gels) and vice versa - this material can oscillate at low frequency ~1-0.01 Hz under an ac square wave input [Shahinpoor et al 1998]. They have been used as fuel cells and water electrolysis and their function as actuators is the reverse process to charge storage. IEPMC materials can be readily formed and cut into any desired shape and configuration. These materials possess ionisable charge groups (polyions) on their molecular backbone which are susceptible to externally applied fields. They are three-dimensional networks of cross-linked 3D macromolecular polyelectrolyte networks with internal electrodes that swell, contract, bend and deform in an electric field forming a gel in water. Strips of this material may be composited with a noble metal such as platinum or gold. If a displacement is imposed, a proportional output voltage is generated. They have good lifetimes of $\sim 10^6$ cycles. Conversely, at a critical applied resonant frequency, maximum deformation is generated with high repeatability and accuracy. Nafion film has the following chemical formula [Bar-Cohen et al 1999]:

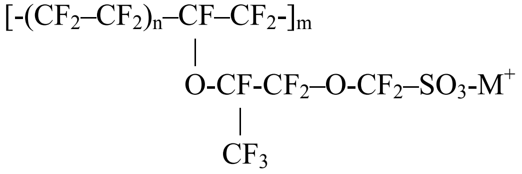

$$[-(CF_2–CF_2)_n–\underset{\substack{|\\ O\text{-}\underset{\substack{|\\ CF_3}}{CF}\text{-}CF_2–O\text{-}CF_2–SO_3\text{-}M^+}}{CF}–CF_2\text{-}]_m$$

where:

$n \sim 6.5$
$100 < m < 1000$
$M^+ = (H^+, Li^+, Na^+)$

These polymers comprise a hydrophobic (fluorocarbon backbone) and hydrophilic ($-SO_3^-H^+$) regions. The proton may be exchanged for other mobile cations such as Li^+ or Na^+. Nafion can absorb large quantities of polar solvents such as water in its hydrophilic channels. EAPs can be rolled into ropes to act like muscle fibres.

Such materials can operate at low temperatures ~-140^{o}C and in a vacuum. They are extremely sensitive to contamination and require clean-room facilities for processing. If electric current is maintained above a critical voltage, charge migration continues until electrolysis occurs which causes irreversible damage to the material. Their chief disadvantage is that they require water to operate but may be sealed. Spraying with polysilicon polymer coatings for environmental resistance (chemical corrosion, moisture, ultraviolet radiation, high temperatures up to 250^{o}C) provides a skin to prevent drying for months but their low adhesion necessitates the prior use of chemical etchant (eg. Tetra-etch) to remove F atoms from the surface chains and subsequently exposed to moist air to form hydroxylm carbonyl and carboxyl groups for adhesion [Bar-Cohen et al 1999]. Flourinated ion-exchange membranes with platinum electrodes chemically deposited on both sides of the membrane are generally adopted. These 0.2mm film polymers bend when exposed to an applied electric field. Back-to-back pairs will contract longitudinally while thickening in diameter. An applied voltage of 5V will contract the electroactive polymer by 10%. A 0.1g film can carry 6g with low power consumption ~20-30 mW. Although their force delivery is limited, they have mass, power and response speed advantages. They offer lower voltage requirements and faster response times than shape memory alloys and have been used in a multi-fingered gripper prototype for a Mars retrieval end-effector with a force density of 400 N/kg [Lumia & Shahinpoor 1999]. Ionic conducting polymer film actuators have been used to generate travelling waves for ultrasonic motors [Wax & Sands 1999] Poly(vinylidene difluouride-co-trifluoroethylene) electrostrictive co-polymers which are based on phase change from ferroelectric to paraelectric states have improved strain (around 4% strain capability) and energy densities (of 3 x 10^5 J/m^3) over more conventional EAPs and piezoceramics. Electroactive polymers may be used as sensors – polymers of the poly(pthalidylidenarylene)s class, eg. poly(3,3'-pthalidylidene-4,4'-biphenylylene), are sensitive to temperature, pressure, electric field, pH, etc which generate an insulator-metal transition in the polymer film whose resistance changes by 10-12 orders of magnitude [Lachinov et al 1996].

Conducting polymers convert electrical energy directly into mechanical energy with high force generation capabilities and high power output densities yet require only low voltage input with much greater superiority over polymer gels [Baughman 1996]. However, they have limited lifecycles ~10^6 and limited energy conversion efficiencies. Common conducting polymers include electrochemically-doped polypyrroles and polyanalines. The dopants are usually alkali metals which modify the polymer's electrical and mechanical properties. They require three components – thin strip anode and cathode electrodes and intervening solid state electrolyte. The oxidation of an ideal α-α'-polypyrrole (pPy) conductive polymer chain film in an electrolytic medium containing ClO_4^- due to an applied voltage between the electrodes is given by [Otero et al 1999]:

$$(pPy)_s + (nClO_4^-)_{aq} + mH_2O \Leftrightarrow [pPy^{n+}(ClO_4^-)_n(H_2O)_m]_{gel} + (ne^-)_{metal}$$

This oxidation reaction causes a significant increase in volume due to inclusion of counterions and water and associated polymeric conformation changes. Thus, electric current makes this conducting polymer act as a molecular motor. Their key differentiating characteristic from electrostatic, piezoelectric and electrostrictive actuators is that the electrically-generated strain is dependent on the applied voltage rather than the applied field. Conducting polymers with a low applied voltage of ~V across them will generate a strain and induce linear dimensional changes of 10% (compared to piezoelectric crystals which offer ~0.1% dimension changes). They have the advantage that they can be maintained at any actuator response position without energy input, though they can suffer drift. Their response times are dependent on molecular diffusion with RC time constant but this can be minimised by using highly conductive thin film electrodes (eg. polyanaline) and electrolytes (eg. hydrated poly(vinyl alcohol)/H_3PO_4). Their response times are ~0.3 ms, much faster than biological muscles ~1 ms, and can exert a maximum force of 80-100 times that of muscles. Conducting polymer electrochemical cycle corresponds to the charge-discharge cycle of a battery. As they are effectively batteries, they can be used as such when not connected to an external power source and to assume a desired mechanical configuration on discharging (similar to shape memory alloy). They can also be used as sensors as thin film cantilever structures with sensitivity to gases. It may also be possible to construct non-faradic electromechanical actuators from carbon nanotubes acting as conducting polymers suspended in an electrolyte which may offer 1% strains with ~ns time constants.

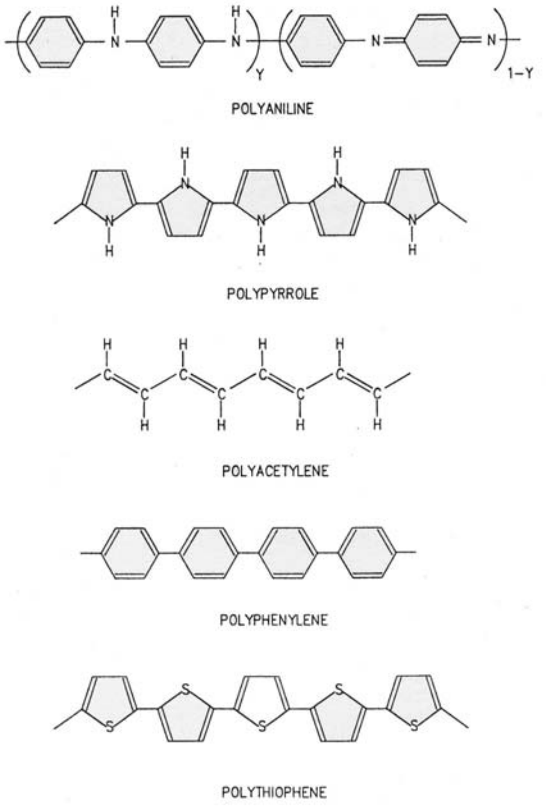

Figure 20 – Example conducting polymers that may also be used as biosensors (Gerard et al 2002)

Table 7 – Review of EAP Actuators (Wax & Sands 1999)

Actuator	Performance	Advantages	Disadvantages
Polymeric gels	Stress: 1-0.3 MPa Strain: 1000%	Large volume changes	Slow response
Ionic polymers	Stress: 0.1-1 MPa Strain: 1-10%	Large deformations	Poor understanding
Conducting polymers	Stress: 450 MPa Strain: 1-10%	Reasonable deformations and forces	Difficult processing
Electrostrictive polymers	Stress: 0.2-2 MPa Strain: 10-30%	Small strains but large forces	Hysteresis

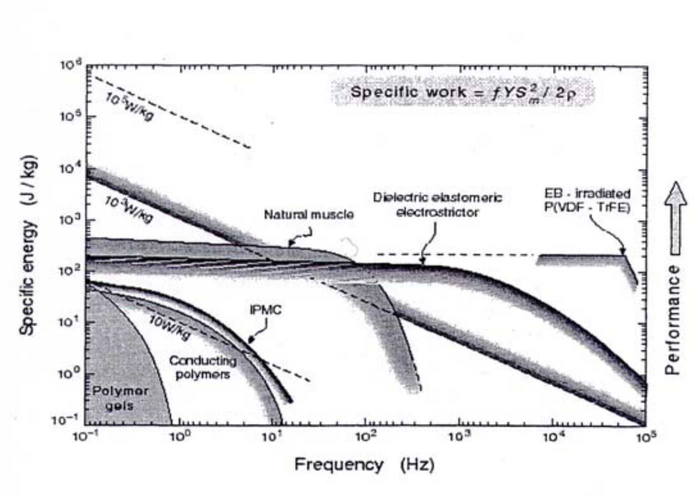

Figure 21 – EAP actuator energy and bandwidth (Wax & Sands 1999)

Many recently developed actuator technologies are termed "artificial muscles" by the ones who create them. These technologies include (but are certainly not limited to) shape memory alloys, flexible pneumatic chambers, electrostatic devices, and a large variety of mechanochemical or electroactive polymers. However, actuators based on these technologies are not like muscles in ALL respects. In the context of biomimetic robotics, Roy Kornbluh and his team define an artificial muscle as "an actuator that embodies certain desirable characteristics of natural muscle".

Biologists have identified some of the characteristics of natural muscle, but many have not yet been identified. It is therefore a complex task to select suitable metrics to determine whether a given actuator technology is "an artificial muscle", as some less obvious muscle characteristics may help to emulate the desirable motion of biological creatures. These include degradation of performance with fatigue, ability to self repair, sensitivity to temperature, the ability to be consumed to provide energy, composition of a large number of fibres acting in parallel, and the diverse performance of individual muscle fibres. Although future biomimetic robots might incorporate many of these features (assuming they are considered important), the role of the next generation of biomimetic robotics will be based almost solely on its ability to produce and control motion. It is to these features focusing solely on motion that this team is focusing on.

How muscles work

Conventionally, muscle is viewed as a linear motor/actuator that provides a motive force to an object; essentially 2-dimensional. However, they also function as brakes and shock absorbers, springs, and even relatively rigid struts. The functional role of an individual muscle may change during the course of a certain task, depending on joint kinematics and external loads. Muscles can also perform more than 1 specific role at a time, such as working as a spring and a motor or brake. A true artificial muscle actuator should emulate all the functions that are important to the execution of motion. Leg muscles, as in cockroaches, also contribute tot damping disturbances, a response now termed a *preflex*. A preflex, in contrast to a reflex, does not use any sensory feedback or closed-loop control and instead acts almost instantaneously on external perturbations based on open-loop dynamic response of the muscle and joints. The spring-like behaviour of muscles is important in making locomotion more efficient, as energy normally used to accelerate or decelerate the body can be recovered during cyclic locomotion and used elsewhere in the cycle. Examples of such include wing flapping in many insects and small birds, kangaroo hopping, swimming fish, and walking or running of legged creatures including humans.

Table 8 – Common Actuator Metrics and Parameters

PARAMETER	SCALE-INVARIANT VERSION	COMMENTS
Energy and Power		
Energy output	Specific energy	Over full cycle
Power output	Specific power	Average or peak
Energy conversion efficiency	Same	Energy output over a full cycle / energy input over a full cycle (excluding recovered energy)
Response time, frequency bandwidth	Same (usually scale dependent)	For one directions or full cycle
State Variables		
Displacement, stroke	Strain	Peak over a cycle
Force	Stress, pressure	Peak over a cycle
Velocity	Strain rate	Peak over a cycle
Impedance and Controllability		
Stiffness	Elastic modulus	Usually non-linear (not a constant)
Damping	Specific damping, loss factor, loss tangent	Usually non-linear (not a constant)
Accuracy (displacement or force) (%)	Percentage of strain or stress	Usually percentage of maximums
Repeatability (%)	Same	Usually percentage of maximums
Linearity or sensitivity (%)	Same	Deviation from linear input-output relationship
Operational Characteristics		
Environmental tolerance	Same	Recommended ranges of temperature, humidity, etc., or effects of variations in temperature, humidity, etc., on the above parameters
Durability, reliability	Same	Number of cycles before degradation threshold or total failure, degradation per cycle or time
Input impedance (power supply requirements)	Specific impedance	Voltage and current requirements or pressure and flow, depending on power mode

Unfortunately, these parameters alone cannot serve to accurately compare actuators to natural muscle nor determine whether or not a given actuator technology can fairly be called an artificial muscle. First, the parameters of different types of actuators and muscle may not have been determined under identical conditions. This is significant because muscle is highly dependent on the magnitude and timing of specific inputs, motions, and loading conditions. Second, the parameters in Table 8may not accurately or completely describe the behaviour of a muscle/actuator in actual usage. Often force relates to the maximum blocked force (isometric – muscular contraction against resistance) and stroke is the maximum “free” stroke at a constant force loading (isotonic – muscular contraction in the absence of significant resistance). However, these definitions do not provide enough information to determine the actual force or stroke output produced but they do directly indicate the amount of work that the actuators can perform. This could fit an actuator between the two extremes of 1) being completely constrained and have a large force but zero stroke and do no external work and 2) a free-moving actuator that may have a large stroke but exert no force and do no external work. A better metric would be to use parameters such as maximum blocked force and maximum stroke at a constant force loading, which would give a better indication of the capabilities of the actuator.

In the study of biological motion, we cannot consider muscle in isolation. Individual muscles work with other muscles, limbs, segments, bones, etc. to accomplish the required task. In the case of biomimetic robots, the artificial muscle is often attached to a skeleton-like linkage, however such linkages introduce significant non-linearities in the forces transmitted to and from the muscle.

Dielectric Elastomers

Dielectric elastomer actuators have demonstrated characteristics such as high strain, high efficiency, high energy density, fast response, good controllability, and the capabilities of impedance selection, and open-loop modulation. Because of these characteristics, dielectric elastomers are a good candidate for consideration as artificial muscles for biomimetic robots. (Kornbluh et al. 2002)

Organic-Inorganic Hybrids

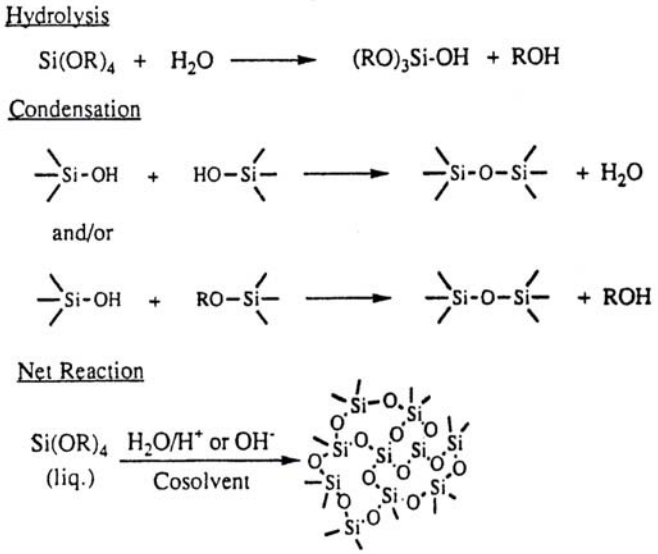

Figure 22 – Hydrolysis and condensation of silicon alkoxide forming silica network in sol-gel process (Ahmad & Mark 1998)

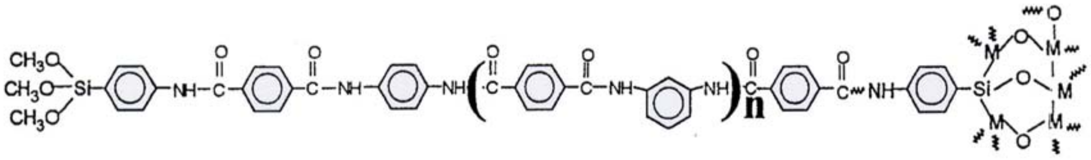

Figure 23 – Structure of aramid chain chemically linked with inorganic network with aminophenyltrimethoxysilane (M=Si, Ti, Zr or Al) (Ahmad & Mark 1998)

3.3.2.6. Conclusions

Animal muscles are capable of exerting high lever torques with great precision. Muscles also provide additional functions as shock absorbers and springs by imparting damping [Kornbluh et al 2002]. Muscles can operate at a wide range of frequencies and speeds for a wide range of functions [Dickinson et al 2000]. They are characterised by highly non-linear multi-parameter dynamic systems. In particular, muscle fibre contraction provides the basis for rotary joint motion – Hill's model comprises a contractile and series elastic components such that the cross-section of the muscle defines its force potential. Current muscle-like actuators include shape memory alloys, dielectric elastomers and electroactive materials [Fijita 2001], though polyMEMS actuators have been proposed to exploit cellular construction [Horning & Johnson 2002]. Shape memory alloys exhibit limited deformation response but electro-active polymers are lightweight plastic strips of high bending capacity. Rather than discrete variability in actuation mechanisms, continuous variability in actuation response offers greater potential and may be accomplished using embedded multiple actuators in conjunction with fibre optic sensors [Ghandi et al 1991]. Electrorheological fluids and piezoelectric materials exhibit more continuous behaviour than shape memory alloys which tend to exhibit discrete behaviour. Generally, it is difficult to obtain simultaneous large deflection with high actuation force as these properties appear to be inversely related.

Table 9 – Comparison of IPMC, SMA and EAC (Lumia & Shahinpoor 1999)

Property	**Ionic polymer metal composites**	**Shape memory alloy**	**Electroactive ceramics**
Actuation displacement	>10%	<8%	0.1-0.3%
Force (MPa)	10-30	~700	30-40
Reaction speed	μs-s	s-min	μs-s
Density (g/cm^3)	1-2.5	5-6	6-8
Drive voltage (V)	4-7	N/A	50-800
Power consumption	mW	W	W
Fracture toughness	Resilient, elastic	elastic	brittle

3.4. BIOMIMETIC STRUCTURES, MECHANISMS & DEPLOYABLES

Biological material is typically "soft" with structural stiffness commonly imparted through mechanical rigidisation of fluid pressure (turgor). Biological materials are generally highly compliant with low Young's modulus and provide toughness for resistance to fracture rather than stiffness for resistance to deformation. Mechanical loads in biological organisms are carried by carbon-based polymer fibres – cellulose in plants and collagen in animals. The macromolecular chains are sometimes twisted into fibrillar structures with high modulus stabilised by side groups (resembling nylon, aramid and polyethylene artificial fibres). Like engineered composites, their physical properties are highly anisotropic offering stiffness and strength matched to the direction of the applied load. The fibres are bonded within a matrix (similar to engineered fibre composites) of polysaccharides, polyphenols or buckling resistant minerals such as calcium carbonate or hydroxyapatite. Fibres can carry only tensile loads as they are prone to buckling when exposed to compressive loads – when used in composites the matrix does not contribute to load-bearing. Biological systems often overcome this deficiency by pre-stressing them to impart tension. Fluid pressure is used in non-woody plant cells and soft animals to pre-stress fibres to provide the structure with resistance to compression. Cytoskeletons comprise chains of proteins which give cells their shape, movement and structural integrity. They are arrays of protein filaments which form networks that give cells their shape and provide for movement. They originate from near the nucleus which contains the cell's pair of centrioles. The main kinds of cytoskeletal filaments are microtubules ~25 mm in diameter and actin filaments ~7 nm in diameter. The cytoskeletonprovides tensegrity (tensional integrity) in which protein chains provide both compression (as in rods) and tension (as in wires) for structural integrity within a stable but flexible structure [Ingber 1998]. Tensegrity refers to mechanical stabilisation through the distribution and balance of tensional and compressive forces and stresses within a structure. Prestressing involves structural members being in tension or compression before the application of mechanical loads – compression-bearing struts resist opposing tension-bearing cables in a mutually stabilised network. Tension is long-range, typically transmitted across the structure with short-range compression being balanced locally at intervals through the structure. The contractile microfilaments exert tension throughout the biological cell. Opposing this are the compressive microtubules, cross-linked microfilaments and intermediate filaments which act as guy wires. The internal cytoskeleton of the biological cell can also perceive applied forces through integrin proteins embedded in cell membranes which are "spot welded" to the extracellular matrix where they possibly detect stiffness of the cytoskeleton.

Plant leaves can react to environmental conditions such as tracking the sun (phototaxis). Leaf structures may be packaged and then deployed which have applications in the deployment of antennas and solar panels (eg. solar power satellites). Osmotic pressure in plant cell walls provides the basis for such actuation mechanisms. Although the osmotic pressure is low, it is amplified by the storage of elastic energy in the cellulose structure. Insects adopt a similar strategy to store energy in resilin within the cuticle to drive their wingbeats during flight. Spiders transmit hydraulic pressure in their blood by compressing their bodies to extend their legs. Adaptive (smart) structures are essentially structures that can alter their geometric characteristics according to either the external environment or to imposed commands. Smart structures generally involve the integration of sensing (strain, temperature, pressure, etc), processing this information, and reacting to that information. The traditional approach to structural design has been to incorporate sufficient robustness to ensure mechanical integrity. Biological materials often incorporate this function – the horns of many mammals are sheathed in fibrous keratin which is insensitive to notch weakness when damp. Smart materials however are multi-functional – multi-functional structures for spacecraft are a step in this direction. Biological materials are often characterised by reactivity to environmental conditions – "smart" materials. Smart materials incorporating sensors such as thermal or pressure sensors offer the potential for distributed sensing arrays. Slit sense organs in arachnids are located in holes in the cuticle of the body. Their shapes vary from circular to elliptical slits ~5-200μm to which a thin membrane ~0.25μm thick cover impregnated with sensory cells. Loads on the exoskeleton are sensed by these strain gauges by deformation of the holes. The elongation of the slit allows preferential directionality perpendicular to its long axis. Such slits are aligned perpendicular to the lines of principal stresses. In addition to sensing the environment, smart materials may move and adapt to their environment such as altering their stiffness within a vibration environment. Such capabilities would be of immense benefit for vibration control of large lightweight structures in space. Such smart materials with actuation capabilities offer the possibility of reacting to the environment and indeed, acting appropriately to environmental conditions. Adaptive or “smart” structures are essentially structures that can alter their geometric characteristics according to either the external environment or to imposed commands. Smart structures generally involve the integration of sensing (strain, temperature, pressure, etc), processing this information, and reacting to that information. Desirable characteristics include adaptive materials (shape memory alloys) for vibration suppression, self-healing materials (ionomers), tactile surfaces (piezolelectric/magnetostrictive/electrostrictive materials), multifunctional structures with embedded molecular wires and optic fibres, structures which can alter their viscosity (electrorheological/magnetorheological fluids), etc. For large deployable structures, gossamer ultra-thin film materials would drastically reduce the weight and volume of antennae and solar panels.

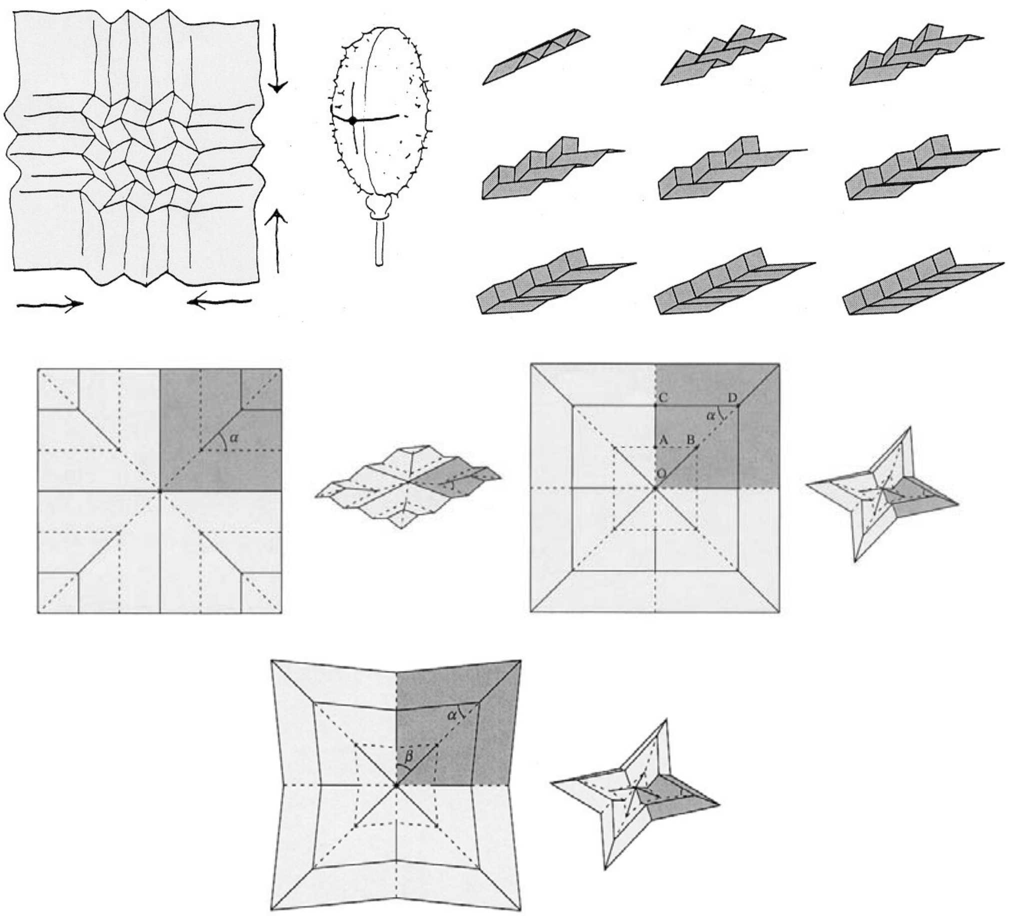

Figure 24 – a) Poppy petal with pleated folding structure similar to Miura-ori (from Vincent 2000); b) nine stages of deployment of leaf-folding pattern which extends simultaneously in two directions; c) leaf-out pattern where leaves point away from centre; d) leaf-in pattern where leabes point towards centre; e. skew leaf-in pattern (De Focatiis & Guest 2002)

3.5. BIOMIMETIC METHODS OF PROPULSION

Movement is a characteristic of the animal kingdom defined by the process of energy intake by feeding. Animals search for food while plants extract food from the environment directly. Movement is a necessary part of exploration and animal locomotion provides the basis for animal survival - the avoidance of threats and the acquisition of resources. The ability to move is the key to intelligence – animals which are mobile (by virtue of their heterotrophic energy assimilation) possess nervous systems, while plants which are static (by virtue of their autotrophic energy assimilation), do not. Molluscs are commonly sessile or slow-moving and have rudimentary nervous systems. Cephalopods however are highly mobile and are characterised by sophisticated nervous systems capable of supporting sophisticated, problem-solving behaviours. Indeed, movement is a fundamental signature of life at microscopic scales. Movement is essentially the actuation component to a control system. In planetary exploration, biomimetic locomotion offers the potential for robust exploration in hostile environments with high scientific returns per platform per mission. We suggest that the key to biomimetic locomotion of all types is the approach to control – robust locomotion requires the implementation of central pattern generation, sensory feedback and preflexes enabled through structural compliance. The problem of structural stiffness/compliance is fundamental in locomotion and manipulation. Robotic control is mainly concerned with precise positional placement or the application of precise forces. Legged locomotion in animals relies heavily on the compliance of the body provided by the spine in the case of vertebrates. In invertebrates such as arthropods this is provided by the joints between the three body segments as well as compliance in the jointed legs. We suggest that “smart” structures may contribute significantly to the implementation of force feedback and preflexive compliance – muscles provide actuation, damping and feedback. We further suggest that a general purpose biomimetically based robotic explorer may be developed which exploits the flight mechanisms of a dragonfly, the legged locomotion of a cockroach, the drilling capability of a mosquito or wood wasp, and the manipulation capability of a beetle.

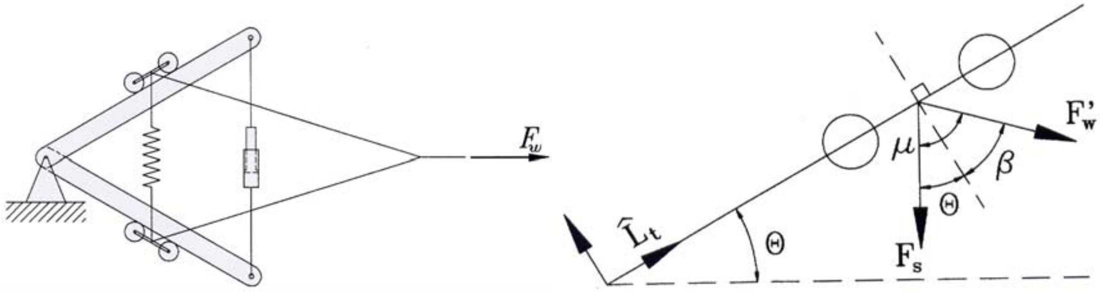

Figure 25 – a) biomimetic actuator stiffness modulation mechanism; b) spring model of stiffness modulation (Kolacinski & Quinn 1998)

Dickinson (2000) states that one of the key challenges in the study of locomotion is to determine how each individual component within a locomotor system operates, while at the same time discovering how they function collectively as an integrated whole. An integrative approach to locomotion focuses on the interactions between the muscular, skeletal, nervous, respiratory, and circulatory systems. These systems possess functional properties that emerge only when they interact with each other and the environment. In his recent investigation into advances in the study of animal locomotion, he has notices four main themes emerge:

1. The spatiotemproal dynamics of locomotion are complicated but understandable on the basis of a few common principles, including common mechanisms of energy exchange and the use of force for propulsion, stability and manoeuvrability.
2. The locomotory performance of animals in natural habitats reflects trade-offs between different ecologically important aspects of behaviour and is affected by the physical properties of the environment.
3. The control of locomotion is not a linear cascade, but a distributed organization requiring both feed-forward motor patterns and neural and mechanical feedback.
4. Muscles perform many different functions in locomotion, a view expanded by the integration of muscle physiology with whole-animal mechanics.

Although we have not considered marine locomotion specifically, Robo-Lobster is a wheeled robot which implements plume tracking in land environments:

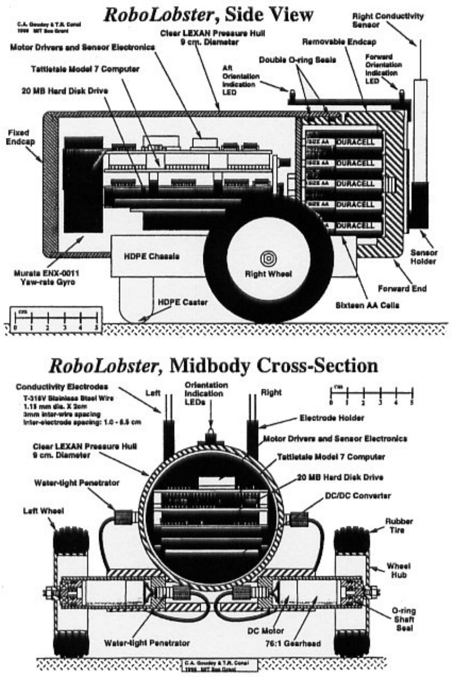

Figure 26 – RoboLobster configuration (Grasso et al 2000)

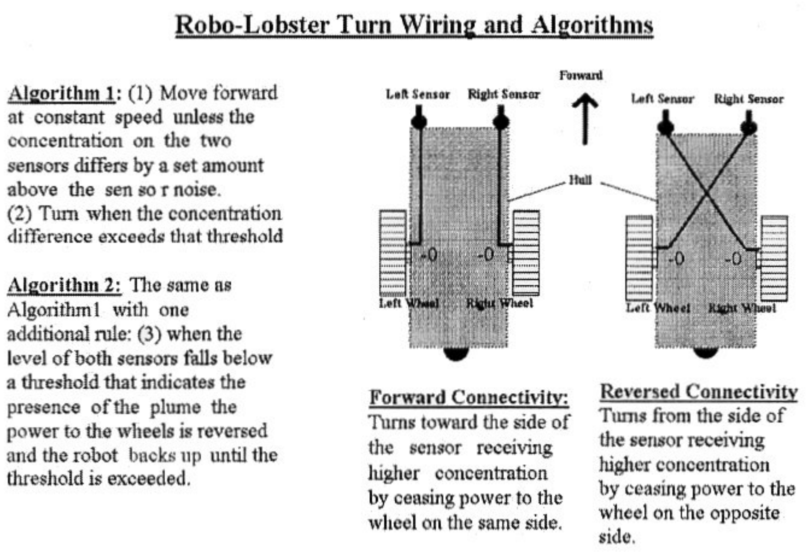

Robo-Lobster Turn Wiring and Algorithms

Algorithm 1: (1) Move forward at constant speed unless the concentration on the two sensors differs by a set amount above the sen so r noise.
(2) Turn when the concentration difference exceeds that threshold

Algorithm 2: The same as Algorithm1 with one additional rule: (3) when the level of both sensors falls below a threshold that indicates the presence of the plume the power to the wheels is reversed and the robot backs up until the threshold is exceeded.

Forward Connectivity: Turns toward the side of the sensor receiving higher concentration by ceasing power to the wheel on the same side.

Reversed Connectivity Turns from the side of the sensor receiving higher concentration by ceasing power to the wheel on the opposite side.

Figure 27 – RoboLobster wiring and algorithms (Grasso et al 2000)

3.5.1. Walking

Engineered systems almost exclusively use wheel-based systems but nature has not invented the wheel (except in bacterial flagella) and relies on linear motion. Biology has not exploited the wheel due to the difficulty in connecting a free-rotating wheel-axle to a fixed support and the requirement of a smooth surface for it to run on. Ambulatory locomotion is based on reciprocating legged motion (walking). Insect walking with six legs is a commonly adopted approach though bipedal walking is also of interest. Walking locomotion has a number of advantages including greater robustness – discrete footholds rather than complete free pathways are required. This is essential for rough terrain typical of planetary environments, such as Mars, which are covered in rocks. Some of the climbing abilities of legs have been invoked with wheeled vehicles through "wheel walking" which allows a vehicle to negotiate slopes at the soil angle of repose. Tracked vehicles are similarly suited to rough terrain but large obstacles and slopes can cause instability and tipping, e.g. Pioneer, which was a tracked adaptation of Sojourner used for sealing Chernobyl in 1999.

Walking overcomes irregular terrain negotiation by dynamically decoupling the body of the vehicle from the ruggedness of the terrain. This is enabled by a large number of degrees of freedom in the legs allowing for simple adaptability to changes in terrain. Terrain contour following is additionally enhanced by body articulation – indeed, much of the locomotion capability of animals derives from the degrees of freedom movement of the spine which provides compliance of the body. During locomotion, mammals essentially behave as inverted spring/pendulum systems in which the structure acts to stabilise the body and keep it mobile at a preferred frequency, as represented in Figure 28. The mammalian leg comprises three jointed segments – thigh, shank and foot connected by the hip, knee and ankle joints, respectively. The action of the muscles results in a force on the ground. In humans, the shank and thigh are approximately the same length. Running is essentially a process of jumping during which the animal has no contact with the ground – the animal's centre of gravity moves along a parabolic curve during the jump. For running, the launch angle α is small to reduce the energy wasted in lifting the centre of mass, and typically has the value ~10-20°.

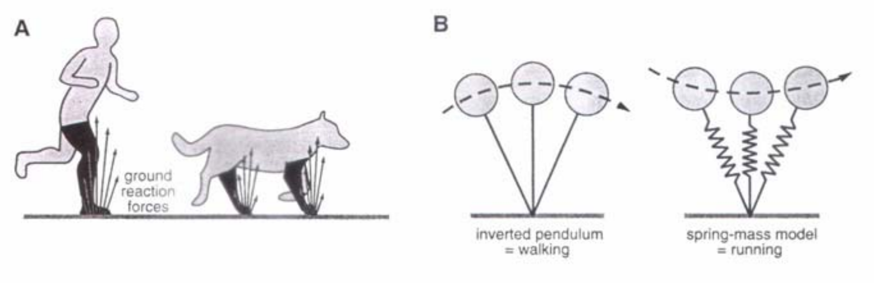

Figure 28 – Ground reaction forces of mammal running and the pendulum representation of walking and running (Dickinson 2000)

3.5.1.1. How Creatures Move

Although it is possible to deconstruct the mechanics of locomotion into a simple cascade – brain activates muscles, muscles move skeleton, skeleton performs work on external world – such a unidirectional framework fails to incorporate essential dynamic properties that emerge from feedback operating between and within levels. One key challenge in the study of locomotion is to determine how each individual component within a locomotor system operates, while at the same time discovering how they function collectively as an integrated whole.

The legs of sprawled-posture animals (insects, crabs, crayfish, spiders, geckos, etc) generate substantial lateral forces while in contact with the ground. The time course of these lateral forces is consistent with the hypothesis that elastic energy storage and recovery may occur with in the horizontal plane, orthogonal to the direction of motion. Species with sprawled postures can actively alter course by changing the orientation of forces generated by a single leg.

Where legs of a runner push against a solid substrate, the appendages and body of swimmers and fliers push against fluids, which distort and swirl to form a complex wake. As these animals move through a fluid (air or water), it creates vorticity, a circular flow of motion in the fluid that represents the momentum imparted by the animal onto that fluid.

3.5.1.2. Walking Devices and Gaits

Legged locomotion control is feed-forward in nature imposed locally at each joint rather closed loop and global. This is particularly the case for rapid and repetitive behaviours (this is the basis for the central pattern generator that generates rhythmic motion in insect gaits).

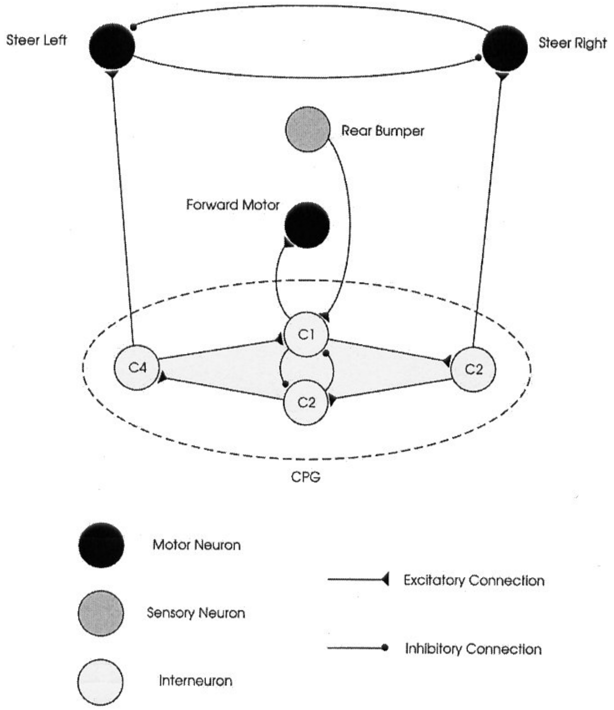

Figure 29 – Central pattern generator of four neurons with mutual inhibition between two neurons (C1 and C3) to generate oscillatory firing pattern. There is excitatory connection from neuron C1 to forward motor and from C2 and C4 to right and left motors (Damper et al 2000)

Behaviours more complex than taxes involve more complex neuronal circuitry which have rhythmic dynamics (central pattern generators). Pacemaker neurons exhibit oscillatory behaviour by virtue of their internal dynamics. Oscillatory circuit networks of neurons have reverberating pathways whereby activity is amplified by positive feedback. Central pattern generators of different behaviours have been built up using neural circuits: exploratory wandering and feeding. Other possibilities for the multiple behavioural suite include flight and fight. There is often insufficient time for feedback reaction to sensory signals. Coordinated control of the legs emerges from interactions between the individual leg controllers. The robustness of this feed-forward strategy is provided by compliance in the structure of the organism's appendages.

Walking devices also require rules that define the timing and synchronisation of placement and lifting of each foot – the periodic gait. Gaits may move one, two or three legs at a time and different gaits are required for different terrain. In general, the number of possible gaits for an *n*-legged robot, $N_g=(2n-1)!$ (where *n* is the number of legs), but few are adopted in biological organisms. The most robust control strategy is a distributed one in which each leg is controlled independently. Many legged robots employ cyclical coordination of two groups of legs. Generally, any reciprocating motions in propulsion systems causes power loss at a rate proportional to the reciprocating mass and square of the peak velocity unless the system oscillates at its natural frequency. For this reason, legs should swing at their natural pendulum frequency. Posture control is concerned with the position of the centre of mass of the body and is central to the static stability of the vehicle [Quinn et al 2002]. Swing control is concerned with the cyclic periodicity of the legs and is the basis of gait. Stance control is associated with the leg configuration. Gait generation is generally distributed being based on local controllers – the velocity of each leg is adjusted as it moves through its swing or power stroke based on phasing signals from neighbouring legs which generate a limit cycle of oscillation.

In walking robotics, mimicking exactly a natural walking system is not practical, as these systems are incredibly complex. Gernot Wendler discusses walking locomotion of insects as consisting of 6 coupled relaxation oscillators, one for each leg. This is based on observations from the 40s, 60s, and 90s from von Holst, Wendler, and Bassler, respectively, examining how each insect leg moves at different frequencies. Recordings of the protractor and retractor nerves in walking animals show rhythmic spike activity, and these nerves produce different spike amplitudes. Each leg muscle is innervated by more than one motoneuron (between 3 and 14 depending upon which muscle). Additionally, individual muscle fibres receive different combinations of motoneurons, making control of muscle contraction extremely complicated. Also, sense organs monitor all leg movements and provide feedback for the coordinated movements required to make a full step.

Cruse (2000) discusses a decentralized, biologically based network for controlling walking robots. In doing so, he first proposes a top-down design, from a more global view to a more local view of the robot. In his consideration, he begins at the movement of the whole body, to coordination between legs, to the control of a single leg, to the movement of each joint within the leg and beyond. He notes that the most common walking patterns in insects are the tripod and tetrapod gaits, shown in Figure 30. These gaits are extreme cases assuming no disturbances in the walking pattern or within the movement of a single leg. The tripod gait involves front and back legs on each side of the body swinging synchronously with the middle leg on the opposite side. This gait is used for faster walking speeds. In the tetrapod gait, one leg on each side moves at the same time, retaining at least four on the ground at any time, which is more stable than the tripod gait and used for slow walking speeds. Sequentially from back to front, the right back and left middle legs move, then right middle and left front, then right front and left back. In either gait, once a disturbance occurs, the insect will slip into a free gait, until it can flow back into its normal pattern.

Insect locomotion is controlled by neurally-based central pattern generators sometimes augmented by feedback mechanisms [Delcomyn 1980, Venkataraman 1997, Klavins et al 2002]. Central pattern generators are groups of neurons which produce rhythmic outputs without sensory input. Central pattern generators are often modelled as sets of bistable oscillators in which reciprocally connected networks have inhibitory connection weights. Individually controlled legs coordinated by an oscillatory pacemaker neuron may determine which legs move forward while each side is phase-locked to each other, thereby implementing a central pattern generator. However, proprioceptive feedback from stretch receptors is essential for maintaining the frequency and duration of rhythms of central pattern generators. Insect locomotion is robust but mechanically complex to control. Although central pattern generators have primarily been used for legged locomotion, they are suited to any type of rhythmic locomotion, e.g. flying, swimming, etc [Cruse et al 1995]. Hippocampal neurons also have the property of firing in bursts suppressed by GABA-mediated inhibitory connections in their normal state. Projects which seek to emulate insect walking include robotic scorpions for rugged environments and robotic lobsters for aquatic environments [Klaasen et al 2001, Ayers et al 1998, 2000]. There are other forms of locomotion including crickets which move by jumping on powerful legs which employ a lever principle – McKibben actuators are suited to this type of biomimicry [Quinn et al 2002]. Locusts also utilise elastic energy stored in the tendon of the hind femur for jumping.

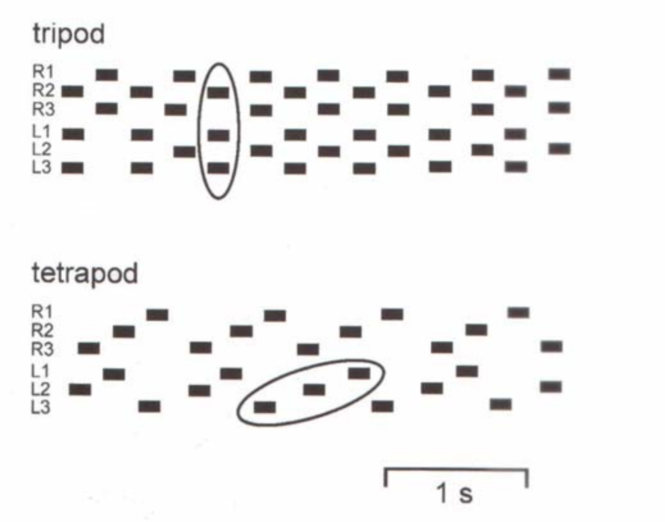

Figure 30 – Tripod and Tetrapod Gaits for a 6-Legged Walking Robot, where the black bars represent the swing movement of each leg during locomotion.

Cruse shows how higher-level coupling rules are only possible after lower level decisions have been made, for example, with a single leg. The transition points between gaits or when a decision to sway from the standard gait are called extreme points. These are the Anterior Extreme Point and the Posterior Extreme Point (AEP and PEP) which are used when transitioning from swing to stance or stance to swing in a forward walking animal, respectively. These are controlled by three separate control networks. The two microbehaviours, swing and stance, are mutually exclusive, meaning a leg can never be in both swing and stance at the same time. Therefore, two separate low-level control networks were developed in simulation, one for stance and one for swing, while transition between them is controlled by a selector network. While the stance and swing controllers are always active, the selector network determines which of them controls the motor output. This removes the need for a central oscillator, as the quasi-rhythmic movement of the leg is based on sensory feedback. Once the swing of a leg is decided, the required movement is sent to the motor controllers in the form of α, β, and γ angles, each representing the joints in each leg. A more detailed breakdown of each controller can be found in the schematic diagram below in Figure 31.

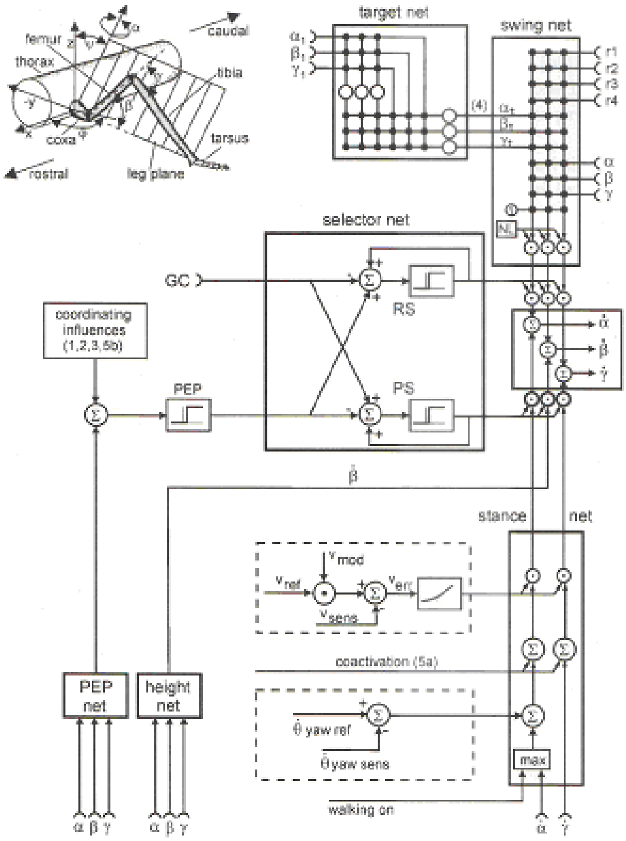

Figure 31 – Leg Controller consisting of three control networks: swing, stance and selector networks.

The control of the stance movements in each leg, or for the system as a whole, can prove to be an incredibly complex process. Constant sensory feedback is required for the angles of each of the joints in each of the legs (18 in total) as well as contact sensor feedback from the impact of each leg on the ground surfaces, and these are further complicated when the insect is on non-level ground or walking around a curve. For high-level machines (not simple insects or small robotic machines) this process can be achieved through large volumes of computations based on all inputs to optimise the stance position of the vehicle. However, real-time control with explicit calculation is far more difficult in a smaller creature or vehicle that cannot handle that complex computational environment.

There is a way to exploit much of the important data in the form of the dynamics of the vehicle's interaction with the environment. This would avoid the slow, computationally exact algorithm, by replacing a central controller with distributed control based on local feedback and the information already provided by physics. There are several problems involved in this process.

- First, positive feedback using raw position signal would lead to unpredictable changes in movement speed, as opposed to the constant speed that is usually desired. The addition of a band-pass filter into the feedback loop will help remedy this problem.
- Second, when using positive feedback for all three legs joints, unpredictable changes in body height would be caused. This is resolved when walking by providing positive feedback for the α and γ joints, but not for the β joints, which are the major determinant of body height as shown in the leg diagram in Figure 31.
- The third problem is when a positive feedback controlled vehicle is pulled backwards by gravity or through a tug from the experimenter, the insect will continue to walk backward even after the initial pull ends. To solve this potential problem, a supervisory system must be in place that is responsible for switching on and off the entire walking system, as well as determining walking direction.
- Finally, in a positive feedback system, the walking speed must be determined. In doing this, a central value is chosen and it is compared with the actual speed (measured by visual inputs or by monitoring leg movement). The error is then signals the positive feedback in the α and γ joints of all six legs.

Determining swing movement is a much simpler process, as it is mechanically decoupled from the body of the insect. In 1995, Cruse determined that a two-layer feed-forward net with three output units and six input units can produce movements that closely resemble the swing movements observed in walking stick insects. In simulation, these outputs are the angular velocities of the joints, $d\alpha/dt$, $d\beta/dt$, and $d\gamma/dt$, and the actual angles are measured and fed back into the swing net. Additionally, it is designed to compensate for external disturbances to simulate an avoidance behaviour observed in insects. For example, when a leg strikes an obstacle during the swing phase, it first tries to avoid it by retracting slightly and elevating briefly before renewing its forward swing from the new position. This requires an additional tactile or force sensors to signal when impact is detected. The end of the swing phase is determined by a sensory stimulus affecting the tarsal mechanoreceptors and possibly the load receptors, too. As soon as the stimulus is above the given threshold, the selector net switches from swing to stance.

They have determined a way that central controllers in the sense of central oscillators are not required in the development of simulated stick insect locomotion. Instead, the oscillations could be the result of the cooperation of different subsystems and the physical world. These results are favourable for slow walkers with a number of legs to provide passive stance stability, though possibly not for fast walkers, which cannot rely on temporally adequate sensory input. However, the major disadvantage to this system is that the simulations are purely kinematic in nature.

Cruse (2000) states that:

> Locomotion using multi-legged walking is a very interesting alternative to wheeled locomotion, in particular when he system has to move in cluttered and unpredictable terrain… because biological systems that can cope with these problems exist, we at least know that there are solutions to these questions.

He has developed neural networking software for model walking locomotion called Walknet. In doing so, he discusses four main complications to support how legged walking is non-linear:

- The combination of legs in stance varies continuously
- During curve walking, legs must move at different speeds

- On compliant substrates, the speed of the individual leg may vary unpredictably
- The geometry of the system may vary through growth and injury or due to non-rigid suspension of the joints.

Walknet is able to control walking at different speeds with an emerging tripod or tetrapod gait pattern, curve walking without any explicit calculation of leg trajectories, walking uphill and downhill, climbing over obstacles, and walking with partly amputated legs.

All of the above work by Cruse and company was performed on their Walknet locomotion simulator, designed after the investigations determined here. Additional tests were implemented into their system, as well. Some of these include an amputation test performed during walking, walking along curved trajectories, and climbing over obstacles. These results were compared to data collected from a stick insect accomplishing each of the same tasks, and the results were favourable. One unexpected behaviour was found, though. If the tarsi of three legs are clamped to the ground, the system will fall, as an extremely disordered arrangement of the six legs is created and the system would be unable to stand up on its own and resume proper walking. Future work in the development of a recovery method after an emergency like that is under development.

Through previous studies, it has been shown that insect walking is not controlled by a hardwired central system such as fixed action patterns or central pattern generators organized in a hierarchical structure. Instead there seems to be a decentralized architecture consisting numerous different and somewhat independent modules that have to be coupled to cooperate in a sensible way. Functional modules performing distinct tasks can be identified on different levels of integration and can be ordered from global to more local:

- Movement of the whole body (control of speed and direction)
- Coordination among legs (gait control)
- Control of different actions by a single leg (stance/swing phase, avoidance reflex, ground contact searching)
- Control of movement of each joint of a single leg performing a given action
- Control of synergistic and antagonistic muscles affecting movement of a single joint

Experiments in biology have shown that 1) each leg has its own control system which generates rhythmic step movements and 2) several coordinating mechanisms couple the movement of the legs to produce a proper gait. To control individual legs, no explicit central controller (such as central oscillators) are used. Instead, oscillations result from the cooperation of different subsystems and the physical world. In fact, using a central controller might worsen performance if such an autonomous central controller sometimes favours actions conflicting with those favoured by the present control model relying directly on the physical properties of the system.

Cruse partnered with M. Frik at the University of Duisburg to develop the TARRY robots. Additionally, there are conflicting experimental results that point to the involvement of negative feedback in some walking situations; therefore, the purely positive feedback simulations may not be entirely accurate. Finally, further work is necessary with respect to the control of swing movement in the simulation and the observations of step height appearing to be more or less independent of step length, which has yet to be incorporated in their simulation. [Cruse 2000].

Locomotion in animals with multi-jointed limbs requires the coordinated action of several appendages and their segmented limbs. The control of the actual motor output encompasses three levels: Intrajoint control, interjoint control, and intersegmental coordination between adjacent limbs. Sensory signalling from leg proprioceptors in the muscle plays a pivotal role in these tasks. These proprioceptors are the sensory receptors found in the muscle that are sensitive to stretching, tension and pressure. They provide information about ongoing limb movement and forces generated by the limbs and is utilised in generating a functional locomotor pattern [Bucher 2003].

3.5.1.3. Bipedal Walking

Bipedal robots require active balancing as they lack static stability. Humanoid robots such as the Honda P3 and the Sony SRD3 projects are bipedal corporate research efforts using body-mounted gyroscopes and accelerometers and force sensors at the wrist and feet to determine the body configuration to ensure that they do not become unbalanced. SRD3 uses two microprocessors to walk at a speed of 0.9 km/h while P3 uses a four-processor computer to walk at a speed of 2 km/h. P3 is 1.6 m tall and has a mass of 130 kg powered by a large battery and has a total of 16 joints. The hip joints provide the basis for shifting of body weight, knees for climbing obstacles and ankles for negotiating slopes. In static walking, the robot centre of mass must lie above the base of the foot in contact with the ground. For rapid walking, the centre of mass must be located between the supporting foot and moving foot.

Until recently, the humanoid robotic gait has been, for the most part, prescribed based on human motion. However, measuring angle trajectories during human walking for a wide range of step lengths and step times is difficult and time consuming. Additionally, the humanoid robot link dimensions and the number of distribution of degrees of freedom are not the same with those of humans. As a result, the angle trajectories recorded from humans need to be manipulated to fit to humanoid robot specifications. Capi and his team at Yamagata University in Yamagata, Japan are studying humanoid robotics and are proposing a new method for optimal gait generation based on genetic algorithms (GA) developed in-house. This gait is generated using two different cost functions: minimum consumed energy (CE) and minimum torque (TC). The results show that minimum CE gait is similar to the human gait.

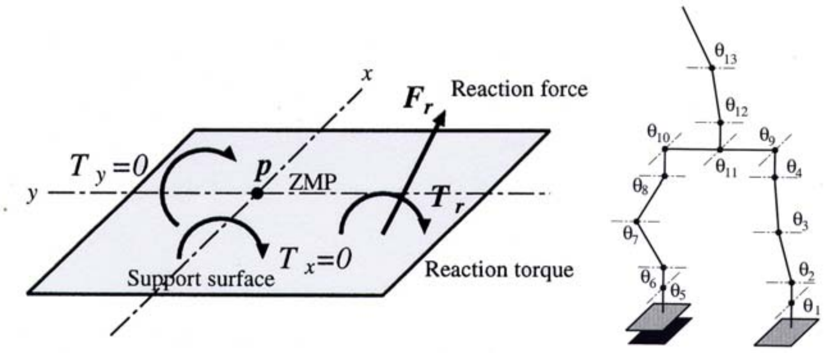

Figure 32 – a) Zero moment point (ZMP); b) Bipedal joint angles (Hasegawa et al 2000)

In real-time situations, the humanoid robot has to change gaits rather quickly, such as when switching from a straight walking gait to a stair climbing gait. GA was able to converge to the optimal gait, but time needed by GA was not suitable for real-time situations. In response, an approximation technique was developed that is computationally and algorithmically simple and can be trained such that over time, it becomes more accurate and faster. The Radial Basis Function Neural Network (RBFNN) is simply a curve-fitting approximation problem in a high-dimensional space. Its "training" is equivalent to finding the surface in a multi-dimensional space that provides the best fit to the data, based on statistics. The simulations and testing were performed using the following programmatic flow shown in Figure 33.

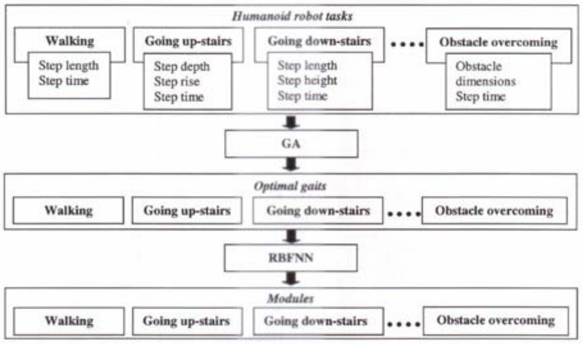

Figure 33 – The Programmatic Flow of the GA and RBFNN Modules to Determine Optimal Gait in Bipedal Walking (Capi 2003)

The cost function to determine energy consumed is:

$$J = \frac{1}{2}\left(\int_0^{t_f} \tau^T \tau \, dt + \Delta\tau_{jump}^2 \Delta t + \int_0^{t_f} C \, dt\right)$$

where:

t_f = step time
τ =torque vector
$\Delta\tau_{jump}$ = addition torque applied to the body link
C = constraint function

The torque vector is calculated from the inverse dynamics of five-link biped robot:

$$\tau = J(\theta)\ddot{\theta} + X(\theta)\dot{\theta}^2 + Y\dot{\theta} + Z(\theta)$$

where:

$J(\theta)$ = 5x5 mass matrix
$X(\theta)$ = 5x5 matrix of centrifugal coefficients
Y = 5x5 matrix of Coriolis coefficients
$Z(\theta)$ = 5x1 vector of gravity terms
τ = 5x1 generalized torque vector
$\theta, \dot{\theta}, \ddot{\theta}$ = 5x1 vectors of joint variables, angular velocities, and angular accelerations

The GA and RBFNN were implemented both in simulation and through the use of the Bonten-Maru I humanoid robot, as modelled in Figure 34. He stands 1.2 metres tall with each leg consisting of three segments (upper leg, lower leg, and foot) totalling 6 degrees of freedom. In experiment, the optimal gait was converged on by the GA after the 33rd generation.

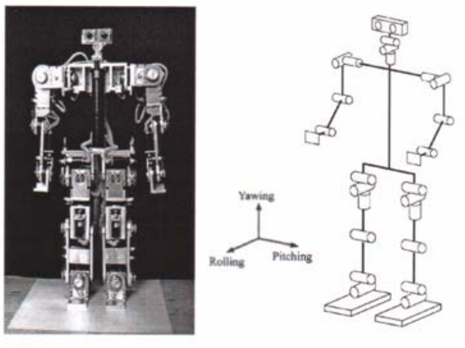

Figure 34 – Bonten-Maru I Humanoid Robot used for GA and RBFNN Control Testing

Based on simulation and experimentation, Capi and his team can conclude that the difference between the RBFNN and GA gaits is very small and the time needed to run the RBFNN walking module to generate a solution is very short. Therefore, the optimal gait generated by RBFNN is stable and accurate for rapid walking gait determination [Capi 2003].

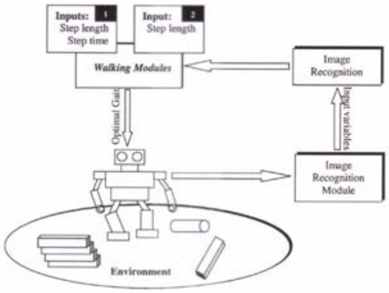

Fig. 4. Real time gait generation during walking.

Figure 35 – Diagram of Real-time Gait Generation During Walking (Capi 2003)

The Japan Science and Technology Corporation is working on a different biped walking robot. Since April of 1996, they have been working on the ESYS humanoid project. Their purpose is to develop and integrate component technologies, to implant them into humanoid robots, and to develop practical humanoid robots that can perform various human-like activities. Their current concentration is on mechanical hardware and control strategies, though they are also actively involved in the development of many component technologies such as real-time robot vision, hybrid behaviour control architecture and control hardware.

When designing a humanoid body, Futura and his team consider the following as the most important factors:

- Limited torque of actuators at the joints in the robot
- Rigidity of hardware such as main frames, arms, legs
- Impulsive loading at landing of legs that can damage hardware

- Interference among links
- Weight
- Ease of handling

The walk controller must have the following characteristics:

- Stability in the sense of continued dynamic walk
- Practicability with starting, stopping, and making turns
- High-energy efficiency
- Versatility in choosing a landing position of the swinging leg

In their design and experimentation, high rigidity and low weight have been accomplished through reinforced plastics and carbon fibre reinforced plastics as the major structural materials.

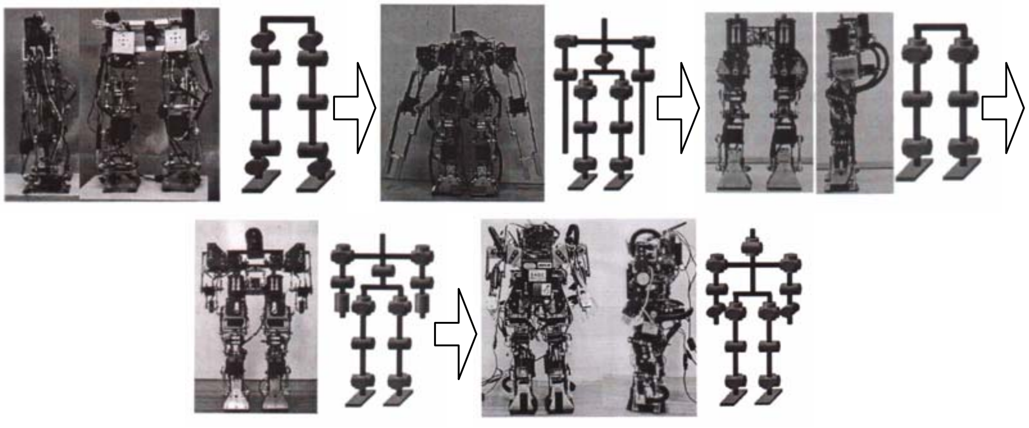

Figure 36 – Mk.1 through Mk.5 Biped Robots (Furuta 2001)

Their robots, none greater than 40cm in height, have been modified since the initial lower-body Mk.1 robot was built in 1996. Most recently, the full-body Mk.5 robot, weighing in at 1.9 kg and standing 35.6cm in height. It can perform all the 3D dynamic walk of the previous robots, as well as transfer objects between its arms. They hope that their development of the Mk.5 robot will become the standard in biped locomotion research and control system testing. Pictures and degree-of-freedom diagrams of each robot are shown in Figure SDFS, as well as a table of their specifications to compare capabilities with each incremental robot [Furuta 2001].

Table 10 – Technical Specifications of Mk.1 through Mk.5 (Furuta 2001)

	Mk.1	Mk.2	Mk.3	Mk.4	Mk.5
Degrees of freedom	10	19	12	23	24
Height (m)	3.25×10^{-1}	3.80×10^{-1}	3.03×10^{-1}	3.63×10^{-1}	3.56×10^{-1}
Leg length (m)	2.81×10^{-1}	2.29×10^{-1}	2.11×10^{-1}	2.09×10^{-1}	2.02×10^{-1}
Mass (kg)	1.286	2.500	1.400	2.000	1.900
Main material	FRP	CFRP	CFRP	CFRP	CFRP

3.5.1.4. Hexapod Walking

The most common configuration is the insect-inspired six legs which always ensure that three legs are in contact with the ground in motion, though Scorpion was an eight-legged planetary rover which was controlled by a central pattern generator type controller [Kirchner et al 2002]. Hexapod vehicles function in a statically stable mode while others have active balancing for dynamic stability. Insects are usually highly statically stable with a low centre of mass and wide support base. They usually have at least three of their six legs in contact with the ground at all times, a tripod-based gait. Much of locomotor control is imparted by the mechanical design of the locomotor system with hopping capacity provided by the storage of elastic strain energy. Six legged walkers – either insect-type or omnidirectional three degree-of-freedom configuration - are a common design approach [McTamaney et al 1988].

Legs may vary in complexity from simple hexapod type designs to complex articulated legs.

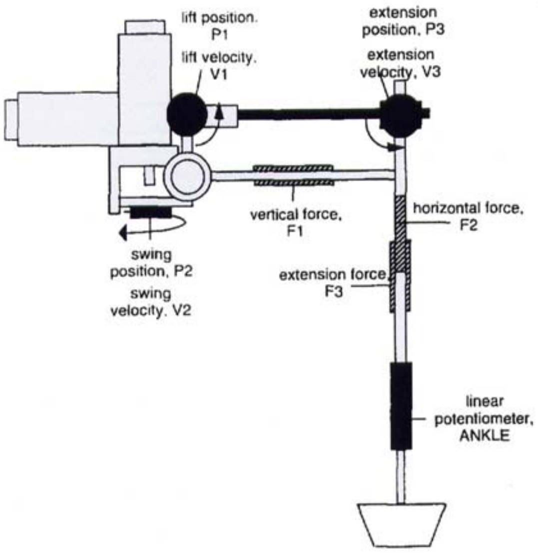

Figure 37 – Hannibal's leg: Hannibal has 60 sensors, 8 computers and six legs each with three degrees of freedom – swing DOF advances and retracts the leg by rotating it about the shoulder P2, lift DOF raises and lowers the foot by rotating the upper link about shoulder P1, and extension DOF lifts foot by extending lower link about elbow (P3) (Ferrell 1995)

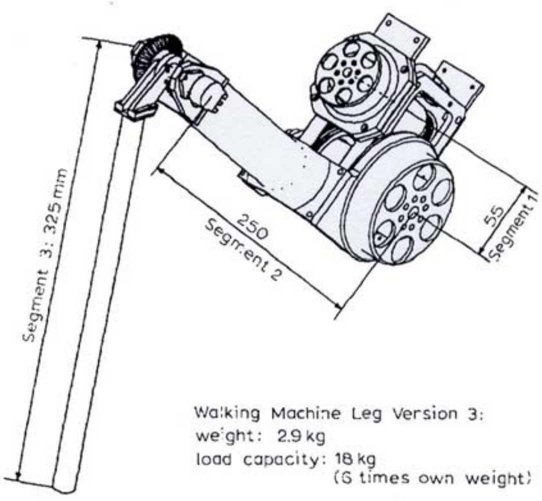

Figure 38 – Optimised walking robot leg (Pfeiffer et al 1995)

The forward extension d of an animal's progression is a fraction η of its length r where η~0.3. The animal minimises the lift height of the leg which is determined by the roughness of the terrain. Terrain traversibility is determined by the size of the animal's body whereby h=l/λ where λ~50 (so the leg cannot be lifted more than 1/50 of the animal body length). Three legs are the minimum required for static stability but a fourth leg provides static stability while in motion.

There has been much research activity in legged robots by a number of research laboratories throughout the world. Most have been concerned with particular issues associated with legged locomotion such as central pattern generators, etc. Some of these projects are discussed below, siting the important gains from each.

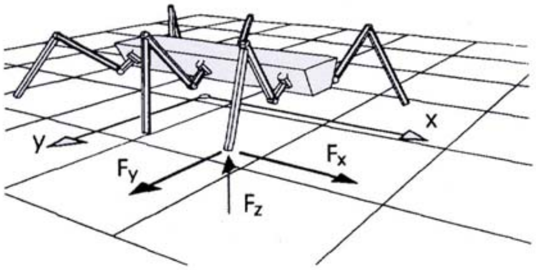

Figure 39 – a) Walking robot model (Pfeiffer et al 1995)

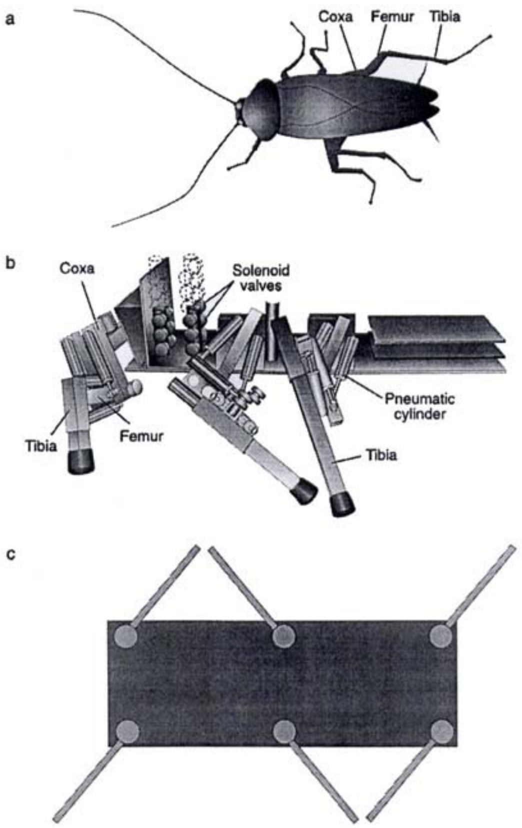

Figure 40 – Position of legs during walking in insects a) in cockroach, angled attachment of legs to the body and different sizes of front middle and rear legs allow the legs to be positioned close to one another on the body without mechanical interference; b) robotic cockroach analogie with similar leg arrangement; c) conventional leg placement which requires significant distance between legs to avoid interference (Delcomyn & Nelson 2000)

3.5.1.4.1. NASA

NASA's hexapod robot Ambler can walk at 1m/minute. Hexapods often adopt a tripod gait where there is a 180^{o} phase difference between the two sets of legs. This mode of walking is inflexible. Both force and kinematic sensing at high sample rates ~20-100 Hz is required for legged locomotion with active suspension [Waldron 1995]. Each foot requires a contact sensor and rotary position sensors are required at each leg joint. Almost all legs require a two degree-of-freedom hip joint (in order to turn without slippage) plus one single degree-of-freedom knee parallel to the hip axis for full terrain adaptability.

Another biomimetic hexapod robot project is being built by NASA's Jet Propulsion Lab: a 6-legged Lemur-inspired robot prototype to perform dexterous small-scale assembly, inspection and maintenance of macro space facilities shown in Figure 41a. The carbon-fibre based structure weighs in at less than .5 kg, half of what a similarly designed aluminium structure would weigh. The high-strength graphite cyanate ester prepreg facesheets contain a honeycombed Nomex core, giving a very strong, but very light (10% of the total system mass) system.

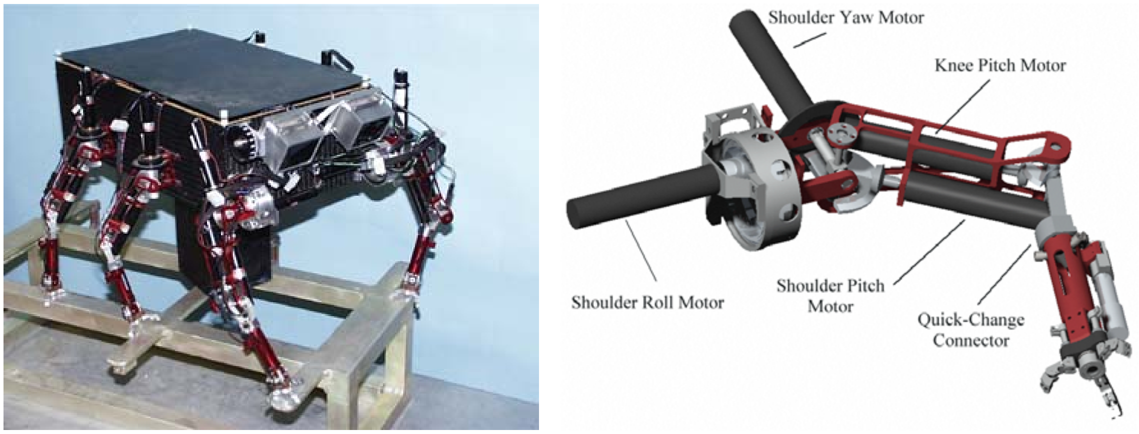

Figure 41 – a) NASA JPL's Lemur Robot; b) Close-Up of the Limb

The front two limbs are 4-DOF while the rear four are each 3-DOF. Although all 6 limbs are designed to have interchangeable "hands" to use a variety of tools, the front 2 limbs are primarily used in the repair and assembly operations; the additional DOF joint allowing for work to be done in direct line to the onboard cameras to assist in remote operations and control. The limb design was dictated by the need to decrease the overall mass of the limbs and create as simple a mechanism as possible. In addition, the limbs needed to incorporate whatever the structure to allow the quick change-out of end effectors. The design of each load-bearing part was passed through several cycles of finite element analysis (FEA), which yielded an extremely lightweight system. Each of LEMUR's degrees of freedom is completely independent, simplifying the control algorithms for movement. In addition, the added degree of freedom in the front shoulder is oriented to create a kinematically spherical joint (3 joint axes are mutually perpendicular and intersect at the same point in space), further simplifying the limb kinematics and minimizing singularities.

An innovation in the limb design is the inclusion of a quick-connect for the tools. While quick-connects are commonplace in industrial robotics, they are a novelty in a system as small as LEMUR. LEMUR's quick-connect is a simple bayonet-style connection that is spring-loaded to both lock the tool in place and to eject the tool when the tool is twisted to the release point. There were several tools that were prototyped into the design. Some examples include the Gripper Palm (Figure 41b) with three fingers used for walking and gripping with each step. Also, a Ball Driver tool can be used to assist the assembly and repair process (like a complex screwdriver) as well as having the ability to act as a load-bearing leg in case the Lemur needs to relocate from its initial position while performing a repair. Finally, a fibre-optic Foot Camera can be used as both a low-volume microscopic and a small video camera [Trebi-Ollennu 2002]

In 2002, Robert Hogg, an engineer at JPL's Mobility Systems Concept Development Section, and his team built Spider-bot, a 6-legged micro-robot prototype. The team envisions this as a prototype for an explorer that can "go anywhere any time" to explore whatever terrain necessary. They see additional benefit of one day hopefully doing maintenance and repairs on the International Space Station. A similar styled vehicle on Earth might be able to fill in for humans by investigating hazardous materials or taking soil measurements on farms.

Figure 42 – NASA JPL's Spider-bot

Like a real spider, this robot has feeler-like antennas to help detect various obstacles. Although the prototype is sized to fit in the palm of one's hand, future visions can be expected to be on-tenth the size and still equipped with cameras and other scientific experiments. Although this robot was developed with 6 legs, depending on the kind of work needed to be performed, the same design can be applied to have orders of magnitude more.

This new approach to robotic explorers allows robots to be reconfigurable, adjustable, and flexible to stand on their own. He notes that "Evolvable, reconfigurable structures are key to changing the paradigm for future space missions by increasing their functionality while reducing cost by becoming a multi-use robot." [Martinez 2002].

3.5.1.4.2. Case Western Reserve University

The Biorobotics group at Case Western Reserve University in the United States has two primary goals in their development of mobile robotics:

1. they strive to design and implement agile and useful legged robots
2. they attempt to further their understanding of legged locomotion in insect species

To accomplish this task, they have selected a cockroach as the baseline biological creature for their main robot project due to the enhanced leg design and increased agility. They used a high-speed camera to examine how the cockroach leg joints worked when running on a treadmill and climbing over barriers. They determined that the forward legs are used to investigate the surroundings and provide braking power, the middle legs provide steering while climbing, and the powerful hind legs are used to push the animal forward. All important natural degrees of freedom within the legs were kept, though for simplicity in the hind legs, some were removed as they are mainly used for forward movement. As a result, they created a pneumatic cockroach robot that measures 30 inches long (which is 17 times the size of a typical Blaberus cockroach) with 24 active DOF controlled by 48 three-way valves.

However, there are two significant problems in their development of a pneumatic robot. First, it does not have the ability to store air, therefore passive stiffness is not available. Second, like most pneumatic devices, when applying air to the system to actuate a "muscle", it expands from the extreme empty to the extreme full very quickly. In order for this to be a more effective model of a cockroach, a control system would need to be developed to smoothly extend the leg joints.

There are three controllers that make up the control system of the cockroach robot. The first is the posture controller, which positions the CoM of the body by directing forces with the legs that are in contact with the ground. While the robot is standing, adjustments to posture do not need to be made exceedingly quickly, however, during walking and running, the posture controller needs to interact with the swing and stance controllers and must act reasonably rapidly. For posture control, the forces and moments on the body required to keep proper posture can be represented by the spring equation:

$$F_i = K_i\left(actual_i - desired_i\right)$$

where $actual_i$ is the actual body position in the i direction and $desired_i$ is the desired body position in the i direction, where i represents any of the three directions x, y, or z. The posture controller then assigns load-bearing responsibilities to the stance legs such that their summation causes F and M to be exerted on the body. However, the number of stance legs typically varies between 3 and 6 during locomotion, therefore the number of equations needed to solve this load distribution problem varies with the gait and the phase of the gait.

The swing controller positions the legs during walking, running, climbing and any other locomotion behaviour, specifically at the onset of the stance. In the case of a kinematically redundant leg (where the same leg forces can be applied on different joints to achieve the same overall stance), the controller is designed to interpolate between the limits of each of the leg joints, providing a smooth, safe, and efficient leg movement. The swing controller incorporates pulse width modulation (PWM) control of the 48 valves, which allows for a smoother.

The stance controller is the third control system and it determines the ground reaction forces to assist the posture controller. The stance controller is designed to react more quickly than the posture controller and is used to cause a forceful extension of joints until the tarsus attains a predetermined posterior extreme position (PEP).

In the future, this group is interested in investigating McKibben artificial muscles or braided pneumatic actuators (discussed in Section 3.3.2.1.1) or other upcoming technologies by SRI and Honeywell, but they are not yet at a point where they have sufficient power for their weight and size. Further fine tuning of the control architecture is also important to the further development of their cockroach robot.

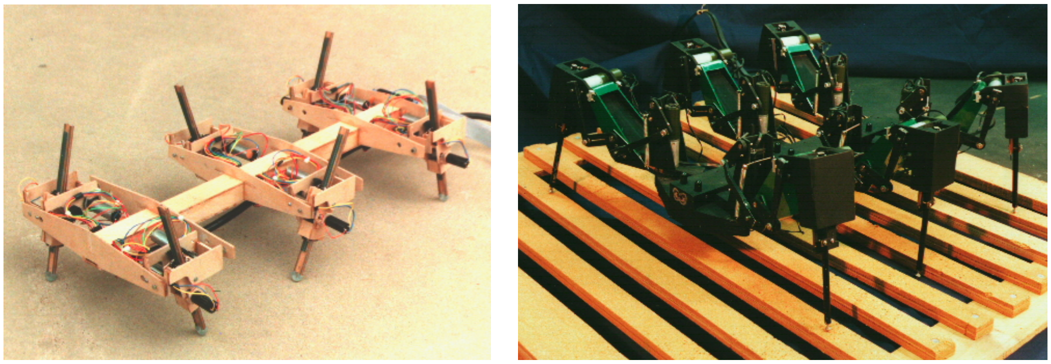

Figure 43 – Robot I and II from Case Western Reserve University

3.5.1.4.3. University of Michigan

RHex is a 6-legged walking robot with a single degree-of-freedom in each leg. In this robot, the single DOF legs are actuated by a motor located at each hip running a very simple clock-driven, open-loop tripod gait. The legs rotate in a full circle and the three legs of a tripod are driven simultaneously through a slow "retraction" phase, as the opposing tripod begins its "protraction" by rotating away from the ground contact.

Figure 44 – RHex from University of Michigan

This robot can move at speeds of 1 body length per second and can climb obstacles well exceeding its body clearance. It uses an alternating tripod gait for forward and backward locomotion, but also allows for more elaborate behaviours such as leaping, stair climbing, etc [Saranli 2001].

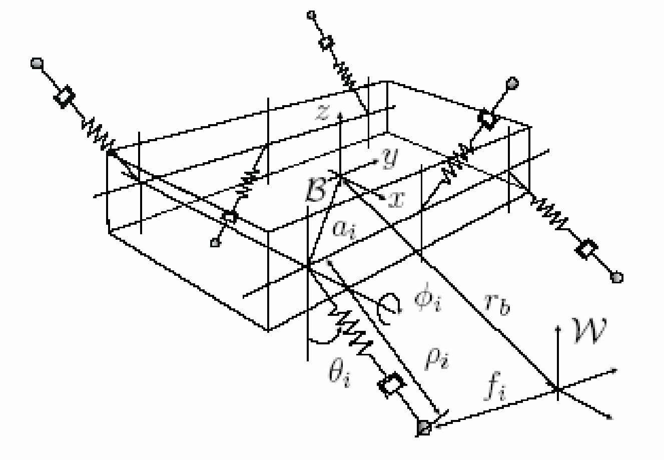

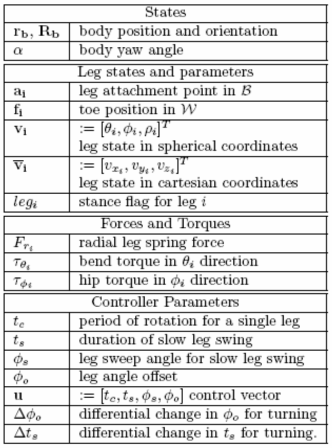

States	
$\mathbf{r_b}, \mathbf{R_b}$	body position and orientation
α	body yaw angle
Leg states and parameters	
$\mathbf{a_i}$	leg attachment point in $\mathcal{B}$
$\mathbf{f_i}$	toe position in $\mathcal{W}$
$\mathbf{v_i}$	$:= [\theta_i, \phi_i, \rho_i]^T$ leg state in spherical coordinates
$\overline{\mathbf{v}}_\mathbf{i}$	$:= [v_{x_i}, v_{y_i}, v_{z_i}]^T$ leg state in cartesian coordinates
leg_i	stance flag for leg i
Forces and Torques	
F_{r_i}	radial leg spring force
τ_{θ_i}	bend torque in θ_i direction
τ_{ϕ_i}	hip torque in ϕ_i direction
Controller Parameters	
t_c	period of rotation for a single leg
t_s	duration of slow leg swing
ϕ_s	leg sweep angle for slow leg swing
ϕ_o	leg angle offset
$\mathbf{u}$	$:= [t_c, t_s, \phi_s, \phi_o]$ control vector
$\Delta\phi_o$	differential change in ϕ_o for turning
Δt_s	differential change in t_s for turning.

Figure 45 – RHex compliant Hexapod Design Schematic and Nomenclature

Using the diagram and nomenclature in Figure 45, the equations of motion for the 1-DOF legs were calculated as follows:

$$\mathbf{F}_i = \begin{bmatrix} -\cos\theta_i \sin\phi_i & \sin\theta_i \sin\phi_i & -\cos\phi_i \\ \sin\theta_i & \cos\theta_i & 0 \\ \cos\theta_i \cos\phi_i & -\sin\theta_i \cos\phi_i & -\sin\phi_i \end{bmatrix} \cdot \begin{bmatrix} F_{r_i} \\ \tau_{\theta_i} / \rho_i \\ \tau_{\phi_i} / (\rho_i \cos\theta_i) \end{bmatrix}$$

$$\tau_i = (\overline{\mathbf{v}}_i + \mathbf{a}_i) \times \mathbf{F}_i$$

$$\mathbf{F_T} = [0 \quad 0 \quad -mg]^T + \mathbf{R}_b \sum_{i=1}^{6} leg_i \mathbf{F}_i$$

$$\tau_{\mathbf{T}} = \mathbf{R}_b \sum_{i=1}^{6} leg_i \mathbf{F}_i$$

We suggest that the elusiveness of robust legged locomotion requires a multiple approach. Unlike traditional engineering approaches, biological solutions are integrative involving simultaneous co-evolution of mechanical structures with control systems to provide their functionality. To that end, we suggest that robust legged locomotion will be possible only with a fully integrative solution involving central pattern generators, sensory feedback and the implementation of preflexes through structural compliance.

3.5.1.5. Other Forms of Walking

Aibo is a 1.6 kg robo-dog with 18 degrees of freedom driven by a 64-bit processor based on software modules to generate movement.

Slightly different to the Case Western Reserve project described above, Kirchner and his team have selected the 8-legged scorpion as the inspiration for their robot. They discuss four major components of a scorpion leg as listed below and shown in Figure 46:

1. Thoracic Pivot: mediates protraction and retraction
2. Basalar Segment: contains the pivot mediating elevation and depression
3. Middle Segment: contains the pivot mediating extension and flexion
4. Distal Segment: terminates in the limb tip and its sensors

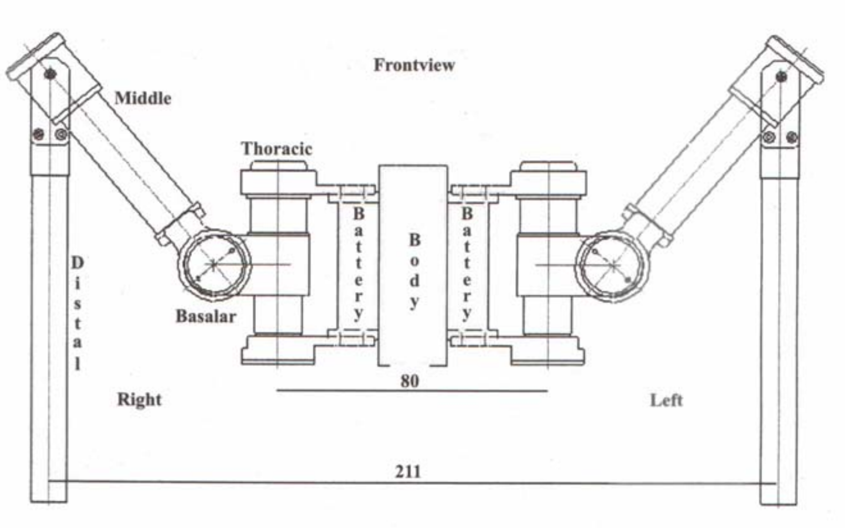

Figure 46 – Scorpion Robot Schematic, Front View

Each leg is the same design that provides 3-DOF, is actuated by standard DC geared motors, and creates a 1:8 weight to lift ratio. This is possible with through the use of high-ratio planetary gears along with the DC motors, which helps the robot climb obstacles that exceed its own height. Leg speed was also a significant factor in the leg design. Reactive speed and reflexes need to be fast in order for the legs to compensate for the robot's own inertial energy when exposed to gravity.

Kirchner's scorpion robot is equipped with two types of sensors: proprioceptive sensors to measure the internal status of the robot and exteroceptive sensors to map states of the external environment onto internal states of the robot.

The scorpion robot is equipped with one main/central computer and many local computing units. This allows the local controllers to perform the low-level tasks (such as driving the legs, reading in sensor data, and executing leg reflexes when needed) which reduce the computing load on the central computer so it is allowed to run the main program. The local and central boards are connected through a serial CAN-bus system, allowing individual addressing and bidirectional data exchange between all components of the network at speeds up to 1 Mbit/s. The local leg boards include high-amp H-Bridge motor controllers (each controlling up to 4 motors at up to 5 amps) and CPLD[2] chips to monitor motor position and current draw, as well as provide real-time feedback control. The central controller runs the master control program to provide high-level behaviours of the robot. In future designs, they are planning to use the Motorola MPC55 microcontroller version of the PowerPC in combination with a free programmable gate array (FPGA) to take over the tasks currently implemented by the local controllers.

Scorpions almost always use the same gait, L4-R3-L2-R1 and R4-L3-R2-L1, as shown in Figure 47a:

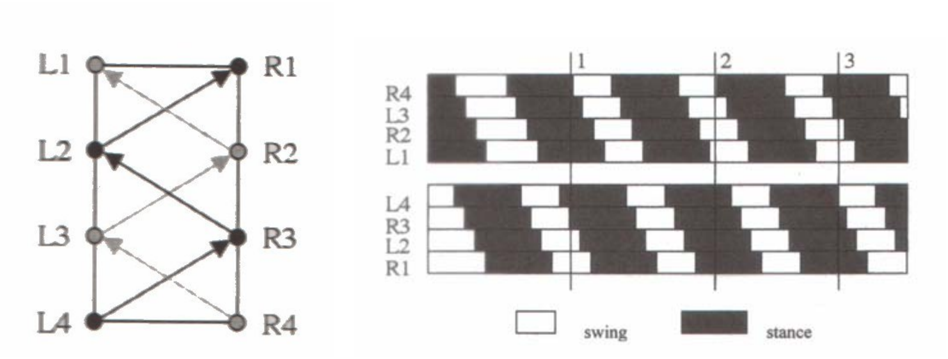

Figure 47 – a) Scorpion Gait Pattern and b) Swing and Stance Pattern

This shows that the rear left leg (4th position) starts, followed shortly afterward by right leg number 3, then left 2, then right 1. This system is repeated on the opposite side starting with right leg number 4 after right leg 1 is placed firmly on the ground.

In the stance phase, forward motion of the entire scorpion is dependant on the ground contact of each leg, i.e. the leg itself moves backward with respect to the body. In the swing phase, the leg is lifted from the ground and moves forward with respect to the body. The sum of time for swing and stance is called the period with the step frequency being notated as 1/period. The phase reflects the delay of motion patterns between two contralateral adjacent legs (i.e. L4 and R3) and is a small fraction of the period (typically 0.1 in scorpions). Figure 47b shows an example of a real scorpion walking gait where swing=1.5, stance=2.5, and phase=0.1. The first step establishes the phase shift from standing still to walking, and the second step establishes the phase shift of 50% between the both groups of legs.

While moving, the robot should be moving in the desired direction, without any shift in any other direction. Therefore, the joints of each leg have to harmonize their rotations at ever moment of a stance. To do this, a straightforward set of equations of motion must be established based on the free body diagram shown in Figure 48.

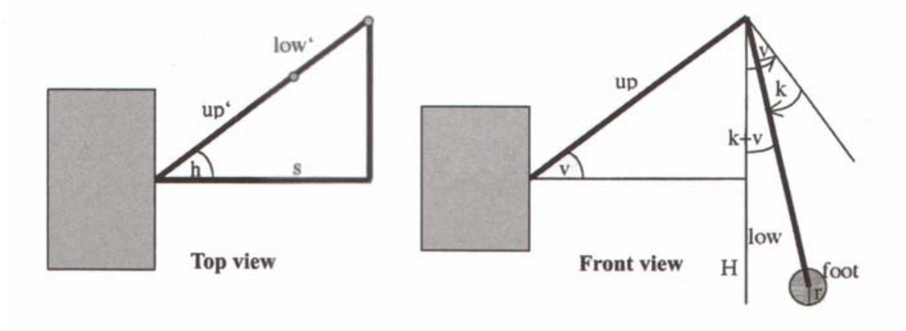

Figure 48 – Free body diagram of the forces, angles, and distances required to determine proper joint angles during locomotion

where:

v = vertical hip-joint angle
k = knee joint
h = horizontal joint
up, low = length of leg
v1, k1 = initial angles
v_i, k_i *= vertical angles for h=0*
r = radius of the foot
s & H = initial posture of R4-L3-R2-L1

The lateral distance between foot and hip is:

$$s = up \cdot \cos(v_i) + low \cdot \sin(k_i + v_i) = \cos(h) \cdot (up \cdot \cos(v) + low \cdot \sin(k + v))$$

$$\sin(k + v) = \frac{s / \cos(h) - up \cdot \cos(v)}{low}$$

The vertical distance between floor and hip is:

$$H = -up \cdot \sin(v) + low \cdot \cos(k + v) + r$$

$$\cos(k + v) = \frac{H + up \cdot \sin(v) - r}{low}$$

Once they determined what forces and movement were required to move each leg accurately, the control architecture was created. This architecture is based on a combination of the Central Pattern Generator (CPG) and a pure reflex-driven approach. The CPG is used to create a rhythmic motor pattern for the high-level movement without any sensory feedback. A general model of a CPG has been found in almost every species, so this baseline is quite accurate.

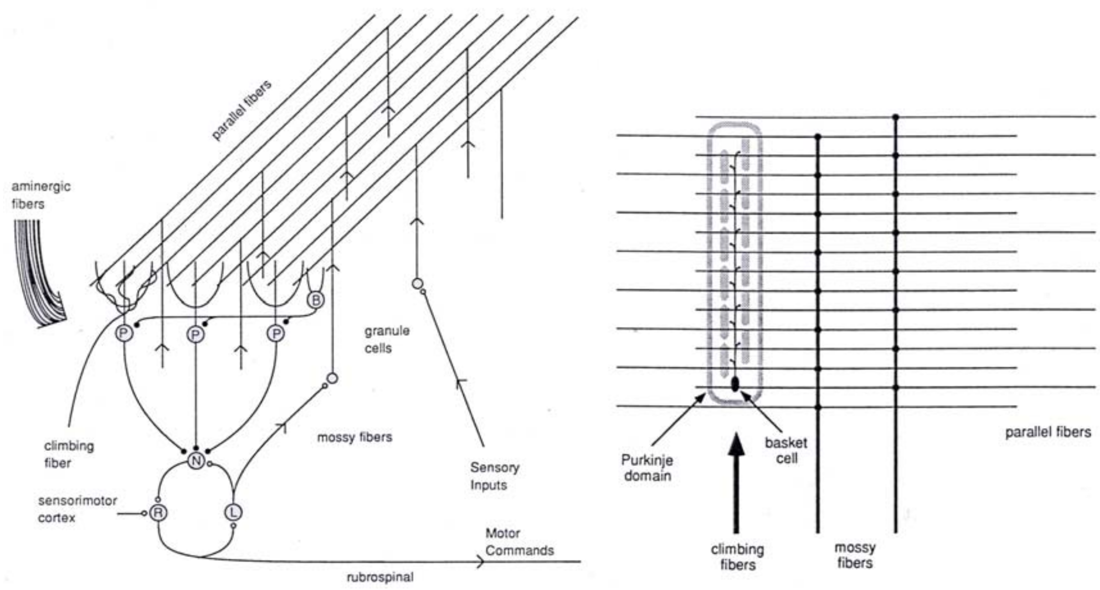

Figure 49 – a) Adjustable cerebellorubrospinal pattern generator neural circuit with positive feedback loop from cerebellar nuclear cells to red nucleus; inhibitory input from Purkinje cells of cerebellar cortex regulates feedback loop intensity while basket cells inhibit Purkinje cells; b) mossy fibres distribute inputs to granule cells; each granule cell axon bifurcates to form parallel fibres and innervate large number of Purkinje cells; dendrites of Purkinje cells fan out to innervate nuclear cells but are innervated by climbing fibres (Houk et al 1996)

CPG is commonly implemented through CPG neuron whereby the neural state is governed by:

$$C_i \frac{dV_i}{dt} = -\frac{V_i}{R_i} + \sum w_{ij} f(V_j) + I_{\text{int}} + I_{ext}$$

where:

Vi = neuron input voltage
Ri = neuron membrane resistance
Ci = neuron membrane capacitance
wij = connections weights
f(Vj) = saturating linear thresholding fuinction, eg. sigmoid
Iint = intrinsic current
Iext = extrinsic input current

Intrinsic current switches between high state of fixed duration (dependent on synaptic excitatory/inhibitory inputs) and low state.

The reflex-driven approach then provides stabilization to the walking pattern provided by the CPG and allows a fixed situation-reaction set of rules based on external disturbances to the high-level pattern.

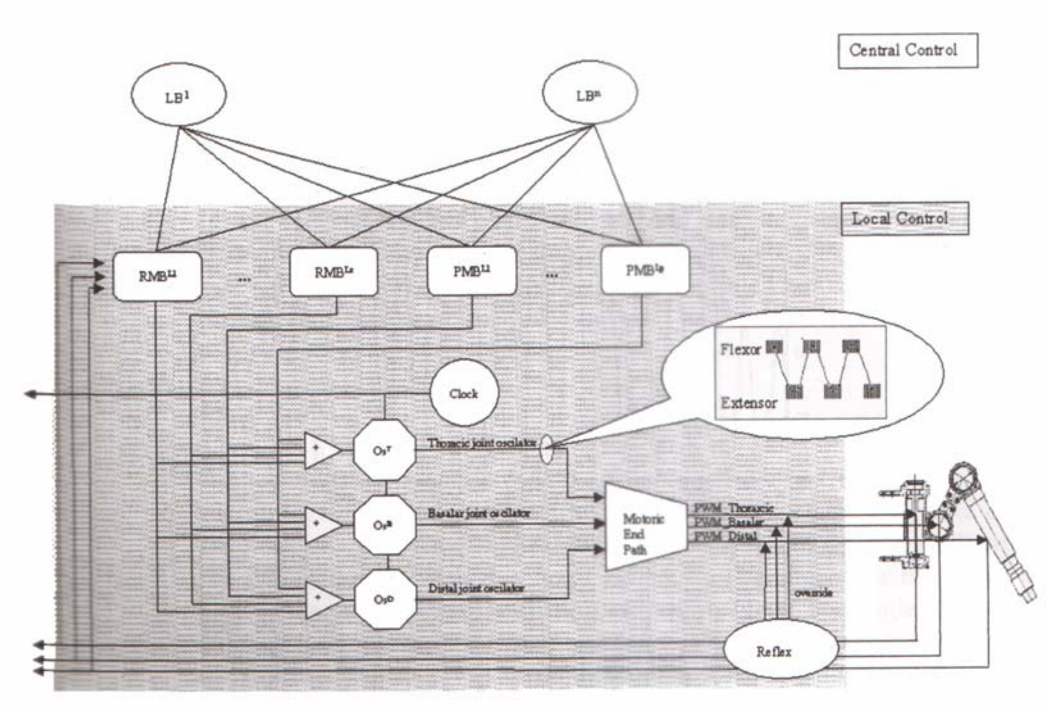

Figure 50 – High-level Control Architecture for the Scorpion Robot

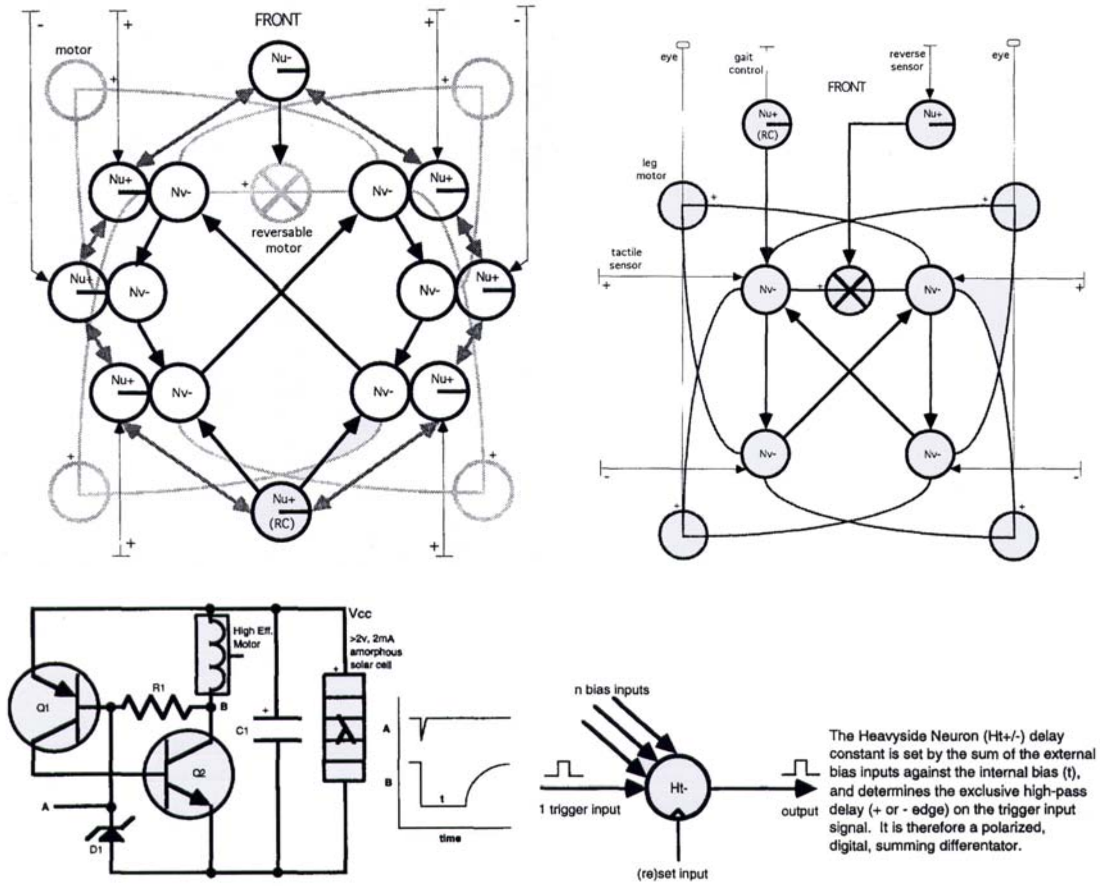

The Heavyside Neuron (Ht+/-) delay constant is set by the sum of the external bias inputs against the internal bias (t), and determines the exclusive high-pass delay (+ or - edge) on the trigger input signal. It is therefore a polarized, digital, summing differentator.

Figure 51 – Tilden's nervous net system for solar powered micro-robots based on CPG generation: a) "lobster-robot" neural structure; b) "walkman" neural microcore structure; c) solar-engine; d) heaviside motor neuron as summing differentiator (Hasslecher & Tilden 1995)

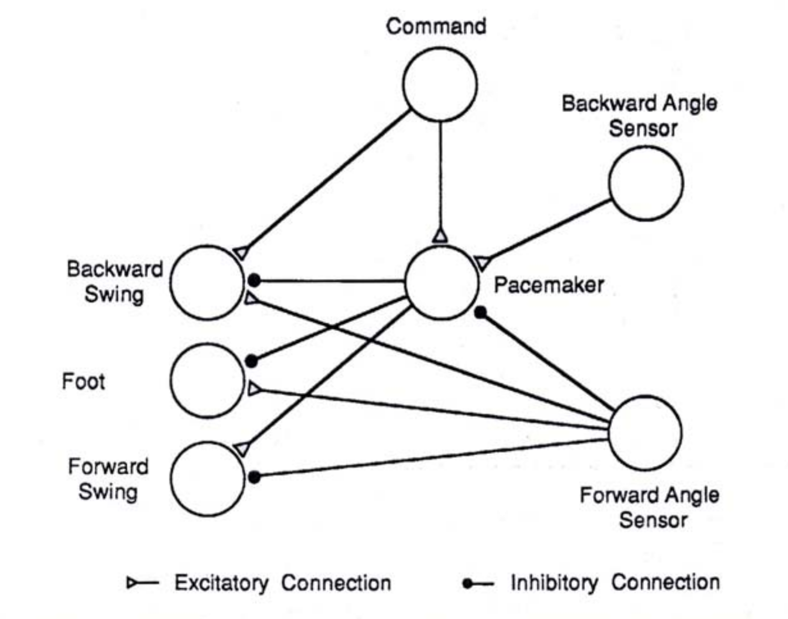

Figure 52 – Leg control circuit – each leg is monitored by two sensory neurons which signal extreme forward/backward positions; three motor neurons control foot state and leg forward and backward swing velocities; motor neurons are driven by pacemaker neuron which has oscillatory output; single command neuron makes same two connections on every leg controller (Beer et al 1992)

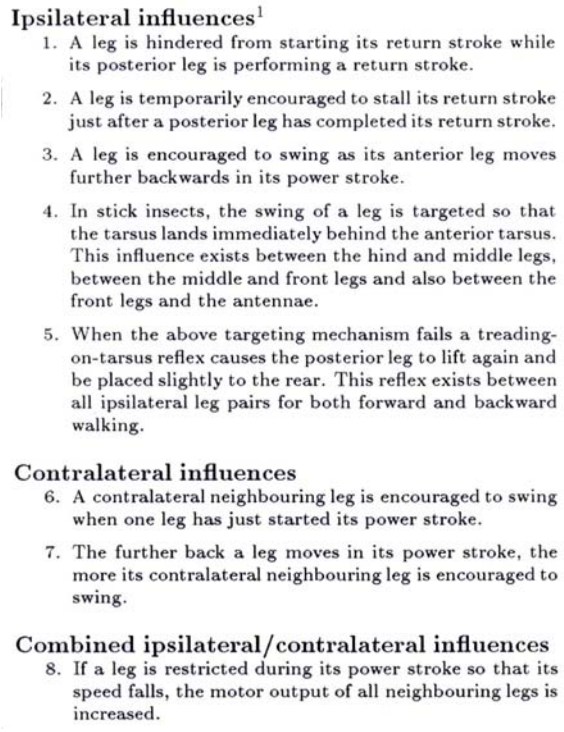

Ipsilateral influences[1]

1. A leg is hindered from starting its return stroke while its posterior leg is performing a return stroke.
2. A leg is temporarily encouraged to stall its return stroke just after a posterior leg has completed its return stroke.
3. A leg is encouraged to swing as its anterior leg moves further backwards in its power stroke.
4. In stick insects, the swing of a leg is targeted so that the tarsus lands immediately behind the anterior tarsus. This influence exists between the hind and middle legs, between the middle and front legs and also between the front legs and the antennae.
5. When the above targeting mechanism fails a treading-on-tarsus reflex causes the posterior leg to lift again and be placed slightly to the rear. This reflex exists between all ipsilateral leg pairs for both forward and backward walking.

Contralateral influences

6. A contralateral neighbouring leg is encouraged to swing when one leg has just started its power stroke.
7. The further back a leg moves in its power stroke, the more its contralateral neighbouring leg is encouraged to swing.

Combined ipsilateral/contralateral influences

8. If a leg is restricted during its power stroke so that its speed falls, the motor output of all neighbouring legs is increased.

Figure 53 – Walking Rules (Barnes 1995)

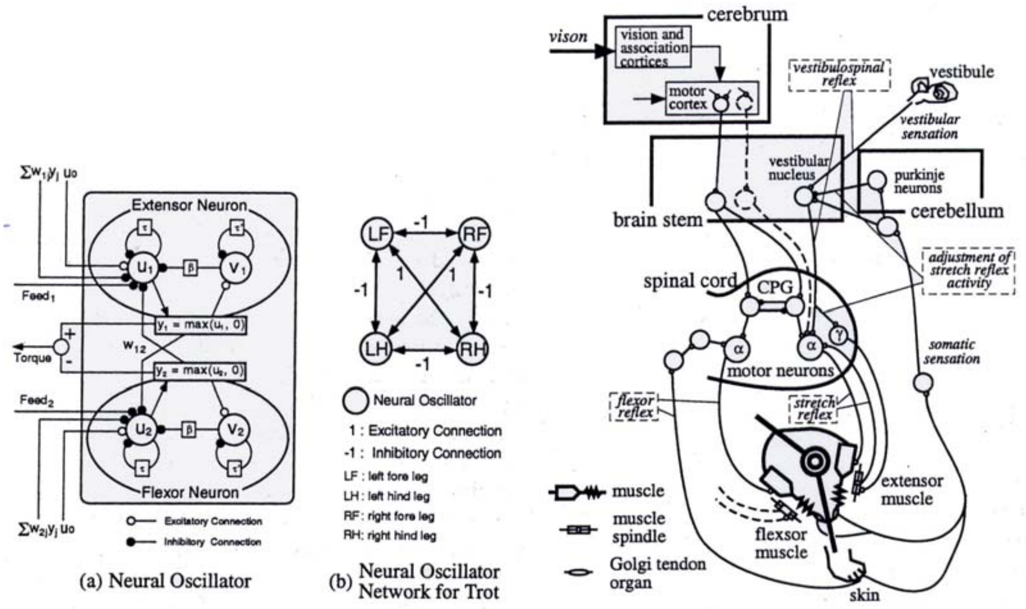

(a) Neural Oscillator

(b) Neural Oscillator Network for Trot

Figure 54 – a) neural oscillator; b) neural oscillator network as a model of CPG ; c) Simplified nervous system for adaptive control of legged locomotion (Kimura et al 2000)

3.5.2. Snake Locomotion

Body compliance is taken to its extreme in the case of snakes and centipedes. A worm moves forward through soil by elongating its front end and pushing forward against the resistance of an expanded rear section. Centipedes in particular are constructed from cuticle (qv) with a large number of compliant joints for high compliance. Following forward motion, the front section expands and extends bristles to resist backward motion. The snakebot concept is the ultimate form of body articulation and has good adaptability to rough terrain, stiffening into a bridge to cross crevasses, access to orifices, winding around obstacles, low path cross-section requirement, redundant structure in case of joint failure and modular construction. The motion comprises consecutive folding/unfolding of the body with high degrees of friction between the body and the ground. The snake moves by passing a sine wave laterally through its body while simultaneously curving the body so that the peaks are lifted off the ground. Locomotion is most efficient with highly elastic bodies which allow longer contraction/elastic waves to be accommodated. The force exerted by a snake perpendicular to the direction of motion is limited by the transverse friction of the body with the ground. This must overcome the friction force in the longitudinal direction. A number of different snakebot "gaits" are possible such as side-winding, coiling or flipping. Furthermore, the snakebot approach offers the capabilities of a manipulator for precise pointing of the "head" which may incorporated sensor heads to scientific instruments, or as a gripper by coiling around an object of almost any shape. Snake locomotion is relatively slow and inefficient. All the modular segments of a snakebot's body are essentially identical, hinged together into a chain and each hinge must typically be powered by a motor, gearbox, control electronics and power source within each segment. However, the use of shape memory alloy wires of NiTi at each joint counter-balanced by springs provide the basis for flexure between segments. Electrically heating the wire causes NiTi wire to contract causing the joint to bend. Cooling assists the spring to return the joint to its original configuration. The cyclicity of such a flexure may be as high as 5 MHz. The addition of strain sensors or artificial skin would enable the snakebot to adopt reflexive behaviours. The chief disadvantages lie in the control of locomotion and friction between the snakebot and the terrain as frictional forces are proportional to the sliding surface area [Miller 2002]. Furthermore, hydrostatic bodies such as snakes cannot exert high point forces.

Hirose (1993) has proposed the active cord mechanism (ACM) as a general form of snake locomotion and/or manipulation, applicable also to aquatic motion and ciliated/flagellate motion. In the snake, the long thin flexible body is used to generate winding paths which distribute the snake's weight over its whole body. Snake movements are of four types as shown in Figure 55:

1. serpentine movement is the most efficient and is exhibited by almost all snakes;
2. rectilinear motion is used by large snakes such as boas and vipers in approaching prey;
3. concertina motion is used by snakes while traversing slippery surfaces but is of low efficiency;
4. side-winding motion is used by desert snakes such as the rattlesnake.

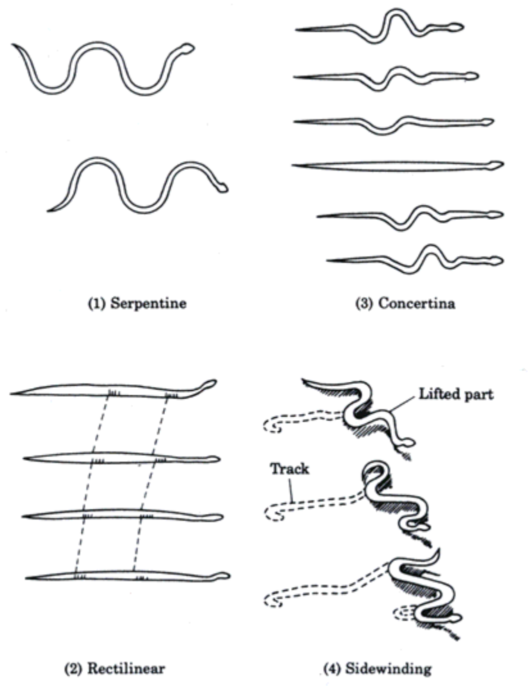

Figure 55 – Characteristic snake movements (Hirose 1993)

The snake body comprises some 200-400 vertebrae connected by joints controlled by agonist/antagonist muscle pairs. The range of movement of each joint is small ≤4°. The ACM may be analysed using serial link kinematics and Newton-Euler dynamics (as used in robotics). ACM actuators generate two rotational torques – one in the vertical plane and the other around the body axis - which intersect at right angles. The tangential joint force is given by:

$$F_t = \int_0^L \frac{dT(s)}{ds} \rho(s).ds$$

where:

$T(s)$ = joint torque
s = segment length between vertebral joints ~5mm typically
$\rho(s) = \frac{\pi\alpha}{2l}\sin\left(\frac{\pi s}{2l}\right)$ = curvature distribution
$\frac{dT(s)}{ds} = \frac{\sigma a}{l} f_o \left(\frac{s}{l}\right)^{\sigma-1}$
a = distance between joint to muscle lever arm
σ = force distribution >1
f_0 = maximum muscular force of muscle cross-section

Normal joint force, $F_t = \int_0^L \frac{d^2T(s)}{ds^2}.ds$

Sufficient propulsive force in the direction of the trunk must be generated without slippage occurring in the normal direction.

The computation of power follows readily:

$$P(s) = T(s)\frac{d\rho(s)}{ds}v$$

The serpentine shape coordinates are given by:

$$x(s) = sJ_0(\alpha) + \frac{4l}{\pi}\sum_{m=1} \frac{(-1)^m}{2m}.J_{2m}(\alpha).\sin\left(\frac{m\pi s}{l}\right)$$

$$y(s) = \frac{4l}{\pi}\sum_{m=1}^{\infty}(-1)^{m-1}\frac{J_{2m-1}(\alpha)}{2m-1}.\sin\left(\frac{2m-1}{2}\frac{\pi s}{l}\right)$$

where:

$J(\alpha)$ = Bessel function of the first kind
α = winding angle on body axis

The sinusoidal form is apparent. The α-adaptive principle states that $\alpha \geq \alpha_0$ where $\alpha_0 = \frac{\mu_t}{\mu_n}$ =friction coefficient ratio. Below α_0 slippage occurs but maximum efficiency demands $\alpha \approx \alpha_0$.

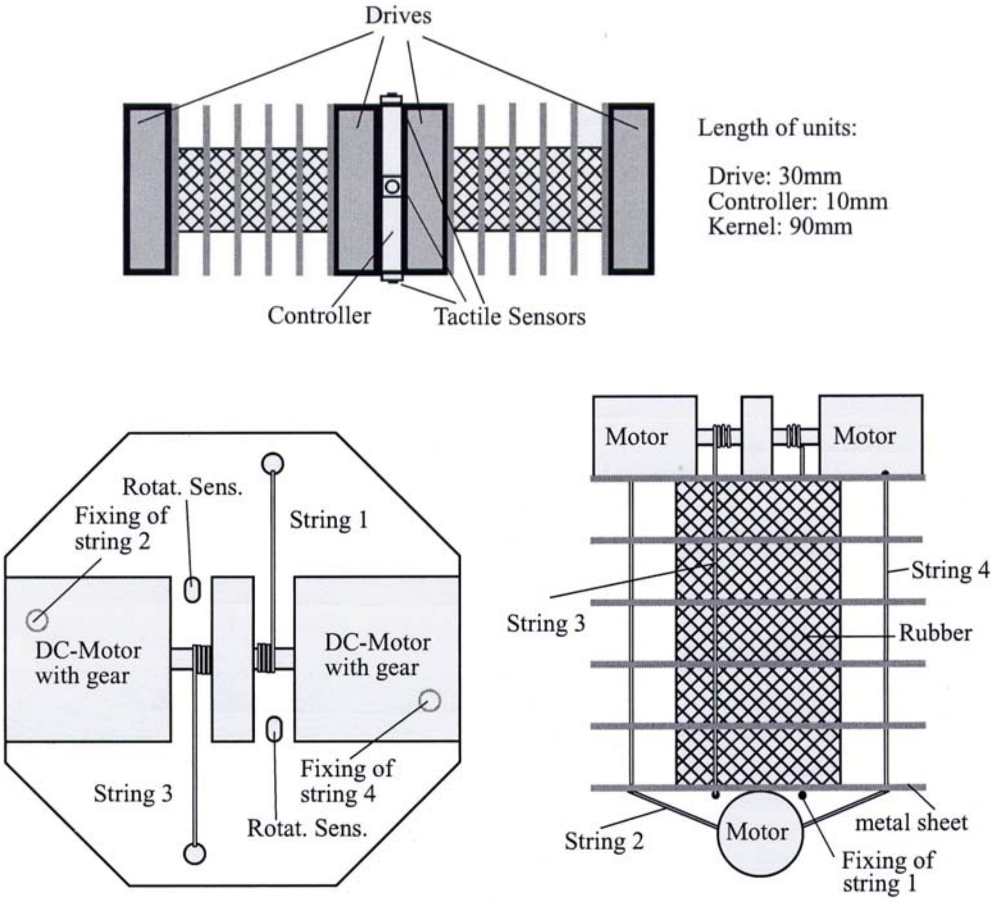

Figure 56 – a) Snakebot full joint section between links; b) snakebot joint top view; c) snakebot joint side view (Paap et al 1996)

A variation on the snakebot is the inchworm mechanism which may be used for ground-penetration, such as the Honeybee Robotics' Inchworm Deep Drilling System (IDDS) mole [Goravan 2000].

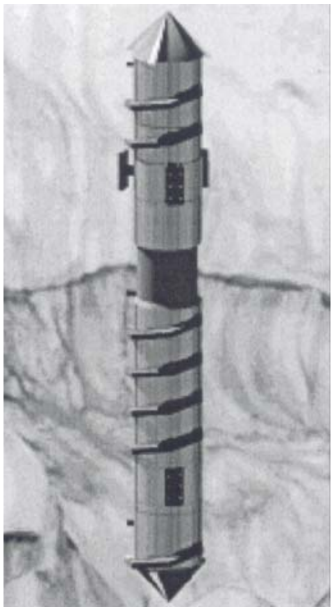

Figure 57 – Inchworm Deep Drilling System mole (Goraven 2000)

It comprises of symmetrical fore and aft sections, each fitted at its end with a motor-driven rotating drill bit pointing in opposite directions, and each with three braking shoes for anchorage to the walls of the hole. The two sections are connected by a telescopic joint to alter the length of the inchworm. A linear motor connects the fore and aft segments which contract and expand the inchworm. The mode of operation involves anchoring the aft section in a deployment tube with the aft braking shoes while the fore section drills to the maximum length of the telescopic tube. The telescopic tubing serves to ensure the integrity of the hole between the fore and aft sections. The fore section anchors itself to the drilled hole with its braking shoes. The aft section releases its brake and slides down the hole to meet the fore section as the telescope contracts. The aft section engages its braking shoes and the fore section releases its brakes and resumes drilling and the cycle repeats. To return to the surface, the procedure is reversed with the aft section taking the lead. Mead & Sedgewick (1993) have proposed a self-propelled mole concept actuated by a group of NiTi wires controlled by day/night temperature variations. A more practical variant would be to use electrical heating cycles to contract/relax the wires against a spring.

Similar in nature to snake-type rhythmic locomotion is swimming. Conventional submarines use screw propulsion with two vertical rudders below the propeller for yaw control, and hydroplanes for pitch control (and depth control). Swimming styles are varied – long fish such as eels undulate their entire bodies; flat fish undulate the long fins up and down; tuna undulate their rear sections left/right. The robotic lamprey swims by flexure of the body by flexion waves which increase in amplitude as they propagate from nose to tail. A central pattern generator generates the wave pattern actuated by segmented muscles constructed from Nitinol [Wilbur et al 2002]. The robotuna is streamlined and undulates the rear 40% of its body for propulsive power [Anderson & Kerrebrock 2002]. Robotuna uses cables as tendons driven by six electric motors to actuate its rear section. Nekton Research (Durham, NC) has developed fin-like propulsors from the "Twiddlefish" concept.

We consider that snake-like devices have limited utility in that they introduce difficulties in instrument packing and limited force exertion. Burrowing however has high utility in enabling access to sub-surface environments which are often of high scientific interest (eg. astrobiology). For the case of Mars, we can infer from models the change in soil parameters with depth.

The goal of Gavin Miller and his colleagues is to develop a biomimetic snake to be used during the aftermath of a major earthquake. This snake would have the ability to send audio and video back to the controller, as well as carry sensors to detect body heat, small explosives to clear rubble, and additional worthwhile payloads. Numerous robotic snakes have been constructed since Shigeo Hirose did the first one in the 1970s, and Miller plans to build on that previous knowledge for his design.

He considers the applicability of current and future snake robots to an imaginary search-and-rescue mission. Moving away from limbed creatures has proved to be an effective survival strategy. Snakes can hunt underground in confined tunnels, above ground in grassy fields, up in tree tops, even through extreme terrains. Through the design of biomimetic snake robots, the possibility exists of creating a useful vehicle to carry sensors and payloads over a variety of terrains.

Snakes have many locomotion gaits:

- Side-winding is popular on crumbly terrain such as hot sand. It may also serve the purpose of reducing contact area with the ground (to minimize heat transfer) while allowing a non-sliding means of locomotion on compliant surfaces.
- Concertina Progression is achieved by using a non-traveling curve along the length of the snake with modulation of the curve amplitude to attain forward motion.
- Horizontal Undulatory Progression is accomplished through a travelling wave along the length of the snake which produces directional friction with the ground (or through pushing against vertical obstacles).
- Rectilinear Motion is the most elaborate mechanism of the snake gaits, requiring the skin of the snake to be pulled forward and backward over the ribs as downward pressure is modulated on the scutes by a bulging and releasing of muscles. This is a slower movement, however, has the advantage of smooth and steady progress for the head over rough terrain.

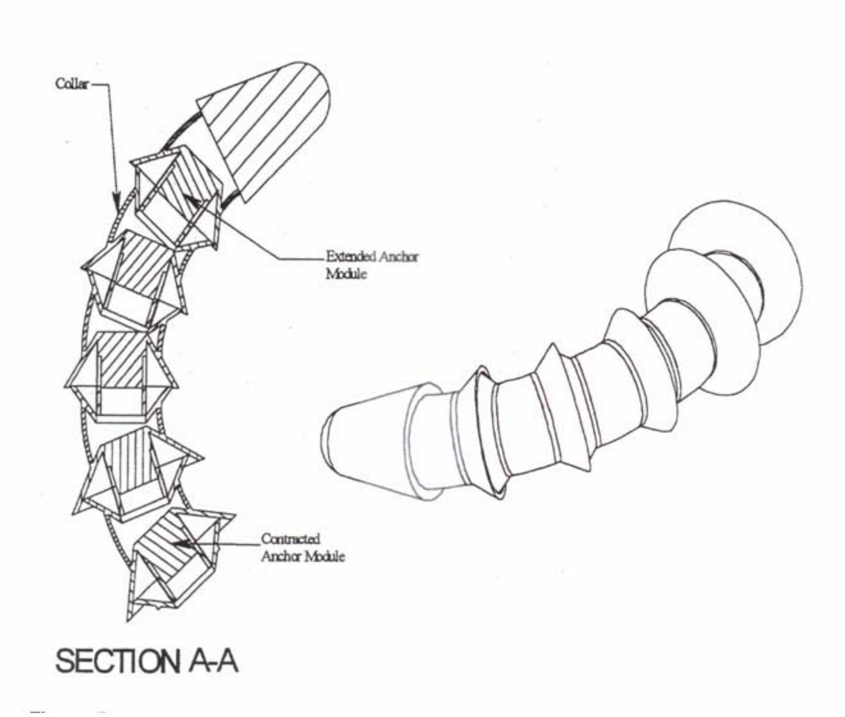

Figure 58 – Initial Burrowing Legless Robot Concept (Thakoor 1999)

An earthworm-like robot was proposed by Sarita Thakoor and the Caltech team at NASA's JPL. The peristalsis mechanism used for undulating movement of the robot is biomimetic to that of an earthworm. The foldable segments that become wider as they are shortened and longer as they are narrowed. As shown in the bottom portion of Figure 59, when the diamond-shaped open segments are contracted, they become longer and narrow, propelling the robot forward over the ground. The modification of the tip of the robot could be used to help burrow into the ground or under rocks. Small payloads could be added within the structure of the robot, or to the tip for sample collection. Each specific-purpose earthworm robot would be developed for its own specific task, as payload space is limited, however, the opportunities of a vehicle such as this are limitless. Designs to include a miniature active-pixel-sensor video camera, temperature sensors, soil or water detection payloads, or any number of scientific experiments would be possible. Also exploration robots to search through earthquake rubble or across uncharted interplanetary terrain would also be possible.

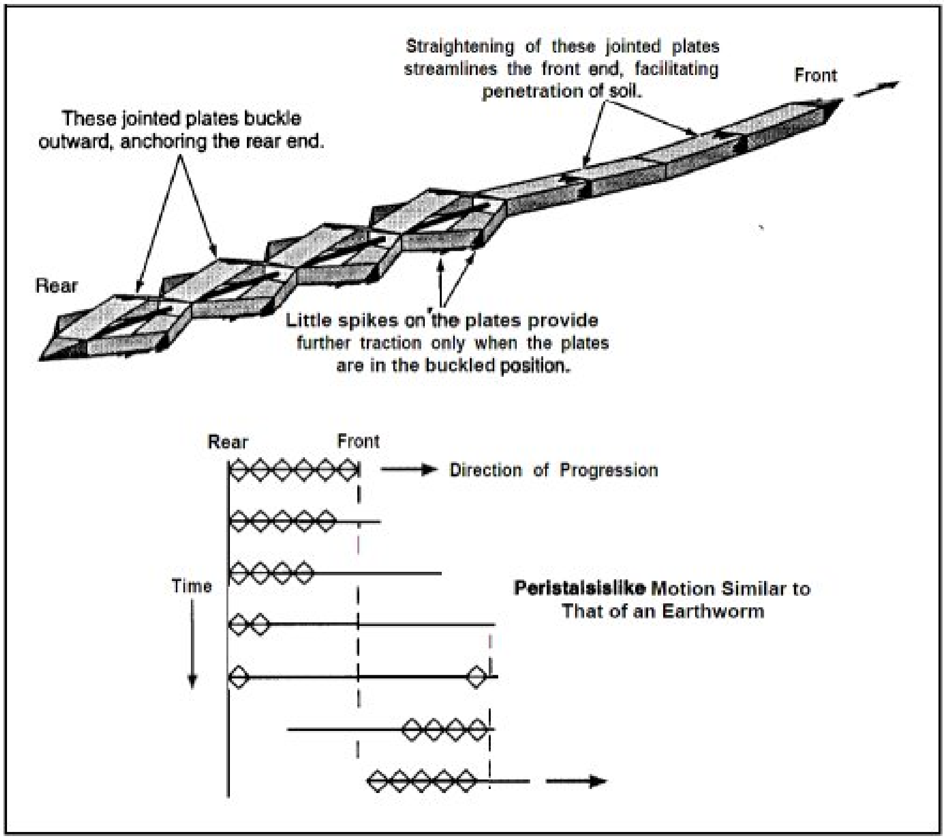

Figure 59 – Earthworm-like Biomorphic Robotics Explorer Concept (Thakoor 2002)

3.5.3. Winged Flight

Insects are a class of the phylum arthropoda (animals with jointed limbs) which also include crustaceans, myriapoda and arachnida. Insects, of which there are 29 orders, comprise 95% of all known animal species on Earth. Insect bodies are divided into head, thorax and abdomen; three pairs of legs are attached to the thorax, which can be considered to be the main locomotory unit. They have an external skeleton made of cuticle (qv). Their main structural material is chitin, a nitrogenous polysaccharide, which comprises ~20-50% of insect cuticle. Chains of chitin molecules are laid down in parallel layers each layer with a different orientation. The proteins of the cuticle matrix are covalently cross-linked forming a rigid three-dimensional structure suitable for the exoskeleton. The legs are structured so that the insect body is supported by the femoral-tibial joint of each leg. The three pairs of legs are located close to the head from the thorax for stable support. Non-flying insects appeared in the terrestrial fossil record around 400 million years ago. Silverfish are evolutionarily the most primitive insects, probably evolved from an annelid worm in which multiple segments fused into three body sections. All three pairs of legs emanate from the thorax suggesting that the thorax resulted from the fusion of three primitive annelid-type segments. Of the arthropods, only insects evolved wings and they were the first animals to evolve flight. Indeed, they are one of only four groups of animal which evolved flight (the others were pterosaurs, birds and bats which evolved independently much later). Indeed, it was the evolution of insect flight that led to the evolution of web-building arachnids. Flying is 10 times as energy intensive in time as ground locomotion (power intensive), but is 4 times more efficient in ground coverage (energy per km). Hence, the flight muscles of insects have the highest metabolic rate of any biological tissue, though they are thermodynamically still only 10% efficient. The elasticity of resilin provides the basis for mechanical energy storage and illustrates the utility of compliant structures in robustifying mechanical actuation. Insects breathe through spiracles which lead into branched tubes – the tracheal system.

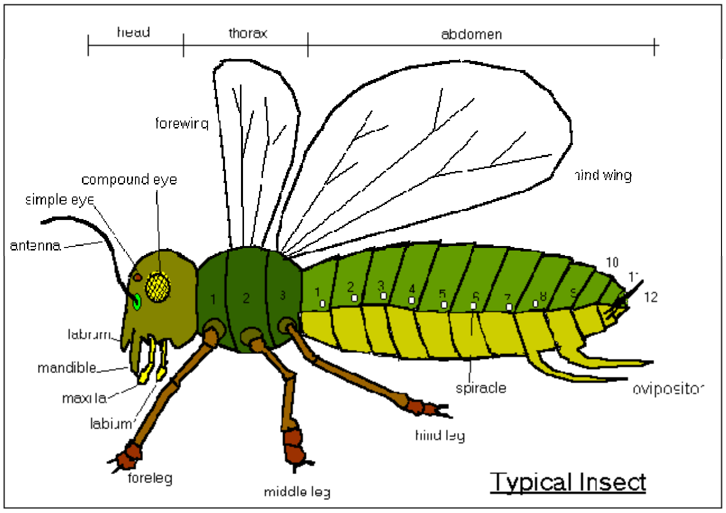

Figure 60 – General Insect BauPlan

Dragonflies and damselflies (Odonata) were among the first winged flying insects. They have two similar pairs of membranous wings. Each wing has five main veins and a nodus where the second vein intersects the wing leading edge. The two sets of wings are flapped 180^{o} out of phase. This ensures that a lift component is maintained through the wingbeat cycle. They have much diminished antennae but very large hemispherical compound eyes (allowing below and above vision) of 28000 lenses each (ommatidae). Their legs project forward on the body for grappling but dragonflies rarely walk as their legs aregeometrically wrongly disposed. They are strong fliers capable of forward speeds of 30-60 km/h, hovering and manoeuvring at high speed. They beat their wings at 20-30 Hz. Their wingspan varies from ~20 mm to 20 cm (average 5-8 cm in temperate climates up to 20 cm in the tropics). They are evolutionarily ancient, the oldest fossil of Protodonata dating from Carboniferous deposits ~325 million years ago. The largest dragonfly *Meganeura monyi* of the Carboniferous period had a wingspan of 75 cm, possibly due to the oxygen-enriched atmosphere ~35% O_2 generated by the carboniferous fern forests of the time. Cockroaches (Blattodea) evolved at around the same time ~360 million years ago as dragonflies and similarly reached gigantic proportions. Cockroaches have long spiny legs with flat broad bodies, usually with wings but they rarely fly. They are among the fastest runners of all insects. Beetles (Coleoptera) are the largest and most diverse order of insects but evolved somewhat later than dragonflies and cockroaches around 250 million years ago. They have hard forewings which protect the membranous hind wings. The forewings are held horizontally during flight aiding stability and the hind wings provide propulsion. Beetles however are not efficient flyers. Beetles are the most common and diverse animal on Earth which suggests the evolutionary success of this body plan which includes significant manipulative capabilities. Blowflies and mosquitoes (Diptera) possess two wings – the forewings – which are used for flight flapping at ~200 Hz – 1 kHz; the modified hindwings are gyroscopes. More evolved ("higher") Diptera are highly manoeuvrable capable of rapid changes in direction. Female mosquitoes and many other Diptera have a piercing proboscis which they use for feeding on the blood of birds or mammals. Bees and wasps (Hymenoptera) have two pairs of membranous wings, the forewings often longer than the hindwings which are held together along their common edge by tiny hooks so that the dual wings act as a single surface. Bees and wasps, like butterflies, evolved around 25 million years ago probably stimulated by the appearance of flowering plants. Ants (Formicidae) are characterised anatomically by a narrow front part to the abdomen – only reproductives are winged. They are, like many bees and wasps, social insects and live in colonies, and are, like bees and wasps, believed to have evolved from sawflies (Hyenoptera) relatively recently around 25 million years ago. Butterflies and moths (Lepidoptera) evolved subsequent to the appearance of flowering plants and possess wings bearing sheddable scales which can act as a defence against arachnid webs. They have a proboscis which coils under the head for extracting nectar from flowers. Grasshoppers (Orthoptera) have highly developed hind legs for jumping – the forewing is a protective cover for the membraneous hind wing used for flying. The praying mantis (Mantodea) have strong spiny forelegs which are used for grasping prey. They usually have hardened forwings and membranous hindwings and rarely walk rather lying in wait for prey.

Many birds and flying insects are unstable flyers which require frequent manoeuvring to remain aloft. There are two types of vertebrate flight: gliding/soaring (such as the "flying lemur"), and flapping flight (such as the bird and the bat). Unlike bats and birds, insect wings are not permeated by muscle but act as levers. Vertebrate wings however have considerable *controllable* flexibility to act as control surfaces. Bat wings (Chiroptera) have greater flexibility of their wing membranes than insect wings, the conformation of which is controlled by muscles within the wing. The thin elastic wing membrane of the bat sandwiches muscle and blood vessels, and is supported by the limb, body and tail. Bats can alter their wingspan by 20%. The bat has four downstroke muscles (compared to one – the pectoralis major – in the bird) and five upstroke muscles (compared to one – the supracoracoideus – in the bird) attached to the scapula. Birds (Aves) have feathers rather than membranes which provide a greater range of geometries. Birds possess 48 muscles within the wing itself to control the wing surface by the orientation of the feathers. This use of internal wing muscles of both the bat and the bird allows greater range of movement than that available to the insect. Both bats and birds exhibit similar wingbeat patterns. In both cases, the leading edge is tilted downwards during the downstroke and tilted upwards during the upstroke. The wing is also extended during the downstroke but in many, it flexes inwards towards the body during the upstroke. Some large birds glide to exploit thermal updraughts to minimise energy consumption [Pennycuick 1996]. The L/D ratio for bats is similar to that of pigeons at 6.8.

Birds possess feathers which maintain a thin layer of warm air around the bird to maintain its temperature. Bird and bat flight is considerably more complex than insect flight – it is suggested therefore that insect flight represents a more tractable solution to the problem of aerial planetary rovers. Generally, with insect locomotion, we have two choices: first, we choose a general purpose bodyplan for general deployment, e.g. the beetle; secondly, we choose the optimal solutions to each locomotion capability. We suggest selection of the last option whereby the optimal flight solution is combined with the optimal ambulatory solution. We further suggest that both may be controlled using similar approaches, i.e. central pattern generators for rhythm production calibrated with sensory feedback and augmented by structural compliance.

Flapping winged flight is the means of locomotion by some insects and birds – wing flapping generates lift whereas engineered aircraft generate lift by the aerofoil shape [Hedenstrom 2002]. Fixed wing aircraft utilise three control surfaces for manoeuvring – elevators (pitch), rudder (yaw) and ailerons (roll); rotary wing aircraft use the pitch of the rotating blades to generate lift while propulsion is enabled by tilting the rotor disc forward (with tail rotor pitch providing yaw control). Fluid inertia provides pressure force required for lift which increases with fluid density. Shear forces generated by fluid viscosity which relates to a particular size scale of the body. This is quantified through the Reynolds number which is the ratio of the inertial forces to viscous forces:

$$\mathrm{Re} = \frac{l v \rho}{\nu}$$

where:

l=characteristic length of flyer
v=velocity of flyer
ρ=fluid medium density
ν=fluid viscosity $\cong$ 100-1000 for insects

A Reynolds number exceeding 4000 indicates turbulent flow. When designing a flying machine, the wings must support the weight of the vehicle such that lift exceeds weight. Wing loading, $\frac{W}{A} = 0.3\rho V^2$ where ρ is the air density (1.25 kg/cm^2 at terrestrial sea level), A is the wing area, W is the weight, and V is the flight velocity. The typical aspect ratio $R = \frac{b^2}{A}$ for a bird with low wing-loading is 7-10, i.e. long slender wings. For hummingbirds and insects, $W = 0.25\rho v^2 b^2$ where v is the downward wing speed (0.45V) and b is the wingspan. For hummingbirds and insects, R~6, which provides a cruising speed of 7 m/s and wingbeat speed of 3 m/s. Some large birds glide to exploit thermal updraughts to minimise energy consumption [Pennycuick 1996].

Table 11 – Flying Species with Weight and Wing Dimensions

Species	**W (N)**	**A (m2)**	**b (m)**
Crane fly	3×10^{-4}	7.5×10^{-5}	0.02
Dragonfly	0.01	0.0018	0.1
Swallowtail	0.006	0.003	0.08
Cockschafer	0.01	0.004	0.06
Marsh harrier	6.8	0.22	1.35

Most insects possess two pairs of wings attached to muscles hinged with resilin to the thorax which flex the wing to manipulate the leading and trailing edges. Insects employ vortex generation to gain lift and are characterised by high wingbeat frequency for a high coefficient of lift. Insect wings are thin with sharp leading edges which generate leading edge vortices which generate most of the lift. The leading edge is stiffer but the trailing wing surface is flexible. Muscles are connected to the exoskeleton of the thorax. Insect wings are constructed from thin membranes often supported with blood-filled veins. They flap at a rate of 20-1000 Hz such that the aerodynamic flow is dynamic which dampens out stable lift exploited by aircraft. Flapping frequency and wing orientation provide the basis of fine control in flying insects - the blowfly has a wingbeat frequency of 150 Hz. Insect wings' leading edge tilts down during the downstroke and upwards during the upstroke. The wing is thus rotated at the end of each stroke. Due to this rotation, the tip of each wing traces out a thin oval tilted at a steep angle. The wing rotation makes the topside face up during the downstroke and the wing underside face up during the upstroke. Insect wings are thin with sharp leading edges which generate leading edge vortices to aid the lift from the wings. Insects employ vortex generation to gain lift and are characterised by high wingbeat frequency for a high coefficient of lift. The delayed stall causes a leading edge vortex of high flow velocity to form just above and behind the leading edge on the downstroke which initially increases lift and spirals to the wing tip and shed before the upstroke. The wings flap so rapidly that the wing flips over and reverses direction to produce a new vortex in the opposite direction once the previous vortex is shed. Furthermore, axial air flow forms from the base to the tip of the wing to keep the vortex attached to the wing throughout the downstroke. Backward rotation of the leading edge of the wing at the end of each stroke produces rotational lift by pulling air faster over the top of the wing. In addition, each wing stroke creates a complex wake behind the wing. When the wing rapidly reverses direction at the start of each stroke, it passes through the wake. By adjusting the timing of rotation, insects can manoeuvre by adjusting the lift on each wing. Earlier rotation of one wing generates more lift on the outside of the turn while later rotation on the other wing generates less lift on the inside of the turn. Such turns are controlled by feedback from optical flow detected by the compound eyes, gyroscopic motion of the hind wings and mechanosensors on the wing to alter the rotational turning and stroke amplitude. The thorax aids wingbeats by elastically deforming such that the wing is stable only in the extreme positions and are dynamically unstable in mod-position – this is enabled by the distortion of the thorax and the hinge articulation. Wing movement is achieved by rhythmic thorax deformation using dorsolongitudinal muscle contraction for the downward sweep and dorsoventral muscle contraction for the upward sweep. Insect landing is accomplished by transforming the leading edge vortex into a wingtip vortex which may be shed without stalling.

Insect landing is accomplished by transforming the leading edge vortex into a wingtip vortex which may be shed without stalling. Robofly based on the blowfly comprises of a two-winged system driven by three stepper motors [Fearing et al 2002]. The Entomopter project has the goal of developing a fully functional autonomous insect capable of crawling, swimming and flying. The Entomopter flight system uses a resonant wingbeat with one degree of freedom (and so no tilt control) and flap 180^{o} out of phase [Michelson 2002]. The wing deformation on the upstroke provide some lift by controlling the angle of attack and camber. The fuselage is constructed from shape memory alloy with torsional memory to store mechanical energy, once again illustrating the importance of body compliance in animal locomotion. It uses a chemical muscle that generates energy by chemical reactions to drive the wingbeats [Michelson 2002]. In terms of control, wing movements are treated in a similar fashion to legged locomotion, typically involving the use of central pattern generators. We similarly suggest that preflexes imparted by structural compliance offer the potential for robust flight.

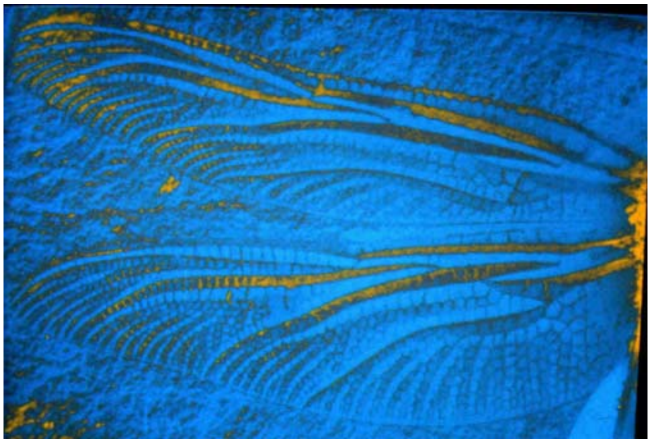

Figure 61 – Dragonfly Wing

3.5.3.1. Muscles and Wing Beats

In most insect flight, the main power muscles rhythmically deform the thorax of an insect, which, in turn, causes a downward rotation of the wings, as shown in Figure 62. These power muscles are innervated only by a single motoneuron. This type of flying is common in two main functional groups of flying insects: neurogenic and myogenic fliers, which are known to have wing-beat frequency ranges between 15-35 Hz and 30-1000 Hz, respectively.

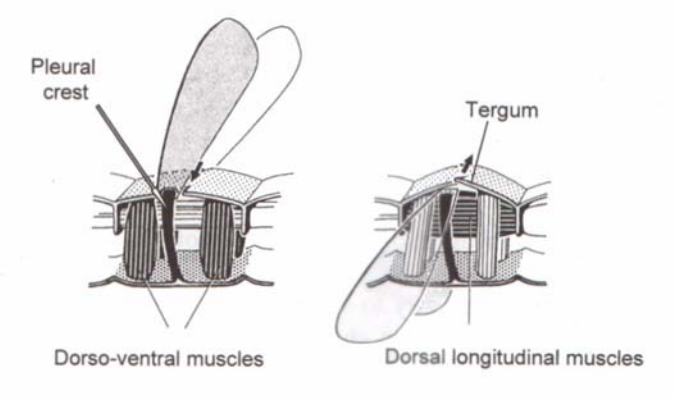

Figure 62 – Muscles used in Insect Flight.

In neurogenic fliers, the wing beat is under full neuronal control, meaning that a spike in a motoneuron of the main power muscles always produces a full contraction, as shown on the top of Figure 63. Interneurons produce a basic rhythmic motor pattern, all of which are spiking and no non-spiking interneurons have been found. These interneurons activate motoneurons, which in turn activate muscles and produce the wing movements. Receptors monitor the movements of each of the wings, including a stretch receptor to detect disturbances in wing speed timing and wind hairs that monitor the rhythmic air stream produced by the wings. These receptors are part of the flight motor and provide very rapid response times, while other feedback systems are not part off the motor itself.

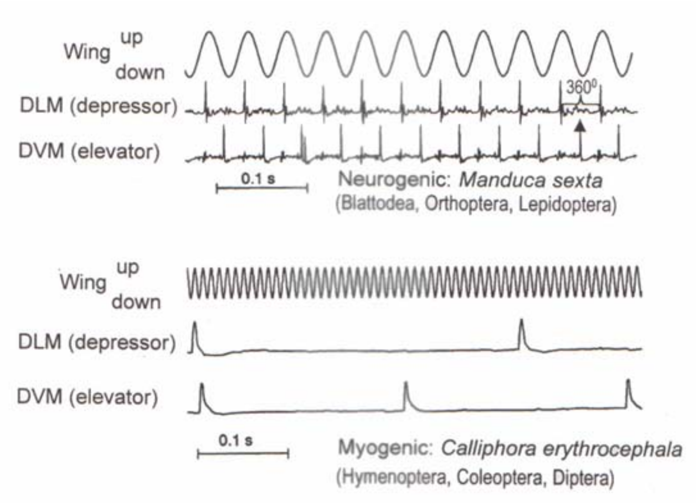

Figure 63 – Wing beats for Neurogenic and Myogenic Insects, respectively

Myogenic fliers use stretch activation, a property of their power muscles, which allows them to respond with an active contraction to a rapid mechanical stretch. When the dorsolongitudinal muscles (DLMs) contract and bend the tergum upward, the dorsoventral muscles (DVMs) will be stretched by the distortion of the thorax. This stretch causes a single contraction of the DVMs, which will stretch the DLMs, whereupon they will contract and so on. This system operates in resonance and the motor neurons fire at rates much less than the wing-beat frequency, as shown on the bottom of Figure 63. Additionally, steering is controlled by separate muscles, activated by motoneurons continuously or only during the flight manoeuvre in wing beat frequency. Though these differences exist for the steering system, much of the feedback system design is similar to that of the neurogenic fliers.

Wendler specifically studied the Hawkmoth *Manduca sexta*, a 4-winged myogenic flier whose hind wing muscles are small and weak, thereby simplifying the analysis to functionally a two-winged insect. The activity of the insect was monitored on three levels. First, motor output were monitored through 12 flight muscles. Analysis shows that in straight flight, the elevator muscles are activated approximately in anti-phase to the depressor power muscles, while in a turn the inside elevators are activated slightly sooner, indicating a reduction of thrust on that side. Second, three forewing movements were measured: flapping, turning, and tilting. Third, the aerodynamic forces and moments are measured with a 6-component flight balance, where the insect is tethered to a floating system that allows only very small movements (0.5mm and 0.2^{o}). The resultant forces and moments are measured by linear photodiodes by the reflection of 3 laser sources reflecting off of tiny mirrors mounted below the insect.

Observations have shown asymmetrical positioning of the wings, showing that the wings on the inner side are shifted far backward while the other wings are drawn forward. Also, sometimes forewing and hindwing on one side are mechanically decoupled, mainly based on shifts in the forewing. Drastic movements in the head and postural changes of the abdomen are also seen. Correlations can be made between motor pattern and wing movement patterns, however, only between complex patterns as a whole. It is not possible to draw any conclusions from the amount of phase shifts in a single muscle.

Further experiments were performed by removing the ventral nerve cord from the insect, while leaving the thorax and wings intact. The main power muscles were then stimulated by a pattern of computer-controlled electrical pulses. Depression and contraction times were determined, as well as frequencies at which the muscle stretches. It was noted that the experiments were carried out at room temperature of 23^{o}C, but they expect that temperatures closer to the insect's body temperature (36^{o}C) will shift the resonance frequencies to match the wing beat frequency of the intact animal. It was deduced that the insect probably flies at or near the resonance frequency of the thorax, which would allow the most efficient use of energy.

These results raise new questions. For example, how is the frequency match between the neuronal and mechanical parts of the system achieved in a neurogenic flier, since the wing beat is fully controlled by the nervous system? Additional further work includes mounting the computer controlled insect body to the 6-component flight balance in order to manipulate the motor system to obtain desired flight manoeuvres.

3.5.3.2. RoboFly

Ronald Fearing (ref.) and his team from University of California, Berkeley are working to apply flapping flight for a microrobot, which they aptly call RoboFly. A conceptual drawing of their micromechanical flying insect (MFI) can be seen in Figure 64. It is planned to have a 25 mm wingspan, be driven by piezoelectric actuators, and will have a flexural four-bar element to provide sufficient wing stroke. It will be modelled after the blowfly *Calliphora*, which has a mass of 100 mg, wing length of 11 mm, wing beat frequency of 150 Hz, and actuator power of about 8 mW. RoboFly is a two-winged system driven by three stepper motors and closely mimic the stroke kinematics of a fruit fly. Strain gauges measure instantaneous wing forces.

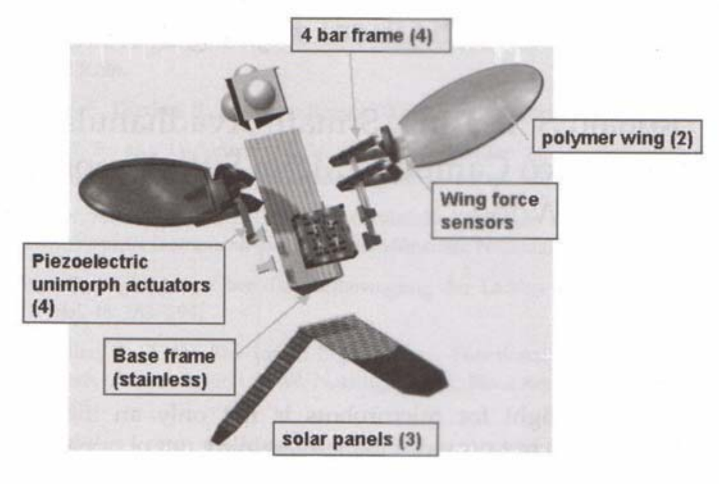

Figure 64 – Conceptual Design of RoboFly Microrobot Flier

Running RoboFly with a wingbeat of 1/6 Hz in oil provides the same Reynolds number as a fruit fly with a wing beat of 220 Hz in air. Results of this test identified three key aerodynamic mechanisms used by insects: delayed stall, wake capture, and rotational circulation.

Insect flight at the centimetre scale requires both large stroke amplitude and wing rotation. The *Drosophila* has a wing stroke of 160^o combined with a wing rotation of more than 90^o, which creates a complicated arrangement of linkages and cams. The RoboFly thorax uses separate four-bar frames to control the leading and training edges of a wing.

For testing purposes, a 5x scale thorax was constructed. In a wind tunnel, the following Reynolds calculation was performed:

$$\mathrm{Re} = Sf\phi / v$$

where:

$f = 17$ Hz
wind tunnel flow = 0.14 m/sec
$v = 1.6 \times 10^{-4}$ m^2/sec (kinematic viscosity of the air)
$S = 32$ cm^2 (wing surface area)
$\phi = 1$ rad (stroke amplitude)

An overall Reynolds number of 340 was determined, which is comparable to the blowfly. Upstroke and downstroke induced velocities were as high as 1.65 m/sec and 0.79 m/sec, respectively, with a peak wing-tip speed of 2 m/sec. Two interesting results are 1) significant momentum change in the particles directly under the wing from the end of the upstroke to the beginning of the downstroke indicate a period of large lift forces and 2) vortex shedding at the trailing edge of the wing is readily apparent as the wing finished the downstroke and begins the upstroke. Also, 8 strain gauges were added to measure actual wing-force measurements, each requiring 2 wires for data transfer. Although wire ribbon was found to be minimally intrusive to the aerodynamics of the 5x scale RoboFly, much smaller and more flexible wire will be required for the original scale vehicle.

The full scale RoboFly prototype was constructed from cuticle with resilin connections, which compares very well in the stress limits with stainless steel. Peak stress must be minimized to maintain system life, which can be done with thin materials and by keeping the overall system Q low (measure of stored energy to energy dissipated per radian). Strength-to-weight ratio must be high n the four-bar frames, so hollow rectangular and triangular beams. Cross-sectional moments of inertial solid beams with width *w* and height *h* are:

$$I_{rect} = wh^3 / 12$$

$$I_{tri} = wh^3 / 36$$

The stiffness of a cantilever beam of length *l* and modulus *E* is given by

$$k = 3EI / l^3$$

A hollow beam has a cross-sectional moment of inertia given by:

$$I_{hollow\_beam} = I_{outer} - I_{inner}$$

So, creating a hollow triangular beam from folding a 12 micron thick sheet provides an I_{tri} of 700x10^{-18} m^4. Additionally, instead of pint joints, which are difficult to place on a 1 mm structure without numerous problems, flexural joints were used, which were created by folding photo-etched stainless steel. Although the prototype used manually folded supports, automatic folding will be used in future builds to increase reliability and decrease construction time.

The next hurdle in the prototype thorax design for the MFI is power distribution. Insects deliver on the order of 100 W/kg to the wings. To avoid significant losses with accelerating inertial loads, insects drive their wings at mechanical resonance. Though, even when using the benefits of mechanical resonance, wing damping occurs that reduces the wing amplitude on each cycle. This damping needs to be overcome through the addition of energy to the system.

Wing damping B_w at the wing hinge can be estimated as:

$$B_w = 2m_{MFI} g l_w / (\omega \theta_w)$$

where:

$m_{MFI} = 0.1g$ *(MFI mass)*
$l_w = 10$ *mm (wing length)*
$\omega = 150$ *Hz (wing beat frequency)*
$t_w = 35^o$ *(wing stroke amplitude {1/2 total wing stroke})*

Therefore, $B_w = 13$ nN-s-m.

To have a low-Q manoeuvrable wing, the wing inertia is:

$$J_w = QB_w / \omega = 3.3 \text{ x } 10^{-11} \text{ kg-m}^2$$

The actuator stiffness at the wing hinge is:

$$K_w = J_w \omega^2 = 30 \mu Nm$$

The actuator stiffness with the wing inertia is:

$$K_a = T^2 K_w \qquad \textit{(where T is the given transmission ratio)}$$

At DC, the sidlacement of the wing is just

$$\theta_{wDC} = T \tau_a / K_a$$

At resonance, the reactance of J_w and K_a/T^2 cancel, so the magnitude of the wing displacement is:

$$\theta_{wAC} = \tau_a / (TB_w \omega)$$

The AC motion amplification at resonance is:

$$\theta_{wAC} = \theta_{wDC} = Q$$

Reducing from above, the required DC actuator displacement is just:

$$\theta_a = \theta_w / (QT)$$

So, with a given transmission ratio *T*, the required actuator displacement θ_a and torque $\tau_a = K_a \theta_a$ can now be used to specify actuator dimensions. Finally, for a unimorph with elastic layer stiffness equal to the piezo stiffness, the required length and width can be found in terms of the actuator thickness *h* as follows:

$$l = 4h\theta_a / \left(3d_{31}E_3\right)$$

and

$$w = \left\{16\tau_a / \left(3d_{31}E_3Y_{11}h^2\right)\right\}$$

where E_3 is the electric field and Y_{11} is the elastic modulus.

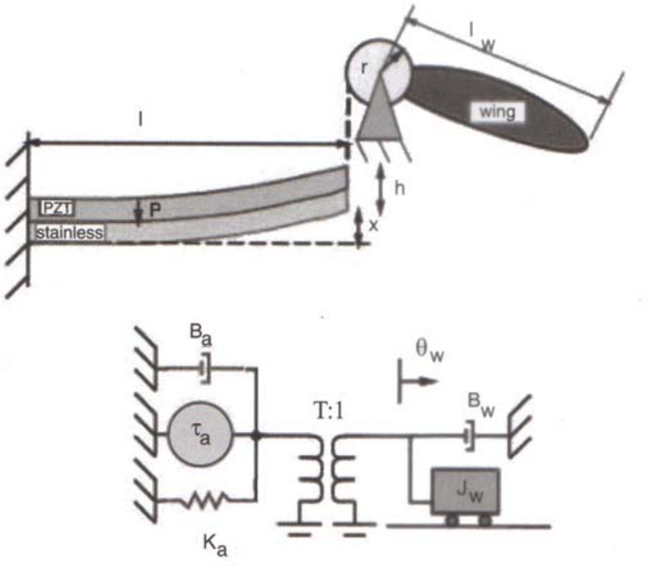

Figure 65 – RoboFly Diagram

With respect to power distribution, Fearing and his team chose high-voltage solar cells (about 100 V) so the piezoelectric actuators can be activated directly without voltage conversion. Piezoelectric actuators have intrinsic electrical and mechanical losses, and transistors are to be used to reduce switching losses. After performing a resonance test and finding a thorax dissipation of 1.2 mW at 150 Hz and 150^o stroke, an estimated overall loss of 2 mW is expected, which compares favourably to gear losses for conventional motors.

Up to this point, they have achieved a good understanding of the necessary wing kinematics, forces, velocities, and power requirements and applications based ton measurements on RoboFly and real insects. Upcoming research includes the development of a closed-loop wing controller that can react to wing forces and modify wing stroke patterns as needed to achieve stable flight.

3.5.3.3. Entomopter

Robert Michelson has developed an Entomopter, a multimode autonomous robot, capable of flight, ambulatory locomotion, and swimming behaviours in a single vehicle (ref).

The Hawkmoth (*Manduca sexta*) was the biomimetic baseline for the winged abilities of the Entomopter. However, the flapping mechanism for the Entomopter extends beyond that found in the hawkmoth, because the flapping structure is a single-piece construction that takes advantage of torsional resonance in the fuselage to recover flapping energy. Though the technique of energy recovery is common in some flying insects, it is applied very differently than what is found in nature.

The reciprocating chemical muscle (RCM) is the heart of the Entomopter system. The RCM is an anaerobic, ignitionless, catalytic device that can operate from a number of chemical fuel sources. It converts chemical energy into motion through a direct non-combustive chemical reaction. Additionally, small amounts of electrical energy are created, which are used for controlling the MEMS devices and nervous system of the Entomopter. Although the RCM creates a constant frequency and amplitude on the wings, directional control is produced through varying the coefficient of lift on each wing. Figure 66 shows a third generation prototype of the RCM that was able to reciprocate at 70 Hz with throws of 0.5 inches.

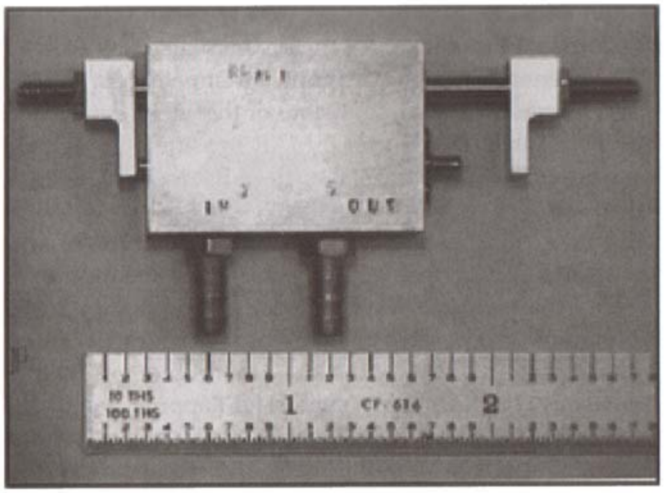

Figure 66 – Third Generation RCM at 2x Scale

The X-wing design employed within the Entomopter provides simplified manufacture process at the micro air vehicle (MAR) scale (under 15 cm in any dimension as defined by DARPA), because the flapping wings and body can be produced as a single torsional part. Another benefit of this design is the additional longitudinal stability along the pitch axis.

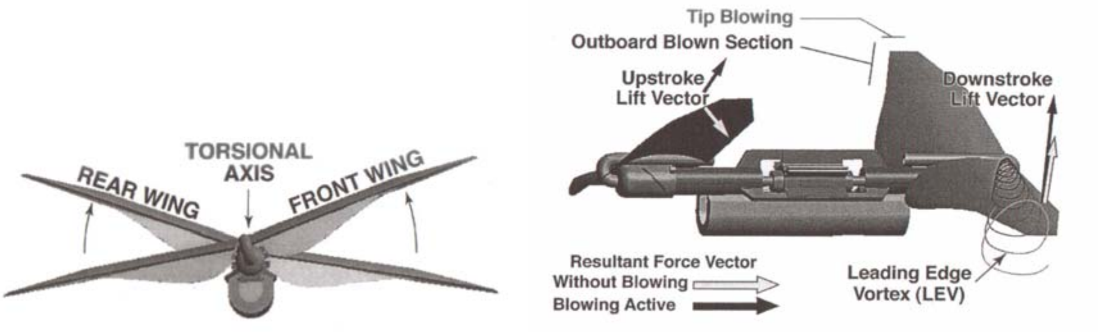

Figure 67 – a) Front view of Entomopter showing the X-wing flapping design; b) Side view of Entomopter showing flight vectors with and without blowing activated

The leading-edge vortex, as shown in Figure 67b augments the circulation and causes lift on the wing. Additionally, waste gas from the RCM can be vented through the wing tips to achieve a level of circulation control and keep the leading-edge vortex attached to the surface longer than naturally possible. This creates minor changes in the lift capabilities of an individual wing, which would provide in-flight turning, eliminating the need to beat a wing faster (disturbing the resonant frequency) or modifying the angle of attack. When the wings are operated 180° out of phase (as shown in Figure 67a) there will always be two rising and two descending wings. The wings were stiffened by materials that react differently to opposite loads, therefore, they can provide lift on both the downstroke and the upstroke. Also, the fuselage is made from a material with torsional memory, which acts as a torsional spring with a resonant frequency when the wings are flapped. This also simplifies the structural manufacturing and lifetime because it can be created as a single part, without the need for hinges or gearing between the fuselage and the wings. The construction of this totally resonant structure is biomimetic as well, as all insects have resonant structures that store potential energy, either in the muscles themselves or in their exoskeletal parts.

Crawling locomotion can be used by the Entomopter when flight is not possible, or to navigate under doors, or especially in a case where the wings are damaged. The RCM is still used for crawling; a set of four legs can be integrated so that they reciprocate in a plane 90° to that of the wing motion (though disabled during flight). The paddle-like legs are developed with as few moving parts as possible to reduce the mechanical complexity of the system. The RCM is lowered from tens of hertz when flying to 1 hertz or less when walking.

Complete autonomous navigation is required for the Entomopter, as the communication hardware required for real-time navigation would adversely affect the weight and control of the system. Of highest importance is short-range obstacle avoidance, which would include a small onboard ranging system. A single ultrasonic source can be vectored by the wing-flapping mechanism of the Entomopter to illuminate both the floor for altimetry and the walls on the sides to detect impending collisions. Simple avoidance reflexes could be programmed with a minimum of onboard processing to have the Entomopter move away from a potential collision.

The Entomopter is a unique aerial robot with biomimetic origins, however it goes well beyond biomimicry to incorporate novel control techniques. The reciprocating chemical muscle can be used to support many forms of locomotion within Entomopter, including flight, walking, and even swimming. The design is ideal for indoor use where wind impacts are minimal, though the design also has promising use on Mars.

3.5.4. Swimming

3.5.4.1. *Dolphin Robot*

The swimming speeds of tuna and dolphins are greater than 30 km/hr and the swimming motion of these animals is thought to be highly efficient. Few studies have examined experimentally the consumed power and propulsive efficiency of a large self-propelled swimming robot. Two simplified models of a dolphin were built with the following features:

1. Designed to study forward swimming at constant speed
2. Self-propelled and self-contained
3. Have two joints (minimum, to actuate the caudal fin). The first joint is actuated by an active motor and the second joint is moved passively by springs
4. No buoyancy control (swims near water surface)

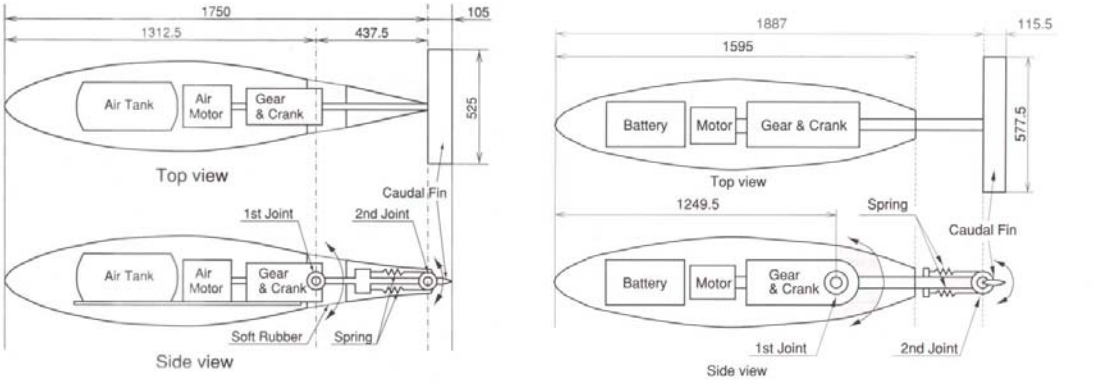

Figure 68 – Dolphin 1 and Dolphin 2 Schematics

The caudal fin was constructed according to NACA 65$_3$-018 airfoil. The body shape was based on the NACA 65$_4$-021 and consists of hexagonal aluminium frame and a urethane foam skin to provide a streamlined shape. The posterior part of the body was not circular but elliptical in order to avoid unnecessary drag that occurs in reciprocal motion.

Power consumed by the fin during swimming is solved by:

$$P = T \times \frac{d\theta}{dt}$$

where:

θ = *joint angle*
T = *torque with respect to time*

Propulsive efficiency is defined as:

$$\eta = \frac{DU}{\overline{P}}$$

where:

D = *towing drag*
U = *propulsive speed*
$\overline{P}$ = *time average of P*

Reduced frequency is defined as:

$$\sigma = \frac{\omega c}{U}$$

where:

c = chord length of the caudal fin
ω = spring stiffness

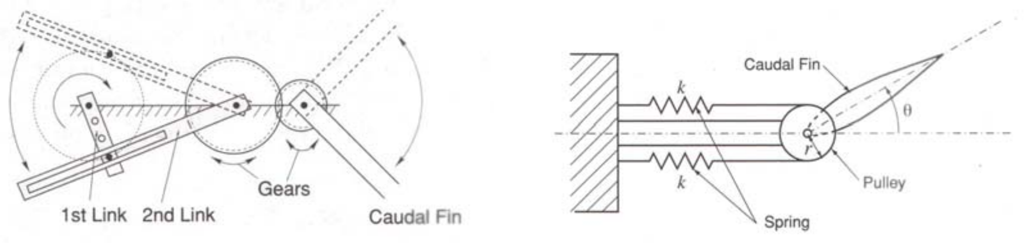

Figure 69 – First Joint and Second Joint of Dolphin 1

A stiffer spring produces lower propulsive efficiency for a constant reduced frequency. The maximum propulsive efficiency reaches around 0.7 and does not vary much with respect to the spring constant except for the case of a fixed second joint. The maximum propulsive speed of the first robot was experimentally found to be 1.15 m/sec

The second robot has a slightly different mechanical construction in the wing assembly than the first robot. This can be seen in Figures 16.17 and 16.19.

A propulsive efficiency of nearly 0.65 was found when a weaker spring was used and the propulsive speed was found to be 1.9 m/sec.

3.5.5. <u>Manipulation</u>

Manipulation by arms is commonly used in robotics but bear little resemblance to their natural analogues – hands, claws, tentacles, etc. Traditional revolute robot manipulators comprise three joints (waist, shoulder and elbow) for three degree-of-freedom positioning and a three degree-of-freedom wrist (roll, pitch and yaw) for orientation. Although manipulators are modelled on the human arm geometry, the human arm has more than six degrees-of-freedom, i.e. it is redundant which provides an infinite number of poses for any given end-effector position. Most robotic end-effectors are parallel jaw grippers though three fingered hands with multiple joints (typically three per digit) have been built, eg. MIT and Utah robotic hands. The human hand is a single 22-degree-of-freedom tool of unprecedented multi-tasking capability by virtue of different grips – it can form a hook, a scissor grip, multi-fingered chuck, squeeze grip and multiple geometric grips allowed through the mouldability of the palm. Prehensile capacity allows the holding of an object in a controlled state relative to the hand. This requires the application of sufficient force on the object to hold it stationary. It imposes constraints by contact and imposes structural and frictional constraints. Although traditional robotics has concentrated on three-fingered hands [Salisbury & Craig 1982], recent fossil evidence suggests that biological evolution converged on five digits as the optimal number for general purpose articulation. The most sophisticated robotic hand currently under development is the Robonaut hand for performing EVA-equivalent functions. An alternative model for grippers is through the use of soft, compliant, universal gripping mechanisms as shown in Figure 70:

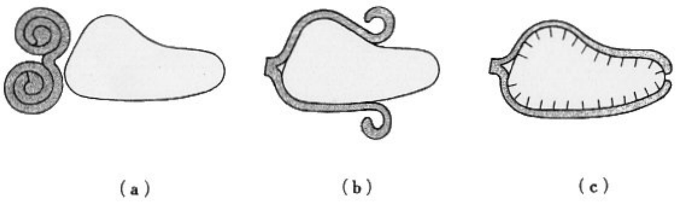

Figure 70 – Soft gripper with flexible grasping mechanism (Hirose 1993)

Robust manipulation requires complete haptic feedback for handling compliant and slippery material, including force feedback for pressure information and tactile imaging for the extraction of shape. Capacitative tactile sensing arrays offer promise for fine motion tactile handling. The greater the dexterity of the robotic gripper, the greater the number and diversity of tasks that can be performed. The more sides that a polygonal shape possesses (i.e. the closer approximation to circular geometry), the greater the appropriateness for a three fingered grip rather than two-fingered. Two fingered grips must have the fingers arranged parallel to each other, while three fingers provide greater stability, typically arranged 120^o from each other. The precision (prehensile) grip involves contact with the object by the fingertips while the power grip uses the fingers and palm to clamp the object - the parallel jaw gripper cannot perform a power grip. More complex grippers involve a minimum of three fingers, each with three or four degrees of freedom (the fourth degree of freedom emulates that of the human palm) to emulate the diversity of grips exhibited by the human hand. The DLR dextrous Hand II which is capable of precision and power grasping - power grasping requires parallel alignment of second, third and fourth fingers while precision grasp requires opposition of first finger and thumb with intersection of other fingers [Hirzinger et al 2004].

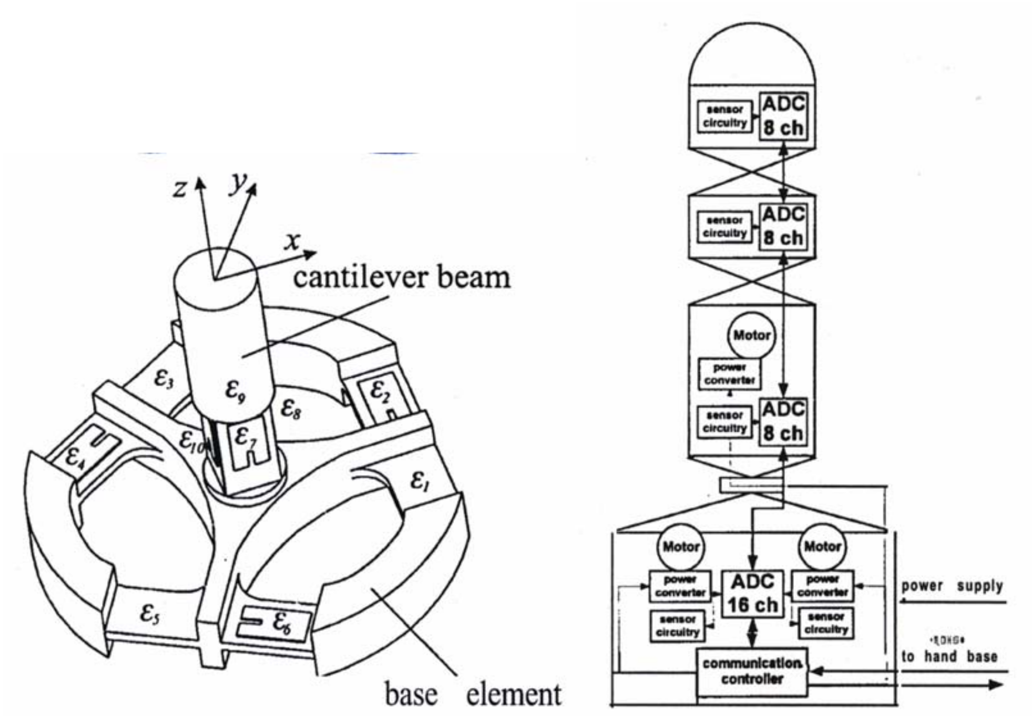

Figure 71 – a) 6 DOF force-torque sensor; b) finger electronic configuration [Hirzinger et al 2004]

Each finger has three degrees of freedom plus an additional degree of freedom at the base joint to provide reconfigurability between grasping modes. Each fingertip is capable of exerting 30 N of force. The hand has a set of joint position and joint torque sensors. This force torque sensor can measures forces and torques of 10-40 N and 150 Nmm respectively. The robotic hand requires force and slip sensors on the fingertips and a force sensor in the palm. If tip force sensors are activated, a precision grip may be adopted, while if palm force sensors are activated, a power grip may be adopted. Slip sensors such as piezoelectric ceramics may exploit the fact that slip generates low frequency ~10-100 Hz vibrations. If slippage is detected by the slip sensors, the hand automatically tightens its grip until the sliding is terminated - this a reflex. If slaved to a human teleoperator, the commanding of such robotic hands may be implemented through myoelectric signals produced by the contraction of the operator's forearm as grip increases, similar to robotic limb prostheses.

Passive compliance yields great physical robustness in locomotion and manipulation capabilities despite highly variable terrain and highly variable payloads respectively. Indeed, this is the basis of the Remote Centre Compliance device used in some manipulators that eases the problem of peg-in-hole tasks (Figure 72).

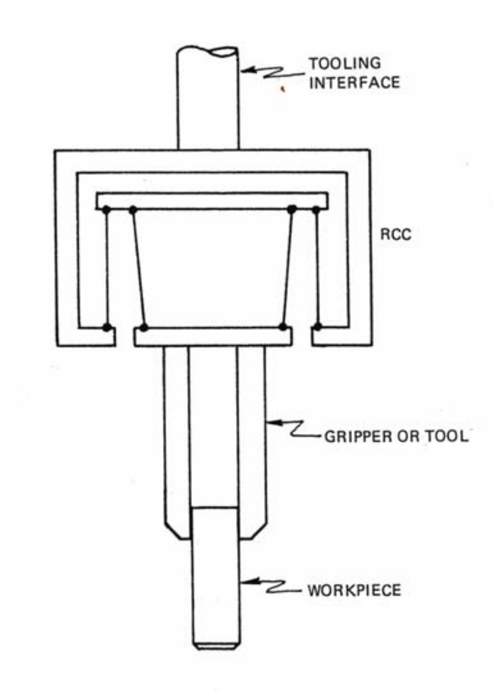

Figure 72 – Remote centre compliance (RCC) device for compliant peg-in-hole tasks (Nevins & Whitney 1994)

The compliance aspect offers an additional passive component to active control that is generally absent in artificially engineered machines. This approach to compliant structures offers greater freedom in the selection of control parameters and the distributed sensor approach offers the controller greater prospects for closed loop controllability.

3.5.6. Compliant Structures in Actuation Control

Engineered robots are highly primitive in comparison to the simplest organisms with regard to their limited sensory capabilities, limited robustness of pattern recognition capabilities, limited adaptability of their control systems, and limited actuation capabilities. The importance of physical bodies for robots is generally under-estimated in control design. A robot is an integrated machine including active sensory input, dextrous mechanical output and a complex control system. Control theory generally encapsulates the physical/structural dynamics of a system within the control gains in only a limited fashion. The dynamics of a physical system are typically assumed to be linear or are linearised through feed-forward compensation. Even adaptive feedback control approaches generally assume linearity and depend on the environment changing more slowly than the adaptation update rate. The evolution of animal control systems occurs in parallel with body morphological evolution – they cannot be treated separately as they co-evolve. Adaptive (smart) structures are essentially structures that can alter their geometric characteristics according to either the external environment or to imposed commands. Smart structures generally involve the integration of sensing (strain, temperature, pressure, etc), processing this information, and reacting to that information. The traditional approach to structural design has been to incorporate sufficient robustness to ensure mechanical integrity. Biological materials often incorporate this function – the horns of many mammals are sheathed in fibrous keratin which is insensitive to notch weakness when damp. Smart materials however are multi-functional – multi-functional structures for spacecraft are a step in this direction. Multi-functional structures integrate electronic, thermal and structural functions – by integrating data transmission, power distribution and heat pipes into the load-bearing structure, the overhead of cabling and connectors can be much reduced. In particular, such structures may be manufactured to provide active vibration damping by incorporating a piezoelectric or electrostrictive material within the structure [Wada 1990]. It is possible that artificial muscles of shape memory alloys, electrorheological fluids, electroactive polymers and other materials with piezoelectric, thermistor and strain-sensitive sensors may be embedded into multi-functional structures augmented by fibre-optics. Such materials may be manufactures using shape deposition manufacturing. Micro-technology based components for signal processing and data processing and transmission offer significant mass reductions making micro-satellites (~10-100 kg), nano-satellites (~10 kg) and pico-satellites (<1 kg) viable. There are fundamental limits to miniaturisation however – solar array area, power storage densities, radiation shield thickness, and resolution limits to optical and radio antennae.

However, it is beneficial to remember that muscles normally consist of a number of muscle fibres that, individually, are very weak, but collectively can perform the required tasks. Through muscle control of a set of these fibres, muscles can perform a wide range of tension and contraction speeds and strengths. Without jumping, the only way to increase walking speed is to increase step frequency (which is impacted through tension changes of the muscles). A range of contraction speeds is absolutely necessary in order to be able to walk at different speeds (Wendler 2002).

Dubowsky (1999) presented the need for self-transforming robots for planetary exploration which alter their configuration to perform a wide range of tasks and overcome a wide range of physical obstacles encountered on planetary surfaces such as cliffs, ravines and craters. These self transforming explorers (STX) were based on Articulated Binary Elements (ABE), compliant structural members (constructed from shape memory alloy wires) with embedded actuators, sensors and data/power conduits. The STX body is composed of a network of multi-facetted nodes constructed from delrin/polyurethane with each face providing a connection point for ABEs. Shape change is implemented by contracting or extending individual ABEs and topological configuration alterations within the structure. Releasing and locking of the bistable compliant mechanisms at the ends of the ABE and the node faces are based on an “electronic handshake” latch. Jacobs (1998) have suggested a biologically-inspired two-legged robot with 18 artificial spring-like muscles controlled by a small number of fuzzy if-then production rules to control both standing (7 rules) and rhythmic walking (4 rules). This Bio-Bot comprises 8 aluminium segments: two feet, two lower legs, two upper legs, a pelvis and a trunk with human scaling. The muscles comprise of flexible tubing within a braided sleeve, one end of which is connected to an solenoid valve air supply. Pumping pressurised air into the tube causes tube shortening and generates torques at the joint. Each muscle group comprised a combination of long one-joint controlling joint position and shorter two-joint muscles controlling joint torque. Each leg comprised 8 muscles – 5 one-joint and 3 two-joint muscles) to give control of speed, height and trunk. Each joint has stretch sensors, a three axis accelerometer package in the trunk gives orientation angles, and four force sensors on each foot complete the sensor package. Bio-Bot maintained standing using hip extensor and knee extensor muscles. Reinforcement learning provided the basis for adjusting the timing of each functional muscle group.

Biological structures tend to be close to critically damped unlike artificial mechanical structures which are heavily under-damped generating the problem of vibration. High control gains tend to generate instability, so are usually avoided. Control of flexible appendages are challenging problem as they require suppression of vibration but high passive damping can substantially reduce this problem. Closed loop control cycles limit the rapidity of response so such compliance is essential for rapid adaptation to environmental conditions, especially as machine-environment interactions are essentially chaotic [Nehmzow & Walker 2003]. Hence, situatedness of an agent interacting with its real world environment makes it behave as a dynamical system with a characteristic phase space trajectory which may be quantified by its correlation integral:

$$C(r) = \frac{\theta}{n(n-1)}$$

where:

θ = *number of times that* $| D_i - D_j | < r$
$D_{i,j}$ = *coordinates of attractor at time t=i,j*
R = arbitrary correlation distance
n(n-1) = maximum number of cases where the condition holds

For an agent interacting with the real world, the trajectory in phase space lies on a strange attractor, so that two trajectories that begin close together will diverge over time. Their separation after time t is:

$$d = d_o e^{\lambda t}$$

where:

λ = *Lyapunov exponent*

The larger the positivity of the Lyapunov exponent, the greater the degree of chaotic divergence. Experimentally, robotic behaviour has indeed been found to be chaotic. Closed loop control cycles – particularly those associated with force control – limit the rapidity of response so structural compliance is essential for rapid adaptation to environmental conditions, especially as force control is essentially chaotic as small changes in applied torques generates divergent sensed forces on contact with stiff environments ~10^4–10^6 N/m which cannot in general be compensated for by closed loop force control. Grasping (and indeed any form of physical environment interaction) is a fundamental component of manipulation involving interaction with the external world with a high variation in impedances from zero (prior to contact) to high values (if the payload is stiff). Passive compliance in the structure absorbs the energy of impact and effectively reduces the stiffness of the environment. In biological structures, low structural impedance ensures that the limb can "give" to unexpected disturbances so that large forces are not generated. In particular, structural compliance provides the basis for “preflexes” within a control system for dealing with interaction forces. The adaptive control system must include modelling aspects of the environment rather than merely the limb itself alone. The internal models of the dynamics of limbs must therefore be based on forces with limb compliance as a variable – they must be based on proprioception rather than external sensory patterns. The most basic form of closed loop control is the reflex which adjusts the compliance and damping of limbs based on stored neural set point models of limb inverse dynamics. However, modelling errors are inevitable suggesting that adaptive feedback is necessary. It is apparent that pure feedback control cannot be employed in limb movement as neuronal feedback loops would be too slow. There are thus two primary feed-forward schemes that may be adopted. Equilibrium based control methods suggest that the central nervous system relies on the spring-like properties of the muscles with control being based on reflexes. Such preflexes maintain stability while minimising the effects of disturbances. Preflexes are near zero-order responses due to the intrinsic properties of the structure due to passive compliance. Disturbances are rejected within ~20-70 ms overcoming the problem of the time delays ~200-500ms inherent in proprioceptive feedback control. This enables the use of open loop controllers such as central pattern generators. It is impedance control that provides the basis for dealing with uncertainty – impedance control utilises an internal model for feed-forward control which is learned by identifying certain parameters such as mass, dimensions, friction, etc based on sensory signals. The control of stiffness provides the basis for using impedance to control externally applied forces [Hogan 1985, Hogan et al 1985]. Impedance control is a methodology for robotic control based on a biological paradigm - the cerebellum stores muscle activation patterns learned from the association of proprioceptive stimuli indicating co-contraction of agonist-antagonist muscle lengths. Modulation of contraction in the timing and amplitude of muscle activity regulates the visco-elastic properties of the joints to move between postures without using sensory feedback.

Muscular activation patterns are generally stored neurally rather than processed anew – in mammals, such patterns are learned from the association of proprioceptive stimuli with co-contraction of agonist-antagonist muscle lengths. These patterns of association are stored in the cerebellum. Muscles behave as tunable springs to regulate the timing and amplitude of the viscoelastic properties of the joints to move between positions. Muscle length-extension defines the equilibrium position and stiffness of each joint – this may be modelled as a potential function of the joint angle (the derivative of which is force). The postural force field is characterised by the product of a displacement vector and a stiffness matrix. Muscles behave as tunable springs by virtue of their arrangement around joints in agonistic/antagonistic pairs. The elasticity of the muscle is determined by its activation level which determines their length-extension α which in turn defines the equilibrium position and stiffness of the joint. This may be modelled as a potential function of the joint angle, the negative derivative of which is the generalised force. The postural force field is characterised by the product of a displacement vector and a stiffness matrix. Activation of agonistic/antagonistic muscles determines the potential minimum (determined by the α ratio) and the potential minimum curvature (determined by the α sum). Displacement from the potential minimum results in a restoring torque independent of feedback. Thus all control acts through a single time-dependent potential function with the muscles themselves computing the joint torques. The function of the nervous system is to transform the desired trajectory into a sequence of joint equilibrium positions and stiffnesses determined by a global time-dependent potential function. This method of hybrid position/force (impedance) control is still under development though offers great promise. It provides a unified framework for all kinds of manipulator motion involving dynamic interactions while preserving stability during interactions with the dynamic environment. Learning impedance modulation is necessary which may be implemented through an objective performance function. It is important to note that biologically, such control of joint viscoelastic properties is achieved by feed-forward open loop mechanisms from learned cerebellar inverse dynamics models due to delays and instabilities in online feedback mechanisms.

Insects store energy in resilin within the cuticle to drive their wingbeats during flight. Spiders store hydraulic pressure in their blood by compressing their bodies to drive their legs. There are four major types of actuation in biological organisms: amoeboid movement by phagocytotic mechanisms, ciliary motion (to convey material along biological tubes), bacterial flagellar movement (which resembles the rotation of electric motors), and muscular contraction. Muscular contraction is of great interest as it is a versatile means of actuation for both locomotion and manipulation. Muscle is powered by the conversion of ATP to ADP with the release of discrete packets of chemical energy. A contractile protein, myosin slides along another contractile protein, actin, to provide mechanical movement – contraction is caused by the sliding motion of the head section of myosin along the actin molecule. Myosin acts as a molecular motor whereby ATP-generated energy provides for the power stroke followed by a recovery stroke. Animal muscle is capable of exerting high lever torques with great precision. Muscles also provide additional functions as shock absorbers and springs by imparting damping [Kornbluh et al 2002]. Muscles can operate at a wide range of frequencies and speeds for a wide range of functions [Dickinson et al 2000]. They are characterised by highly non-linear multi-parameter dynamic systems. In particular, muscle fibre contraction provides the basis for rotary joint motion – Hill's model comprises a contractile and series elastic components such that the cross-section of the muscle defines its force potential.

Current muscle-like actuators include shape memory alloys, dielectric elastomers and electroactive materials [Fijita 2001], though polyMEMS actuators have been proposed to exploit cellular construction [Horning & Johnson 2002]. Shape memory alloys exhibit limited deformation response but electroactive polymers are lightweight plastic strips of high bending capacity. Electroactive polymers and ceramics are characterised by an electromechanical deformation response when a voltage is applied. Electroactive polymers offer greater deformation response than electroactive ceramics (which include piezoelectric ceramics) and can provide high strains [Bar-Cohen et al 1998, 1999]. They provide fast response and good controllability. Electronic electroactive polymers deform by virtue of electronic carrier motion but this requires high voltages, while fluorinated ionic electroactive polymers deform by virtue of ionic carrier motion and require only low voltages ~1-3V and currents ~mA. Ion-exchange polymer-metal composites offer potential application for large bending displacements under ac inputs [Shahinpoor et al 1998]. Their chief disadvantage is that they require water to operate but such materials can operate at low temperatures ~-140^{o}C. Electroactive polymers include ion exchange membranes, gel polymers, and perfluorinated sulphonic polymers – conducting polymers appear to offer greater capabilities than biological muscles. In addition, there are electrorheological and magnetorheological fluids which alter their viscosity in response to applied electric or magnetic fields. The McKibben artificial muscle is based on inflation of an inflatable bladder with a compressed gas [Klute et al 1999].

3.6. BIOMIMETIC METHODS OF ENERGY GENERATION & STORAGE

One of the most promising potential developments in biomimetics for space applications is artificial photosynthesis and the development of related redox reactions. Artificial photosynthesis offers a number of advantages over existing solar panel technology. Firstly, the efficiency might be much greater than traditional solar collection technologies. As it is based on a biochemical architecture that has evolved over 3.5 billion years, it is a much more refined method of gathering solar energy that is much more efficient per unit area than current solar panel technology. Secondly, it opens itself up to possible improvements. Molecular biology can be used to make changes in the basic building blocks of the photosynthetic apparatus in order to make more efficient electron donors as I will discuss in later sections and furthermore, photosynthetic reactions can be linked into other forms of energy production and utilization as will also be reviewed.

Plants do not need to be 100% efficient at photosynthesis. Indeed, the highest efficiency in plants is around 1-2% (Wittwer, 1980). They only need to be as efficient as is required to gather energy for their metabolic processes. Any higher efficiency does not provide greater competitiveness in the natural world. However, for applications in space, the more efficient the solar collection technology is, the less area it will take up and the more energy per unit area will be gathered, and so there are good reasons to want to make photosynthesis more efficient. Directed alterations in the architecture of the photosynthetic structure could be used to make more efficient solar collectors and this is an area of emerging focus.

Artificial photosynthesis could be used to provide power for a number of space systems including main power using systems (life support, scientific instrumentation etc.) and systems on the smaller scale such as nano-machines used to care of the exterior of spacecraft where sunlight impinging on the craft could be used as the source of energy. Construction and repair could be powered by artificial photosynthesis.

To build more efficient solar collectors it is first necessary to understand the structure of the photosynthetic apparatus to be found in plants and to consider the parts of the apparatus that might be improved. Plant photosynthesis has to date provided the most information for the development of artificial photosynthetic systems.

Photosynthesis has at its core some basic reactions sequences (Glazer and Melis, 1987). The principle behind photosynthesis is that water is oxidised to produce oxygen and carbon dioxide is reduced to form carbohydrates according to the reaction below:

$$n\mathrm{CO}_2 + 2n\mathrm{H}_2\mathrm{O} + \text{light} \rightarrow (\mathrm{CH}_2\mathrm{O})_n + n\mathrm{O}_2 + n\mathrm{H}_2\mathrm{O}$$

In plants these reactions occur in chloroplasts, essentially sacks containing the photosynthetic apparatus. The principle behind the chloroplast is important because a similar system might be used to contain artificial photosynthetic apparata in space.

The chloroplast contains the enzyme-rich gel, the stroma, which contains the enzymes that convert carbon dioxide to carbohydrates. Throughout the stroma are thylakoids in which light energy is used to oxidise water and to form ATP and NADPH, the energy rich compounds needed by the stroma to convert the carbon dioxide to carbohydrates [Anderson 1986].

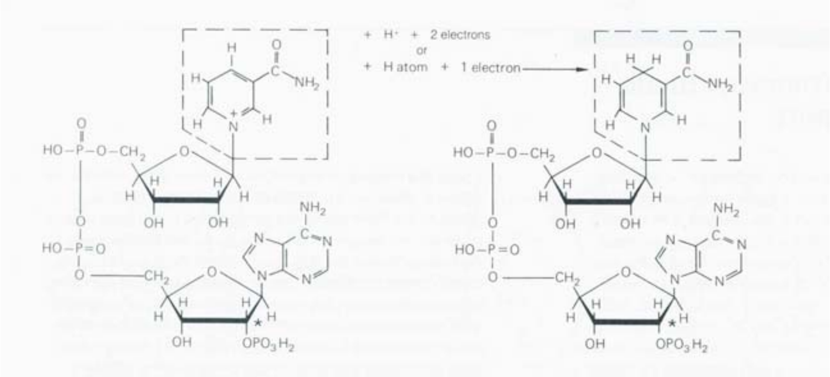

Figure 73 – NADP+ and NADPH. The energy rich compounds used to store energy gathered by photosynthesis.

Within the thylakoids are chlorophyll a and b, which are the light harvesting compounds responsible for gathering sunlight. About a half of the total membrane concentration of thylakoids is accounted for by chlorophylls.

A principle of light absorption important for natural and artificial photosynthesis is called the Stark Einstein Law. It is a law that governs the uptake of photons of sunlight by light harvesting pigments. It essentially says that a pigment can only take up one photon at a time. This photon imparts its energy to one electron which can then rise in energy level creating an ‘excited state’ (see Figure 74).

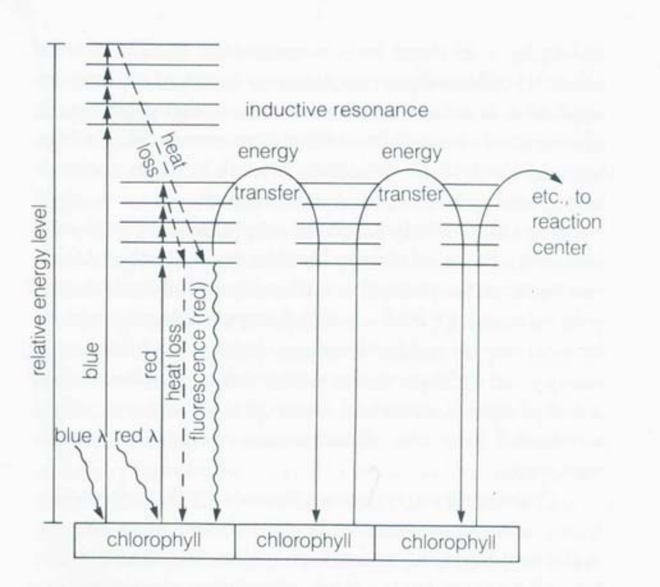

Figure 74 – Electron excitation states in chlorophyll caused by light absorption. Each reaction takes up one photon governed by the Stark Einstein Law.

The Stark Einstein law is important because it limits the rate of light uptake by a system of light harvesting. If, for example, one had an artificial photosynthesis system in space and one wanted to double light harvesting it would be necessary to have twice the reaction centres. It is not possible to get one molecule to absorb two electrons at a time. Here the design of artificial photosynthetic power systems in space is constrained by a fundamental law of physics and this is an important factor to take into account when designing artificial photosynthesis systems.

In natural photosynthesis, light is first gathered by chlorophylls and carotenoids. Chlorophylls absorb in the blue and red regions and carotenoids only absorb in the blue region of the spectrum. The energy of the absorbed light, after being absorbed by an electron, now in a more excited state, has a number of possible fates. It may be re-radiated as fluorescent light and it may be re-radiated as heat. In the case of blue light, some of the energy is always reradiated as heat when the electron drops to a lower state because blue light, being a shorter wavelength than red light, is more energetic.

The absorption of light as well as its re-radiation as heat can generate free-radicals that can damage a wide range of molecules. Ingeniously, carotenoids, as well as being light absorbers, also quench reactive oxygen states. The production of free radicals by molecules designed for artificial photosynthesis systems should also be a consideration when selecting molecules for use in an artificial system. It is unlikely that the energy for light absorption can be harvested with 100% efficiency and the potentially deleterious effects of this lost energy must be considered. Thus, in artificial photosynthesis systems in space independent molecules might be used to quench reactive oxygen states or molecules used to gather light might also have a quenching role in analogy to carotenoids.

Another way for the energy gathered from sunlight to be used is for it to be moved from the light absorbing molecule to another by 'inductive resonance' instead of being fluoresced or lost as heat (Malkin, 1988). This is the most important process for photosynthesis as it represents the principle behind the transfer of the energy from the light harvesting apparatus to the reaction centre where the reactions of photosynthesis occur. One of the major challenges in artificial photosynthesis is the transfer of electrons from the pigment harvesting the light to the reaction centre. In natural photosynthesis this charge transfer occurs from the chlorophylls or the carotenoids to the chlorophyll a in the reaction centre.

The process of converting the light harvested by chlorophyll into the photosynthetic reactions is carried out by two photosystems which are currently the focus of active efforts in artificial photosynthesis research and would be necessary for artificial photosynthesis systems in space. The first of these photosystems is photosystem II (PSII). Within PSII is the P680 reaction centre. About 250 chlorophyll a molecules and numerous carotenoids (xanthophylls) gather the sunlight and transfer the energy by inductive resonance to the reaction centre. These primary reactions of photosynthesis can be summarised as follows and they occur on picosecond time scales:

$$\text{P680} + \text{pheophytin a} + h\upsilon \rightarrow \text{P680}^{+} + \text{pheophytin a}^{-} \quad \text{(3ps)}$$

$$\text{P700} + \text{chlorophyll a} + h\upsilon \rightarrow \text{P700}^{+} + \text{chlorophyll a}^{-} \quad (<10\text{ps})$$

These light harvesting pigments themselves make up what is termed the PSII light harvesting complex (LHCII) [Smeekens et al. 1990].

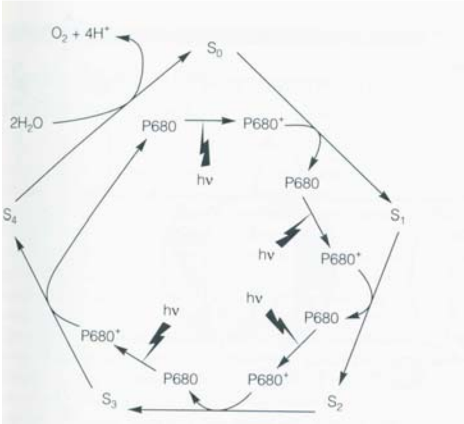

Figure 75 – Steps in the electron transfer chain for P680.

The light harvested and the energy then transferred into PSII is used to oxidise water [Lagoutte and Mathis 1989, Rutherford 1989]. Each water molecule oxidised requires two photons to be harvested.

$$2H_2O + 4\text{ photons} + 2PQ + 4H^{+} \rightarrow O_2 + 4H^{+} + 2PQH_2$$

The electrons from the water are then used to reduce a molecule plastoquinone (PQ) to its fully reduced form (PQH_2). In some ways, the plastoquinone can be regarded as a mechanism to 'fix' the light that has been harvested in PSII. Once the plastoquinone has been reduced, the first stage of photosynthesis is complete and the molecule can then be used to drive then next set of reactions.

The next stage of the process is to get the electrons from the PQH_2 to Photosystem I (PSI). This function is carried out by an iron-sulfur protein – the cytochrome b_6-cytochrome f complex. It does this by oxidising the PQH_2 and then by reducing a small copper-protein named plastocyanin.

$$2PQH_2 + 4PC(Cu^{2+}) \rightarrow 2PQ + 4PC\,(Cu^{+}) + 4H^{+}\text{ (lumen)}$$

The next stage of the photosynthetic system involves the use of the electrons by PSI. PSI is complex because as well as accepting the electrons that have come from the harvesting of light in PSII it also can harvest light itself [Evans and Bredenkamp 1990]. In the PSI is yet another light harvesting complex like in PSII. The PSI complex has about 100 molecules of chlorophyll a. The overall reaction of PSI can be written as follows:

$$\text{Light} + 4PC(Cu^{+}) + 4Fd(Fe^{3+}) \rightarrow 4PC(Cu^{2+}) + 4Fe(Fd^{2+})$$

The electrons gathered by light absorption in the photosystems are transferred to a molecule called ferrodoxin [reviewed in Pschorn et al. 1988]. This is a small low molecular weight protein attached loosely to the stroma side of the thylakoids.

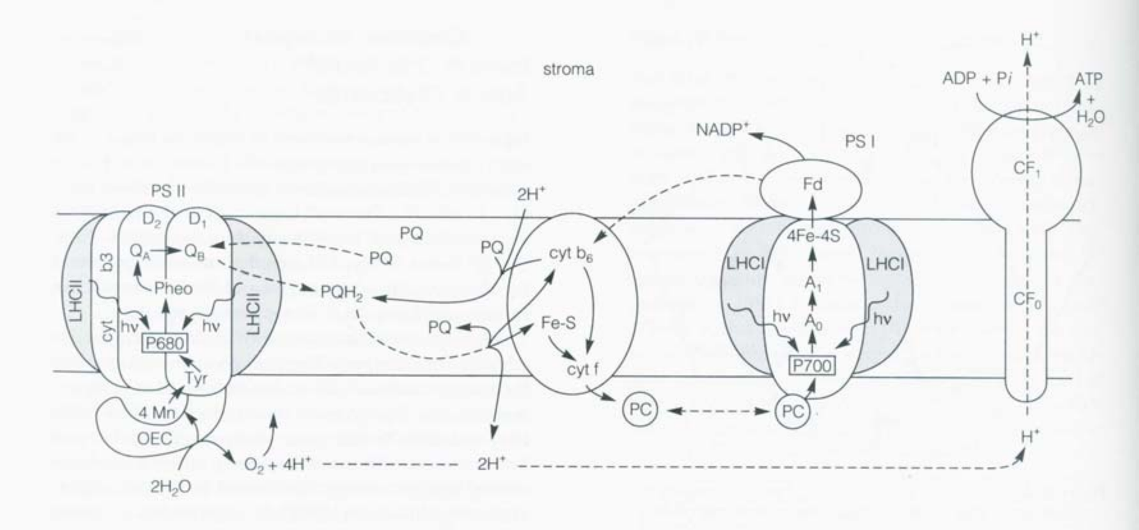

Figure 76 – Summary of the major pathways of natural photosynthesis. The basic of this natural system are replicated in artificial photosynthesis.

Ferrodoxin is involved in the final electron transport step of photosynthesis to reduce the molecule $NADP^+$ to NADPH.

$$4Fd(Fe^{2+}) + 2NADP^{+} + 2H^{+} \rightarrow 4Fd(Fe^{3+}) + 2NADPH$$

This final reaction creates the NADPH that will then be used to reduce carbon dioxide. The carbon dioxide is reduced to a compound, 3-phosphoglyceric acid, the first stage in the famous Calvin cycle. Once formed, this compound is the precursor to a whole series of more complex carbohydrates that form the basis of biosynthesis. It is not the purpose of this paper to discuss the various steps in the formation of carbohydrates, but rather the principles of gathering sunlight as a source of energy. However, it is useful to note that once the first carbon compounds have been produced from carbon dioxide there are potentially limitless ways in which they could be used in various artificial biosynthetic reactions to create compounds of use, such as carbohydrates for food or complex carbon molecules for fabrics. Perhaps chemical reactions could also be used to transform these first stages of the Calvin cycle into alcohols, thus generating methanol and other reactants for propulsion systems.

Artificial photosynthesis can therefore be seen as the early stages of a complex manufacturing process. One could imagine vats of chemicals with artificial leaves generating the reducing potential to provide organic molecules for any number of organic chemical syntheses. Once simple carbon compounds are available then they can be transformed according to industrial need.

The reason for using light energy in photosynthesis is that water is thermodynamically difficult to oxidise. Light energy provides the necessary energy to make this transition occur and this is the basis behind photosynthesis. We might also wonder why it is necessary to have such a complex arrangement of photosystems and would this be necessary in artificial systems used in space? The reason for the stepwise transfer of electrons through the system is to reduce back reactions that might reduce the efficiency of photosynthesis and prevent the oxidation of water.

Another by-product of photosynthesis is ATP, adenosine triphosphate, an energy rich molecule that can be used in a host of biochemical reactions.

$$\text{ADP} + \text{Pi} + \text{light} \rightarrow \text{ATP} + \text{H}_2\text{O}$$

The energy in the molecule resides in the phosphate bonds. It can be used to fuel bacterial flagella and might find use in future nanomachines as a source of energy for driving tiny motors and other mechanical devices as discussed in more detail later. Here we see how artificial photosynthesis could be used to make the energy to drive tiny machines in space stations and other facilities. The production of ATP is made possible by the protons that are generated during the various photosynthetic reactions. Protons are produced during the oxidation of water and the oxidation of PQ and these protons are released into the lumen side of the thylakoid. The protons then pass back through the ATP synthase molecule that straddles the thylakoid membrane [Anderson and Anderson 1988]. The ATP synthase has two parts to it, a stalk and a spherical headpiece.

The movement of protons through the ATP synthase results in the generation of ATP from ADP and P. The system is clever, because it ensures that the protons produced during oxidation reactions are not wasted. Thus, as well as using reduced carbon dioxide to produce compounds for biosynthetic reactions, photosynthesis is also a source of energy. These two abilities might also be paired in artificial photosynthetic systems.

The passage of electrons through the whole photosynthetic system gives some indication of the complexity of the natural systems. Systems of artificial photosynthesis are envisaged to be much simpler and would involve just one or two molecules to complete the electron transfer reactions. However, artificial systems still rely on the same principles – the transfer of electrons to molecules to prevent back reactions – in other words the partitioning of the various reactions.

To understand the complexity of the natural system it is worth reviewing in more detail the way in which the electrons are transferred through the entire system. As one water molecule is oxidized, two electrons are released for transport. The first compound to receive these electrons, one at a time, is tyrosine, an amino acid in the D1 protein of the second photosystem. The tyrosine then passes them to the P680 molecule in the reaction centre that was described above. However, P680 can only accept this electron if it has been excited by light. The light received by the light harvesting compound which is channelled to P680 causes it to become oxidised and it can now accept the electron from the tyrosine, hence the electron now ends up in the reaction centre of the photosystem.

P680 then transfers its electron to phaeophytin which itself passes the electron onto a specialized plastoquinone that is attached to D2, another polypeptide bound to the P680 reaction centre. The plastoquinone then transfers the electron to yet another type of plastoquinone. To reduce both of these plastoquinones requires two electrons and also two protons. This turns out to be important because these two protons will eventually be released again to drive the formation of ATP. Each of the reduced plastoquinones carries the electrons to the cytochome complex. Once accepted by the cytochrome complex these electrons are passed, one at a time, to the ferrodoxin.

So far we can get some idea of how photosynthesis can create energy. For every molecule of water, four protons are formed ($4H^+$). As two water molecules are required to release one molecule of oxygen and fix one molecule of carbon dioxide, then the accumulation of eight protons in the lumen will result from the fixation of one carbon dioxide molecule. Furthermore, we should remember that each two photons are required for each electron transported so that eight photons are needed to release one oxygen molecule and oxidise two molecule of water [Osborne and Geider 1987].

Once the cytochrome complex or the ferrodoxin accepts an electron, then the electron is transported to cytochrome f. The electron is then donated from there to the copper in plastocyanin. Each mobile plastocyanin then carries an electron to photosystem I. The first molecule to accept an electron from the plastocyanin is P700. Just as for P680, P700 cannot accept an electron until it has lost an electron and like P680, this occurs after it has lost an electron after receiving light energy from the light harvesting complex. Once this has occurred then P700 can take an electron from the phyocyanin. The electrons then pass through the various molecules within P700 including a special chlorophyll and a phylloquinone. Then the mobile ferrodoxins accept the electron and transfer it to $NADP^+$.

The reactions by which the electrons ate transferred across the thylakoid membranes are called non-cyclic electron transport reactions because they do not cycle back to the water molecules. This is in contrast to cyclic electron transport reactions where electrons are fed back from P700 into the various electron transfer reactions.

We can see that the natural process of photosynthesis is a complex one, but it has quite well defined principles by which it operates. The principles of artificial photosynthesis are much the same. The idea is to use sunlight to drive the otherwise thermodynamically unlikely oxidation of water, the result being the generation of energy in the form of ATP and the reduction of carbon dioxide into the early stages of complex carbohydrate production.

One promising avenue of artificial photosynthesis research is the use of porphyrins and fullerenes as building blocks for artificial photosynthesis systems [Imahori et al. 2003]. Photodynamical studies on these systems show that fullerenes accelerate photo-induced electron transfer and slow down charge recombination. They slow down charge recombination much more effectively that the usual electron acceptors used in these systems such as the aromatic compounds, quinines and imides. These properties are attributed to small reorganization energies of fullerenes. Research has focused on the development of electron and energy transfer reactions on gold and indium-tin oxide electrodes. When porphyrin- modified nanoclusters are used as electron acceptors artificial photosynthetic systems can be constructed and they might be used as the basis for other molecular devices.

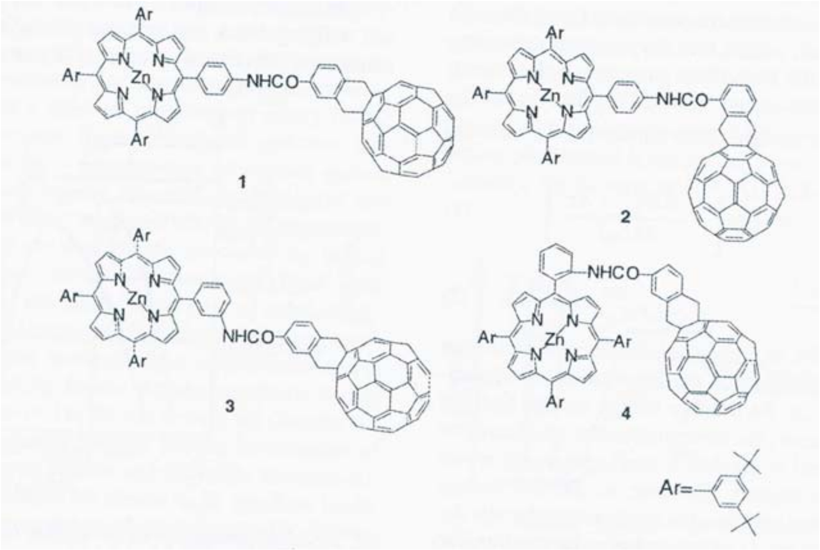

Figure 77 – Porphyrin-fullerene systems can act as artificial electron donor-acceptor systems for artificial photosynthetic systems (Imahori et al. 2003).

Artificial photosynthesis systems are based on some fundamental principles of engineering design discussed here in the case of the porphyrin-fullerene system. The key requirement is to achieve multiple electron transfer so that each charge separation prolongs the life time of the final charge-separated state. This same approach is used in natural photosynthesis using the various intermediate electron transfer states described earlier. In artificial photosynthesis, at least three different components are covalently linked to create a sequential electron transfer. The first artificial systems used a porphyrin tethered to two quinines [Migita et al., 1981]. Carotenoid-porphyrin-quinone systems have also been devised.

These systems showed early promise. However, to achieve efficient charge separation requires an empirical approach to understanding the problem. The electron transfer rate constants that ultimately determine the effectiveness of the reaction are determined from the equation,

$$K_{ET} = (4\pi^3/h^2\lambda_B T)\, V^2 \exp\,[-(\Delta G^0_{ET} + \lambda)^2/4\lambda k_B T)$$

where the free energy change ΔG is given by:

$$\Delta G = (\Delta G^0_{ET} + \lambda)^2/4\lambda,$$

where:

λ = *reorganization energy*
V = *electronic coupling matrix element*
k_B = *Botzmann constant*
h = *Planck constant*
T = *absolute temperature*
$\Delta G^0_{ET,}$ = *free energy change.*

In natural photosynthesis the λ value is optimised for each reaction that occurs and this is defined in curves known as Marcus functions. The goal in artificial photosynthesis is to similarly optimise the reorganization energies and this has conventionally been done by choices in the solvent in which the reactions occur. Furthermore, the process can be improved by seeking molecules with small reorganization energies that minimize the back reactions that lead to inefficiency in electron transfer.

Fullerenes possess small reorganization energies compared to aromatic molecules such as quinones and hence their interest in artificial photosynthesis schemes. The reduction potential of fullerenes (C_{60}) is similar to the quinones. These fullerenes electron acceptors have been coupled to porphyrins as the electron donors in artificial systems.

Some examples of porphyrin-fullerene systems are shown in Figure 77 above. These systems are typical dyad systems. Time-resolved transient absorption spectroscopy has confirmed that photo-induced electron transfer occurs in these systems. The reaction occurs from the porphyrin excited single state to the fullerene moiety and during transfer the zinc porphyrin radical cation and the fullerene radical anion can be detected, demonstrating charge transfer. Derivations of these first systems to be used have mainly involved changing the positioning of the porphyrin on the fullerene which optimises electron transfer properties. By changing the fullerene for quinone, the superior performance of the fullerene as an electron acceptor has been demonstrated, showing how it has a smaller charge recombination time than the older systems that use quinones.

The superior performance of the fullerenes is explained by the fact that the charge in the fullerene anion radical is distributed over the entire C_{60} structure, rendering the reorganization energy smaller than for the corresponding quinones, where the charge is localized mainly to two oxygens. The delocalised electrons in the fullerenes therefore assist in reducing the charge recombination and increasing the effectiveness in electron transfer.

As stated earlier, one means to reduce the reorganization energies is to change the solvent. The most effective solvents are non-polar solvents. The reorganization energy is also a function of the separation distance between the electron acceptor and donor. Therefore, small separation distances within non-polar solvents should be the most effective systems for artificial photosynthesis. Studies of the fullerene-porphyrin systems discussed here in benzene have shown that that this non-polar solvent is indeed highly effective as the smallest values yet achieved in reorganization energies have been observed. Thus, fullerene and porphyrin systems, when coupled, could be used as the basis of artificial photosynthesis systems in space.

The fullerene and porphyrin systems and other systems like them can act as the basis for the electron transfer reactions. However, the electron transfer reactions must then be converted into electrical energy to complete the artificial photosynthesis system. One means of doing this is to construct artificial cell membranes, in analogy to natural cells. However, these membranes have been notoriously unstable and unreliable, particularly as the basis of long-lasting artificial photosynthesis systems. Other methods are required. These problems are particularly acute when applied to space, because conditions inside of spacecraft and certainly outside may be hostile to cell membranes in an isolated state. The approaches being used to make these systems for potential commercial use on Earth might have great promise for use in the extreme conditions of space.

The self-assembly of molecules onto stable surfaces is a promising new area, because the self-assembled structures can become highly ordered, similarly to the order within a cell membrane and they are more likely to be long lasting. Recent studies have focused on self-assembly on gold electrodes. Porphyrin-linked systems can be successfully bound to gold electrodes using butyl groups to increase the solubility of the porphyrins in organic solvents and to suppress the quenching of the excited states of the porphyrin on the monolayer because of their aggregation, which could suppress the efficiency of the electron transfer.

Fullerenes can also be assembled on gold electrodes and using an electron donor in solution such as ascorbic acid, it can be demonstrated that electron transfer can be successfully achieved. The key to the success of the overall system is to be able to put both porphyrins and fullerenes onto a gold electrode. This has now been successfully achieved [Imahori et al. 2003] and porphyrin-fullerene systems can be packed nearly parallel on gold surfaces leading to highly ordered structures. The highest quantum yield in these devices so far achieved is over 50%, but it is still less than natural photosynthesis.

One way to reduce the quenching of the porphyrin charges, a great problem with the gold electrodes, is to use indium-tin-oxide (ITO) electrodes to attach the molecules. This may allow for an increase in the quantum yields achieved. ITO is a semiconductor that suppresses the quenching of excited states of adsorbed materials and may be better as a system. However, the assembly of structures on the surface of these electrodes is extremely difficult and as yet building a complete artificial photosynthesis system on these electrodes has not been possible. Experiments using ITO systems show that photocurrent generation on ITO is over 250 greater than that on gold electrodes because of the suppression of quenching.

A further problem with the existing systems is the difficulty of light capturing on electrodes. A more promising technique is the use of gold clusters, small gold particles of size <3 nm on which the electron acceptors and donors are arranged. Recent research studying assembly of these artificial photosynthesis systems on gold clusters looks promising. The development of monolayer protected cluster technology may offer a means for production of high quantum yield artificial photosynthesis systems.

Above I have described the efforts to create artificial systems that can take light energy and turn it into electron transfer reactions, the first and fundamental reactions of photosynthesis. In space, these same artificial systems would be used to gather extraterrestrial visible light and use it to trigger electron transfer reactions. However, once we have initiated electron transfer, the next challenge is to do something useful with it.

Thus, a further challenge in artificial photosynthesis is to replicate the reactions on the photosystem II side of photosynthesis – that is the reactions involved in the oxidation of water and the use of light driven charge separation to drive water oxidation and thus the production of oxygen. The purposes of these systems for the exploration of space are clear. Artificial photosynthesis systems linked to the PSII reactions could allow for the oxidation of water into breathable oxygen used in life support systems. Essentially therefore in a simplified way, one can regard these reactions as light driven electrolysis. These systems could be useful on spacecraft or on planetary surfaces where oxygen will be required for life support.

Furthermore, the formation of oxygen, a powerful oxidant which can be used in fuel cells or even burned, can be regarded as a way to fix light energy into a useful product. As oxygen has such vital life support uses as well as its possible uses as an oxidant in fuels, it seems a logical choice for the end product of artificial photosynthesis in space. It also has the advantage that it is the end product of natural photosynthesis and so by understanding these reactions in nature we will be more likely to be able to replicate them in artificial systems in space.

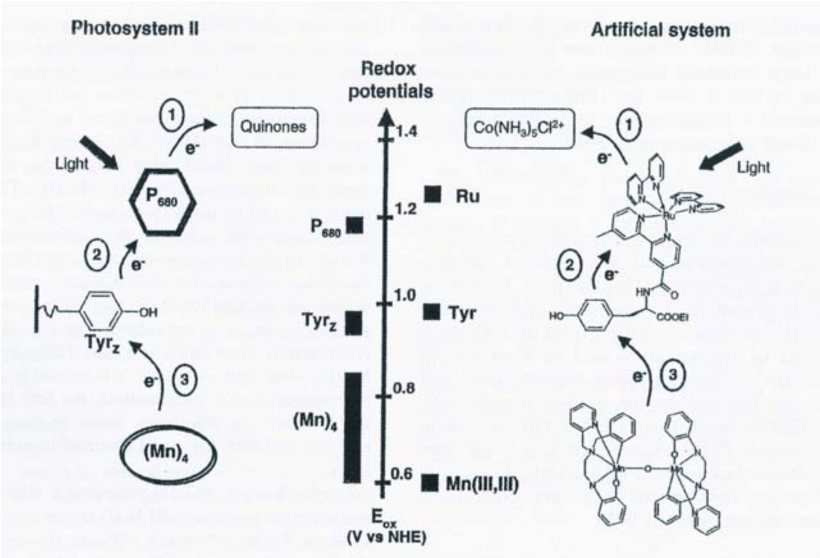

Figure 78 – An example of an artificial photosynthesis system compared with the natural PSII system.

As described earlier, the reactions leading to water oxidation start with the excited P_{680}. This highly oxidising state takes an electron from a nearby tyrosine residue (Tyr_Z), generating a neutral Tyr_Z radical. This radical then oxidises the manganese cluster that comprises the four manganese ions. The oxidation state of the manganese cluster is described by the S-state cycle. Each photon-induced charge separation advances the cluster one step in the cycle. After four of these cycles molecular oxygen is released and the cluster reverts back to its original state.

Four protons are released for each turn of the cycle and at each stage in the cycle the manganese cluster is stable on the order of several seconds until the next photon arrives. The detailed structure of the complex is not known, but the crucial step in the whole process is the formation of the O-O bond in the oxygen from water [Yachandra et al. 1996]. There are some models of how this might work. One suggestion is that the O-O bond is formed between an electron poor oxygen bound to a high-valent manganese and an electron-rich oxygen of an external water molecule [Limburg et al. 2001]. Recent studies on artificial manganese complexes show that oxygen can be released by reaction of water with a high-valent Mn=O species, itself formed by reaction of oxone or hyperchlorite with a dinuclear manganese complex [Limburg et al. 2001] giving early insights into how oxygen is actually formed from these reactions.

Efforts to replicate water oxidation to date have involved the linking of manganese to a photosensitizer [Sun et al. 2001]. The best photosensitizers to date are rubidium complexes as they are stable and can be made synthetically with relative ease compared to some other systems [Abrahamsson et al. 2002]. Thus, efforts have focused on Ru-Mn complexes as a means for artificial water oxidation after the electrons are made available from artificial photosynthesis. Similarly to the light harvesting reactions, one of the major challenges is to achieve small recombination energies. In analogy to the natural systems, the use of a tyrosine radical between the photo-oxidised Ru^{III} from the manganese dimer has proven to be a successful approach to achieving separation [Magnusson et al. 1997]. The kinetics of these tyrosine-including reactions are also found to be similar to the natural systems. To achieve a similar system as the natural PSII, more than one manganese ion must be used to create an analogous S-state cycle system, where electrons can slowly be gathered for each photon received until the water can be oxidised. Dimeric and trimeric Mn complexes linked to Ru(bpy)$_3^{2+}$ photosensitizers have been used [Burdinski et al. 1999]. These systems are still in the early stages, but linked to porphyrin-fullerene or similar systems, it may eventually be possible to realize an artificial system for light driven production of oxygen from water in space facilities.

Biotechnology may also assist in bionic development. Recent work on altering the reaction centres from purple bacteria has been successful in elucidating how charge transfer occurs in the reaction centres [Kalman et al. 1999]. The addition of manganese and similar artificial compounds allows the interaction of these genetically modified systems with artificial compounds to be investigated [Kalman et al. 2003]. Mutations have allowed for different amino acids with stronger manganese binding groups to be added. Either these genetically modified systems will elucidate the means by which water oxidation occurs and how it can be done artificially or they will themselves become incorporated into artificial systems.

Perhaps one of the most significant recent developments is the successful creation of a Ru-Mn complex that can be oxidised in three subsequent steps similarly to the natural PSII system. It was shown that in the presence of an irreversible electron acceptor the a Ru-Mn$_2^{II,II}$ complex could be oxidised to a Ru-Mn$_2^{III,IV}$ by repeated laser flashes [Huang et al. 2002]. This is the first major step towards a light-driven multiple oxidation system. This was close to the four electrons needed for water oxidation, but not yet sufficient. No water oxidation has yet been observed with this system, but ligand exchange and the binding of water molecules were observed, giving the tantalising possibility of water oxidation. Modified Mn complexes that come closer to the high-valent Mn-oxo species required for water oxidation have been prepared, but they appear to be quite unstable and further work will be required to being them nearer the natural system to be found in PSII.

As the reader can see, the systems under development are bringing us closer to the time when oxygen could be produced in space from artificial photosynthesis, but there also other things that we can do with the by-products of photosynthesis, the most notable of which are protons. Protons (or hydrogen ions) may seem an innocuous product, but they themselves can be used to drive the synthesis of chemical energy.

The oxidation of water generates protons which can be fed back through the membrane to generate the energy rich compound, ATP. This ‘proton gradient’ is the basis of energy production in organisms and understanding how protons are naturally pumped across membranes will be essential for designing artificial energy systems for space application, including fuel cells and tiny motors that depend upon proton pumps and could be driven by artificial photosynthesis.

We can gain insights into how proton pumps can be manufactured for space bionics by understanding how natural proton pumps work. Perhaps the most well-known natural proton pump is bacteriorhodopsin, the retinal protein of halophilic archaea that live in naturally salty habitats. These proteins are light-driven pumps and when coupled to artificial photosynthesis systems described earlier, could be the basis of energy generation and any number of small molecular devices.

The pump is initiated by light driven alteration of the all-*trans* retinal to the 13-*cis* form. This alteration creates a high-energy state and the relaxation of the molecule back to its original state generates a cascade of events that causes the proton pumping across the membrane. These reactions are dark reactions that do not require extra photons. The first key event is the protonation of the aspartate at position 85. This protonation causes a cascade of hydrogen-bond rearrangements that culminate in the release of a proton on the outside surface of the cell (extracellular side). The movement of the 13-methyl group of the isomerised retinal causes the rearrangements on the cytoplasmic side that result in the uptake of a proton. In other words, a proton is expelled and then taken up on the inside of the cell – the basis of pumping. In general this direction transfer of protons is believed to be the same mechanism in most proton pumps. Bacteriorhodopsin is highly thermally stable and has high quantum efficiency, but it is still not at the level of efficiency required for commercial applications or space systems. Efforts to improve the effectiveness of the protein have been successful and suggest that systems much more useful to space bionics will be developed soon using site directed mutagenesis, whereby single residues of the protein are altered to improve the systems [Wise et al. 2002].

The ATPase proton pump is also well characterized and is the basis of most ideas on artificial proton pumps that might be used in bionics. The pump is made up of numerous subunit assemblies with different functions. The larger part projects out of the membrane and the protons are translocated through a rotor that has many proton binding sites. These two parts are connected through a stalk which contains a stator. Rotation of the rotor depends on protons passing through the stator. The rotor movement is linked to the binding of the energy-rich compound ATP and the ADP and phosphate.

The design of bacteriorhodopsin and ATPase provides an understanding of how to build an artificial proton pump. The fundamentals are that a proton source would be attached near to the membrane across which proton pumping is to occur. This could be an easily ionisable species or a chromophore of some kind so that the process can be light driven. After protonation, the source transfers the proton to an acceptor further down the membrane, this could be the side chain of an amino acid such as glutamate or histidine or another negatively charged amino acid. This process continues through other acceptors until the proton is captured by a water molecule and carried into the aqueous environment on the other side of the membrane. Once the proton is released then the original source must become reprotonated so that the process can be repeated. This in itself will depend on the design of the pump and the pH difference between both sides.

An example of an artificial energy-producing system using bacteriorhodopsin and ATPase can thus be suggested. Bacteriorhodopsin and ATPases can be added to lipid membranes, where they will spontaneously insert. Bacteriorhodopsin can be easily extracted in the laboratory, as can ATPases. When inserted into membranes the basis of a simple ATP factory can be demonstrated – light is shone onto the membranes, the bacteriorhodopsin pumps protons into the membrane vesicles and the protons then escape back out through the ATPase, generating ATP from the ADP and phosphate that is also added to the system. The light range used in these systems is consistent with space applications. Bacteriorhodopsin absorbs red light between 500 and 630 nm, with a maximum about 570nm, in the visible region. Chemical analogs could also be used that have other light ranges that fit more to the maximum output of the Sun in the yellow region. One of the problems with these systems is that the membranes into which they are inserted are very pH sensitive and therefore not very robust. Artificial membranes constructed from robust artificial materials might provide more mechanically and chemically robust systems for space use.

The use of these ATPases linked to light driven proton pumps in this way offers a means to use light energy in space to produce continuous chemical energy that can be used to power many systems, particularly those that operate in darkness. The chemical energy produced by these systems can, for example, be used in biocatalysis to drive enzyme energy requirements. At maximum light saturation, these artificial systems can produce 0.1 to 0.8 nmol of ATP $\mathrm{min}^{-1}\mathrm{mg}^{-1}$ of protein. ATP gathered in this way could be used to drive chemical reactions in life support systems.

One of the most promising applications of these systems is in batteries. Oxidation and reduction reactions driven by different enzymes could form the basis of small batteries. Light driven ATP production is used to recharge these batteries. These small devices could be used to drive nanobots and other small molecular devices on spacecraft. The mechanics of these devices are in the early stages.

Light driven ATP production can be used to drive nano-motors. The ATPase itself is a tiny motor. The F_1 motor generates a torque during the hydrolysis of ATP that is nearly 100% efficient. The force generated is 100pN and when attached to a solid surface, such as copper, the tip of the motor rotates at about 3-4 revs/sec. Already, these F_1 motors have had tiny nanopropellers attached to them to generate nanodevices

[Soong et al. 2000]. When coupled with light harvesting devices, these tiny motors could be used to drive molecular devices in space stations and other orbital facilities that would move throughout the crew space examining critical life support systems.

Photosynthesis provides one of the most useful methods of electron transfer using the abundantly available light from the Sun as the power for the redox reactions to generate chemical energy. However, electron transfer reactions have many other uses in bionics and it is important to review some of the general processes by which proteins can be used in biomimetic systems to generate redox reactions that could be used for many devices in space.

As we have seen, electron transfer is key to the reactions of photosynthesis, but it is also key to many other reactions of interest including food processing, respiration, drug delivery and other power-related functions. The proteins necessary for electron transfer are not simple predictable entities. For example, the cytochromes, which are proteins commonly found involved in natural redox reactions, can have redox potentials that vary from 800 mV to -400 mV, showing how the protein matrix can have quite dramatic influences on the redox potentials that can be achieved. Engineering systems for artificial use have either involved attempting to engineer existing proteins and improving them or beginning from first principles and designing new proteins.

The design of completed new proteins to drive redox reactions must begin with protein components that can be readily artificially created. Alpha-helical bundles are easier to design than beta-structures. The latter have highly complex interchain interactions and have proven difficult to replicate. Synthetic proteins have allowed the interactions between the haem groups – involved in the actual electron transfer, and the proteins to be studied. A four-helix bundle with a 31-amino acid helix was used to show that the iron-porphyrin moieties that donate electrons or withdraw them have a large influence on the redox potential of the protein. pH was also shown to be critical to the function [Wise et al. 2002].

One way to engineer new proteins is by combinatorial synthesis, whereby mutations are randomly created within specific regions of the protein to generate a battery or library of mutants which can then be tested for physiological properties. Using this technique the electron transfer protein cytochrome b_{562} was engineering to create a range of mutants that had diverse redox properties ranging across 105 mV (Springs et al., 2000). Often these proteins have less useful properties than their wild type originals, so this method is not necessarily useful for throwing up improved proteins, but it can help in elucidating the structure-function of proteins and methods to improve the understanding of which residues to try and improve.

A more directed approach to protein design is to directly engineer proteins that are designed to link together into redox chains. For example, disulphide bonds can be engineered between domains of proteins that would not normally be linked in nature, thus bringing two components of an electron transport chain closer together. In this way, electrons can be channelled through pre-conceived systems in predictable ways. In many ways nature already does this. For example, the electron transfer reactions of photosynthesis are exquisitely designed to channel electrons through particular proteins in a well-defined sequential order. However, engineered electron transfer reactions can be designed to do specific tasks and the engineering of the genes encoding these proteins allows them to be artificially linked. Recently, a genetically engineered link was created between cytochrome c553 and flavodoxin from the bacterium *Desulfovibrio vulgaris* using a disulphide bridge [Sadeghi et al. 2000]. The bridge was used, not to create a new energy system, but to study the how these proteins interact in nature. Nevertheless, the work provides an insight into how genetic manipulation could be used to produce designer electron transfer systems.

As with artificial photosynthesis, one major challenge in the production of artificial electron transfer systems is attaching them to surfaces. Gold and similar substrates have been used and remarkably, these surfaces may give many of the proteins more stability than in their natural state, enhancing their tolerance of extreme pHs and otherwise deleterious solvents. The enhanced stability of proteins in this way may be important when they are applied in new devices where the environmental conditions may be more extreme than found inside a biological cell. The conditions in space are just one obvious example relevant to this study.

A particularly intriguing potential application of redox reactions in space is their use in biofuel cells. Redox couples can be used to turn chemical energy into electrical energy and might be coupled to the photosynthesis/ATP generating reactions discussed earlier. One approach to the biofuel cell is to oxidise organic material with the release of energy. The organic material is biocatalytically oxidised at two compartmentalised electrode interfaces using artificial redox compounds such as those described above. Such a fuel cell has already been proposed, using glucose as the fuel, hydrogen peroxide as the oxidising agent, an anode comprised of glucose oxidase and a cathode functionalised with microperoxidase-11. This fuel cell generated a power of 32µW [Willner and Katz 2000].

This review shows that the field of artificial power generation is an emerging and complex one. The reactions of photosynthesis and the basis of redox reactions are complicated and involved. Yet despite the complexity, principles are beginning to emerge that make it possible to begin to design artificial systems of photosynthesis, proton transport and redox reactions. When combined, it will be possible to gather sunlight and use it to power a diversity of devices in space stations and other platforms. Requirements as fundamental as the production of oxygen might even be driven by artificial photosynthesis and the production of fuel might depend on it. It is likely that the research conducted to date is just a small contribution to the total knowledge that will eventually emerge on this subject and which will make it possible to harness nature's energy capturing systems for space exploration.

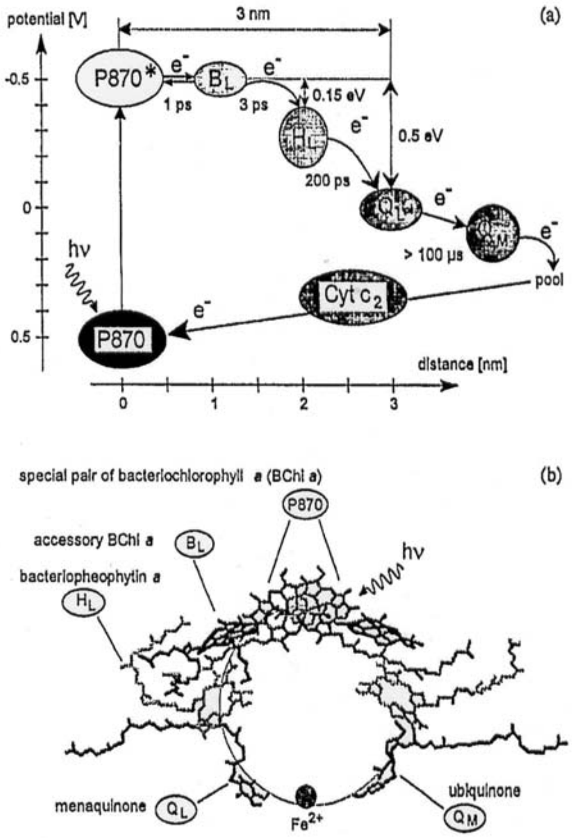

Figure 79 – Natural photosynthetic reaction centre as a typical bioelectronic device: a) energy of electron transport whereby photo-excited electron is removed from P870 vis bacteriophyll B_L and bacteriopheophytin H_L to menaquinone Q_L and transferred to ubiquinone Q_M; b) Optimum spatial arrangement of components of reaction centre (Gopel 1996)

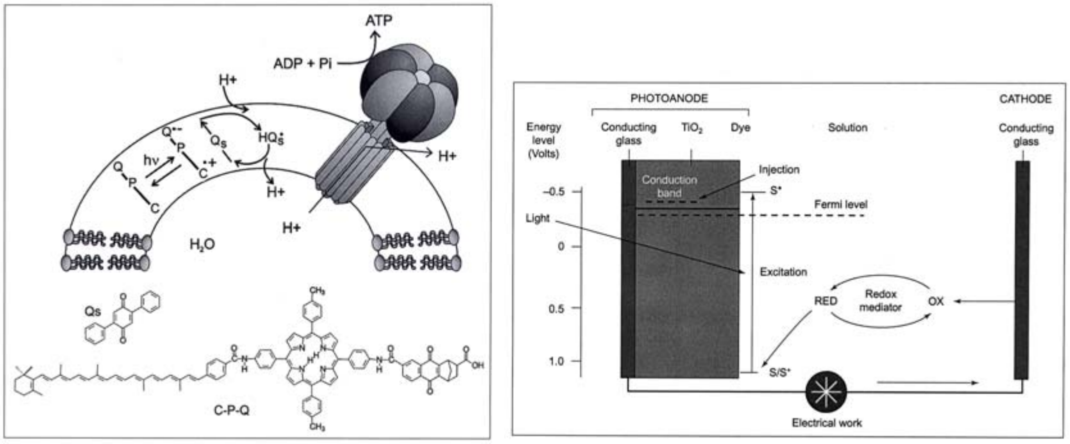

Figure 80 – a) carotenoid-porphyrin-quinone-triad RC mimic and its interaction with a liposome to produce light-induced phosphorylation via the transmembrane proton gradient; b) Gratzel solar cell in which light is absorbed by photoanode and electrons injected into conduction band with circuit completed by electron flow from cathod via redox mediatior to oxidised photosensitive dye (Cogdell & Lindsay 1998).

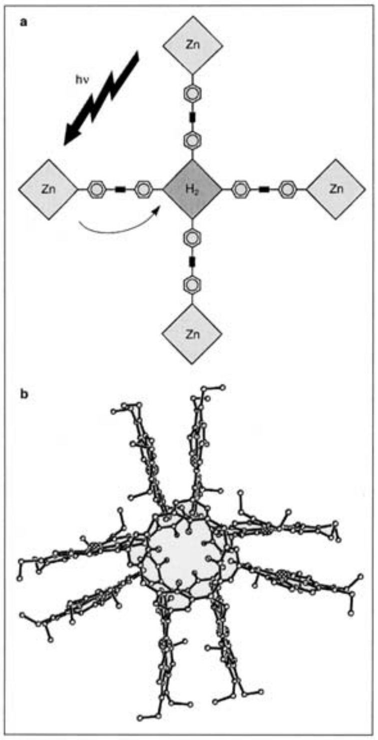

Figure 81 – a) antenna-mimic molecule based on covalent network of porphyrins where energy is transferred from Zn porphyrins to central free-base porphyrin; b) antenna mimic based on purple bacteria LH2 structure such as calyx^3 arenoporphyrin which dimerises similar to bacteriochlorophylls (Cogdell & Lindsay 1998).

3.6.1. Biological Fuel Cells

Fuel cells are an established method of power generation using "cold combustion" of a fuel (typically hydrogen, methane or methanol); the oxidiser is usually oxygen, but other substances such as sodium ferricyanide solution are often used for the biological type. They can generate electricity without significant temperature rise in a similar method to batteries, avoiding the efficiency penalties of heat engines arising from the Laws of Thermodynamics – efficiencies of up to 80% are possible in practical applications. They were invented as long ago as 1843, but had little practical use until the Apollo space missions; they also currently power the Space Shuttle [Smithsonian Institution 2001]. Since they are non-polluting (their exhaust is water vapour, which can be harvested for astronauts to drink in spacecraft applications) and have the potential to generate large quantities of power they are also gaining in popularity for terrestrial products such as electric cars and mobile phones.

The analogy with the processes of biological respiration led researchers in the 1960s such as Yahiro et al [1964] to attempt to use enzymes from the metabolic pathways of respiration for electrical power generation, leading to the development of so-called "enzymatic" fuel cells.

For some time, anaerobically-respiring micro-organisms had been used to convert waste organic matter into flammable gases (principally hydrogen and methane), which could be burnt to produce power; Suzuki et al. (1979) showed that the gases could also be fed to a conventional fuel cell to generate electricity, though this approach used well-established technology and will not be discussed further. Subsequently synthetic chemical mediators were used to transfer energy from the microbes to the circuit in true "microbial" fuel cells [Bennetto et al. 1983]. Recent research [Park et al. 2001; Chaudhuri & Lovley 2003] has shown that some species of bacteria can produce electricity directly without needing a mediator. A more comprehensive review of past developments in this field is given in Katz (2001).

Table 12 summarises the advantages and disadvantages of biological fuel cells (also known as biofuel cells) compared to conventional technology, with highlighting to show where either biological or conventional fuel cells have the advantage over the other type.

Table 12 – Comparison of the Properties of Conventional and Biological Fuel Cells

Property	Conventional fuel cells	Biological fuel cells
Operating temperature	High (>90°C)	Ambient (20°C) – but can be damaged by high temperatures
Operating pressure	Atmospheric to high (>140KPa)	Atmospheric
PH	Acid or alkaline	Neutral – but can be damaged by extremes of pH
Fuel	Restricted number of gaseous/liquid fuels – hydrogen, methane, methanol	Wide range of soluble/liquid fuels – sugars, organic alcohols and acids, biological material, wastewater
Power output*	High (kW-MW)	Low (<0.1W)
Electrode material	Nickel, graphite, perovskite ceramics	Graphite, graphite felt, reticulated vitreous carbon, platinum
Electrolyte material	Liquid or solid	Liquid
Ability to regenerate fuel	Yes (if hydrogen-powered)	No

*The output of a fuel cell stack depends on the number and size of fuel cells used, as well as on their technology – this row gives the scale of output of integrated fuel cell power generation systems that have been developed to date.

It is worth noting that little work has been carried out to optimise biological fuel cells for high power output or low weight – with research this could probably be improved considerably. Also, genetic engineering could make use of extremophilic organisms (those that inhabit inhospitable environments), either whole or reduced to single enzymes as appropriate, to improve the fuel cells' tolerance of high temperatures, high or low pH values or ionising radiation (e.g. *Deinococcus radiodurans* for irradiation survival [Anderson et al 1956]).

Selected data for biological and conventional fuel cells are given in Table 13.

Table 13 – Power Output of Selected Biological Fuel Cells and Two Conventional Fuel Cells

Source	**Date**	**Fuel cell type**	**Fuel**	**Output voltage (V)**	**Output current (mA)**	**Power (W)**	**Power area density (mW/cm^2)**	**Reference**
University of Texas at Austin	2004	Enzymatic	Glucose solution	0.52	0.0083	4.3x10^{-6}	0.54	Heller (2004)
Helsinki University of Technology*	2000	Enzymatic	10% methanol solution	0.5	160	0.081	1.8	Halme et al. (2000)
Helsinki University of Technology	2000	Mediated microbial	Glucose, fish meat	0.5	26	0.013	0.11	
University of the West of England	2002	Mediated microbial	Sucrose solution	2.4	7.5	0.018	0.35	Melhuish (2002)
Ghent University	2003	Mediator-less microbial	Glucose solution	-	-	0.018	0.36	Rabaey et al. (2003)
Space shuttle	1988	Alkaline	Hydrogen gas	0.86	150000	125	92	NASA (1988)
NASA JPL	1997	Direct methanol	3% methanol solution	0.5	50000	25	161	Halpert (1997)

*Predicted values

The main types of biological fuel cells are described below.

3.6.1.1. Enzymatic Fuel Cells (EFCs)

EFCs are very similar in principle to conventional fuel cells, consisting of two electrochemical half-cells separated by a semi-permeable proton exchange membrane (PEM) which allows the passage of hydrogen ions only. The fuel is oxidised at the anode, producing electrons which travel around the circuit (performing work) before they are consumed by reduction reactions at the cathode. Hydrogen ions are also produced during the oxidation of fuel – they diffuse through the PEM to the other half-cell where they react at the cathode with the oxidiser and the electrons. The difference between conventional and biological fuel cells lies in the substances that catalyse the chemical reactions, which in the former type are usually transition metal alloys (such as platinum-rhodium or nickel-tin), whereas EFCs use enzymes – protein-based biological catalysts – usually in the anode compartment but sometimes also in the cathode compartment [Laane et al. 1984; Palmore & Kim 1999]. Since the enzymes cannot transfer electrons directly to the electrodes, a synthetic chemical "mediator" is often added to the electrolyte to ferry the energy released from the fuel to the electrodes through redox reactions. The operation of an EFC is shown in Figure 82.

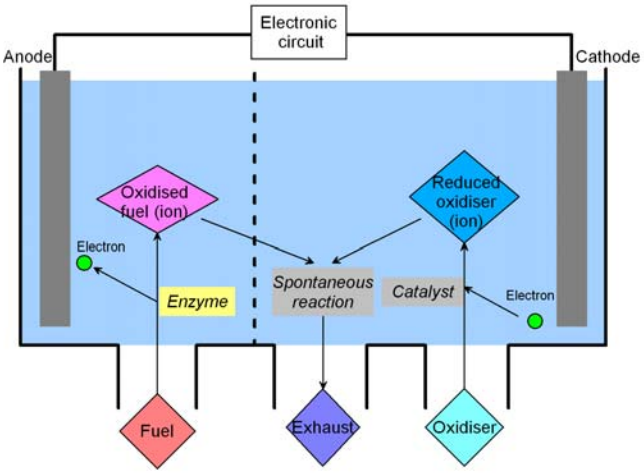

Figure 82 – Diagram of the operation of a typical enzymatic fuel cell

The enzymes can be present either as a component of the electrolyte mixture or attached to the electrodes ("immobilised" enzymes). Such immobilisation, which can give better results, has been accomplished by applying solutions of the enzyme to the electrode and drying them, by immersing an electrochemically pre-treated electrode in a strong solution of the enzyme [Dimcheva et al. 2002], by using a porous polymer membrane [Minteer et al. 2003] or by chemical means such as "wiring" the enzyme to the electrode with a redox hydrogel polymer [Heller 2004]. EFCs using chemically immobilised enzymes can usually avoid the need for separate mediators.

EFCs share the advantages of generalised biological fuel cells shown in Table 12, with one specific addition. Enzymes are such effective catalysts that the fuel (e.g. glucose) does not react spontaneously with dissolved oxygen at the operating conditions – unlike conventional fuel cells where the hydrogen and oxygen would combust explosively if allowed to mix at the operating temperature, consuming fuel without generating power and possibly damaging the device. Moreover, if enzymes are used on both electrodes then (thanks to their specificity for their substrates) it is not possible for unwanted chemical reactions to occur. Thus there is no need for the PEM that is usually required to separate the two half-cells, simplifying the construction of such fuel cells to the extent that it is possible to power a physiological monitor in a living organism by simply implanting two electrodes into a blood vessel [Katz et al. 1999; Heller 2004].

3.6.1.2. Microbial Fuel Cells (MFCs)

In contrast to the EFC described above, the catalyst for the oxidation of fuel at the anode of an MFC is a culture of living microorganisms – either bacteria or yeast. The processes of respiration within these organisms (as described above in Section 3.6.1.1) typically transfer the energy from the oxidation of food to the compound NAD (nicotinamide dinucleotide) to produce its reduced form, denoted NADH. This NADH represents the reservoir of energy that can be tapped to produce electricity. A wide variety of organic fuels can be used for MFCs – to date, sugars (glucose, sucrose, fructose, xylose [Chaudhuri & Lovley 2003]), organic acids (lactate, acetate [Bond & Lovley 2003]) and biological tissue (fish meat [Halme et al. 2000]) have all been used in various experiments with some success.

The different ways in which the transfer of energy is accomplished allow the further classification of MFCs into two types.

First, in "mediated" MFCs, as in mediated EFCs, a synthetic chemical compound is added to the culture whose purpose is to ferry the stored energy of the NADH from inside the microbes to the anode – the operation of a mediated MFC is shown in Figure 83 Combinations of microbe and mediator studied include baker's yeast (*Saccharomyces cerevisiae*) with 2-hydroxy-1,4-naphtho-quinone [Zhang & Halme 1994] and the bacterium *Escherichia coli* with thionine [Bennetto et al. 1983]. Wilkinson (2000) gave an extensive review of the microorganisms that had been used in mediated MFCs, and suggested the following criteria for mediators:

- Penetrates the microorganism cells when oxidized.
- Readily reacts with the electron source.
- Escapes rapidly from the cell when reduced.
- Is electrochemically active at the anode surface.
- Is stable over long time periods.
- Is soluble in the anolyte medium.
- Is non-toxic to the microbes.
- Is not metabolized by the microbes.
- Has a similar redox potential to the reaction that is providing the reducing power.

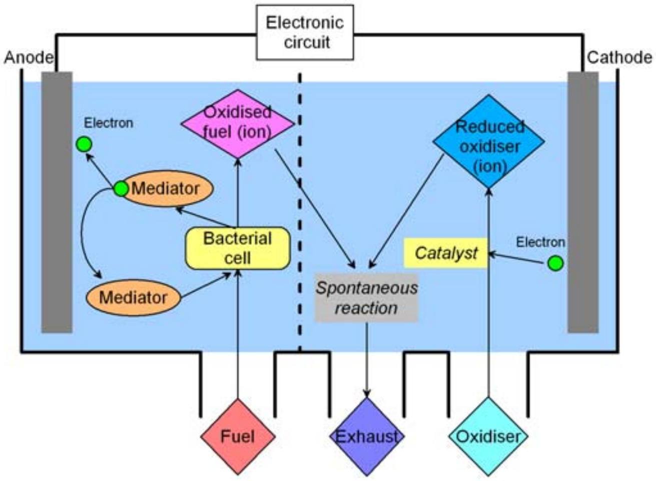

Figure 83 – Diagram of the operation of a typical mediated microbial fuel cell

In contrast, some forms of bacteria are able to transfer electrons directly to the anode without the need for any intermediate stage, and so can be used in "mediator-less" MFCs whose operation is shown in Figure 84. A significant proportion of such organisms, including Rhodoferax ferrireducens [Chaudhuri & Lovley 2003] and the Geobacteraceae [Bond & Lovley 2003], evolved to use extracellular ferrous oxide (rust) as their electron acceptor instead of oxygen; it is conjectured that this explains their ability to adapt readily to using graphite electrodes as an electron acceptor instead. Bond et al. (2002) have demonstrated that mediator-less MFCs can be used to produce power from the organic matter in marine sediments using the bacteria native to the substrate (largely Geobacteraceae), while Liu and Logan (2004) have shown that the same can be accomplished with sewage treatment plant wastewater. In a similar vein, Professor Bruce Rittman at the Environmental Biotechnology group of Northwestern University, Illinois, is working on a NASA-funded project to develop a wastewater-powered MFC with a novel structure that could be used for spacecraft: it uses fibres with concentric layers for the anode, PEM and cathode of the fuel cell, immersed in a vessel containing the bacterial culture. However, no work has yet been published.

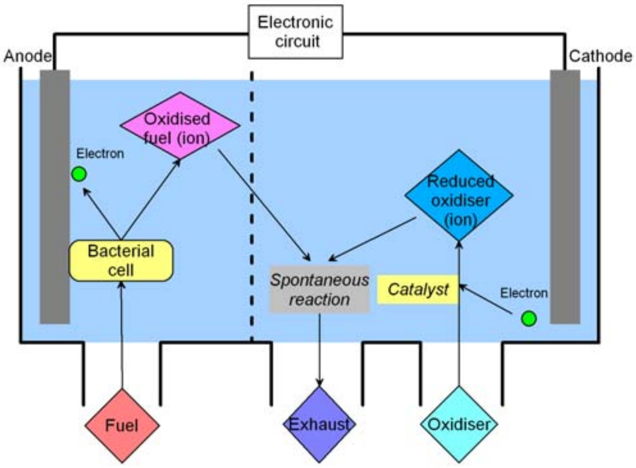

Figure 84 – Diagram of the operation of a typical mediator-less microbial fuel cell

A consideration in the design of MFCs is the fact that many of the microorganisms used in both types of MFCs need to be kept in conditions where there is no oxygen (anoxic). This may be either because they are poisoned by oxygen (such organisms are termed “obligate anaerobes”) or because they use oxygen as an electron acceptor if available in preference to the mediator or electrodes (“facultative anaerobes”).

3.6.1.3. Photosynthetic microbial fuel cells (PMFCs)

Although heterotrophic microorganisms (those that consume an external food supply to produce energy) are most frequently used in MFCs, autotrophic microorganisms (which produce their own food from light by photosynthesis, described previously) have also been investigated for electricity generation [Yagashita et al. 1997]. In most cases the PMFCs have been very similar to mediated MFCs in construction, with a sodium ferricyanide cathode and a synthetic chemical mediator at the anode. The mediator transfers the chemical energy from the products of the photolysis of water to the electrode, and thus is subject to the same criteria as a mediator in a conventional mediated MFC [Wilkinson 2000].

Popular experimental subjects have been *Synechococcus* species with mediators such as 2-hydroxy-1,4-naphthoquinone [Yagishita et al., 1997], 2,6-dimethyl-1,4-benzoquinone (DMBQ) or diaminodurene (DAD) [Tsujimura et al., 2001], though others such as *Anabaena variabilis* and *Synechocystis* have also been studied [Yagishita et al., 1996]. In some experiments the usual sodium ferricyanide cathode was replaced by an enzymatic cathode using bilirubin oxidase to catalyse the reduction of oxygen to water [Tsujimura et al. 2001]. A typical power area density observed was 1.9W/m², with a light conversion efficiency of 3.3% [Yagishita et al. 1997].

3.6.1.4. Comparison of the different types of biological fuel cells

Despite the fact that some of the fuel in MFCs is used for the microbes' survival and growth as well as the generation of electricity, they have the potential to be more efficient than EFCs because a larger fraction of the chemical energy stored in the fuel is available to power-producing reactions (in EFCs, oxidation of the fuel is only partial and a significant amount of energy is discarded in the fuel cell's waste products). In addition, the variety of types of fuel that an MFC can use is far higher, meaning that it can be adapted to use whatever fuel is available or even the waste from another process. Since the microbes are alive, the catalyst is self-regenerating making the system more robust to damage. Some bacteria could be used for secondary functions as well – for example, the oxide-reducing ability of Rhodoferax ferrireducens could be exploited for metal ore processing (though not simultaneously with the production of power).

Mediator-less MFCs are less complex than mediated versions as there is no need to maintain a concentration of mediator in the microbial culture. They can also be more efficient since no energy is lost in the mediator stages, as long as a microorganism can be found that transfers a sufficient proportion of the electrons liberated from the fuel to the anode. Power output is observed to be governed by the volume of the anode compartment in mediated MFCs [Halme et al. 2000] but by the surface area of the anode in mediator-less MFCs [Chaudhuri & Lovley 2003], due to the fact that only the bacteria attached to the electrode surface produce power in the latter case.

Since PMFCs do not consume pre-existing fuel but rather utilise (sun)light to produce electricity, they are not strictly comparable to the other forms of biological fuel cells. They share many of the advantages and disadvantages of mediated MFCs, though the need to collect light means that they cannot be fabricated in a compact stack like the other types. However, if the system containing them can be made sufficiently robust they can function indefinitely, though at present their light conversion efficiency is below that of conventional photovoltaic cells.

3.6.1.5. Applications

To date, little use has been made of biological fuel cells, apart from laboratory and classroom demonstrations of electricity generation, due to their low power generation capabilities at present; nonetheless, PowerZyme Inc is developing commercial enzymatic fuel cells for consumer electronics [Schuilli 2002]. Another area of interest has been the use of micro-scale biological fuel cells, fuelled by glucose from the bloodstream, to power implanted medical devices – one MFC prototype uses the yeast *Saccharomyces cerevisiae* on a micromachined silicon chip and is 0.07cm² in area [Chiao et al 2003]. An even smaller EFC version consists of two 7µm diameter and 2mm long carbon fibres coated with immobilised enzymes, producing 1-10µW [Heller 2003].

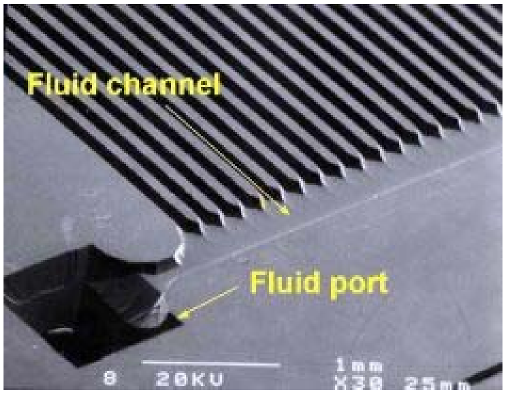

Figure 85 – The yeast-powered MFC constructed by Liwei Lin

At least two examples of macroscopic robots powered by biological fuel cells have also been constructed. The first, "Gastronome" at the University of South Florida [Wilkinson 2000], uses an onboard stack of six mediated MFCs to charge a NiCd battery pack, which then powers a motor and gearbox to drive the robot forwards. The organism used is *E. coli* feeding on sugar; reduced mediator from a separate "stomach" chamber containing the bacteria is pumped to the anode chamber of the fuel cell which uses a sodium ferricyanide cathode. The robot alternates between 18 hours charging the batteries and 15 minutes of forward motion.

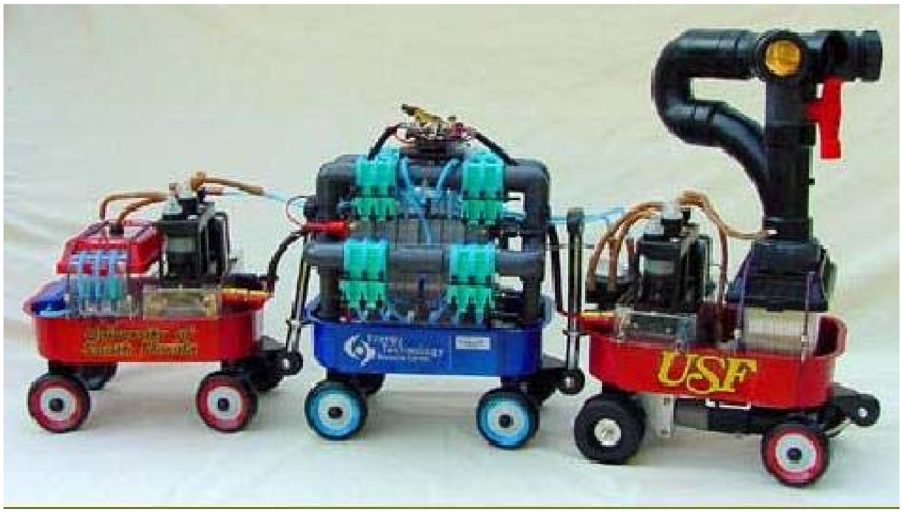

Figure 86 – Gastronome Fuel Cell Powered Test Vehicle

The second, "EcoBot" at the University of the West of England [Melhuish 2002], uses eight mediated MFCs with the same combination of anode and cathode reactions as Gastronome, carried by the two-wheeled robot. In this case, though, the bacterial cells live in the anode chamber of the cell itself and so no energy-requiring pumping is needed. The MFCs charge capacitors that supply power to the robot's motors and sensors, allowing it to perform a simple task (following a light source) – 30 seconds spent charging the capacitors provide enough energy for 2 seconds of motion.

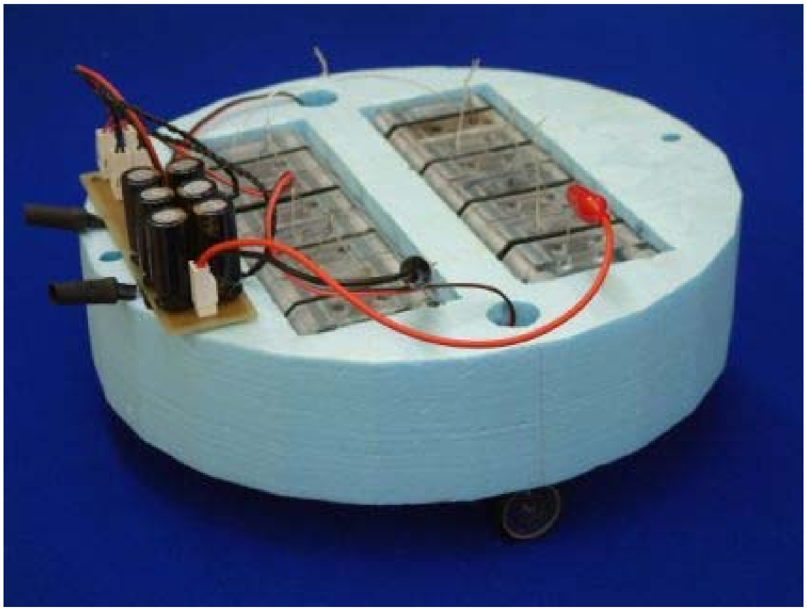

Figure 87 – EcoBot Fuel Cell Powered Micro Robot

3.7. IMMUNOLOGICAL APPROACHES TO SELF-DEFENCE

What is unique about biological systems that cannot yet be replicated in engineered systems is the ability to self-repair. The capacity for regeneration of specialised tissue depends on the ability of its cells to undergo mitosis. In tissues such as skin, muscle, cartilage, tendon and epithelial surfaces, active regeneration occurs by growth inwards from the edges of the wound. Endocrine and liver glands, although they cease to multiply in adults, have the potential for mitosis initiated by stimuli. The central nervous system has no regenerative capability after birth. Peripheral nerve cells cannot multiply but their axons may regenerate as long as they entire the appropriate Schwann cells peripherally. However, self-repair of macro-scale wounds is a general property of living systems. The edges of the gap within a wound fills with blood clot (formed by platelets) and serum in which the fibrin (derived from the coagulant factor fibrinogen) forms a matrix in which fibroblasts and angioblasts grow from the wound edge. Eventually, granulation tissue is formed between the wound edges and the fibroblasts contract into collagen (scar tissue) by the action of platelets on fibrin. Granulation tissue is comprised of fibroblasts and angioblasts infiltrated inflammatory cells. Inflammation is the reaction of tissue to injury. Acute inflammation involves initial dilation of blood vessels due to the release of histamine, 5-hydroxytryptamine and other vasodilators. Serum, white blood cells and antibodies pass through the capillaries into the tissue. Chronic inflammation involves the proliferation of fibroblasts and fibrocytes, macrophages and plasma cells.

The immune system consists of a large number of cells of several types dispersed through the body, which collaborate to protect an organism from a variety of threats both external (disease-causing agents such as parasites, bacteria, viruses and toxins) and internal (mutant cells that might become cancerous), collectively termed “pathogens”. This takes place in two stages: first, recognition of substances that are foreign and therefore could cause harm (termed “non-self”, as opposed to the “self” substances that make up the organism), followed by their destruction. Computer programs have been written using software analogues of the recognition mechanism, termed "artificial immune systems" or AISs. They have been used for a variety of applications, ranging from controlling robots to securing computer systems against viruses. The threat destruction mechanism has not yet been imitated since more suitable alternative methods already exist.

3.7.1. Computer Viruses

The increasing connectivity between computers represents a potential threat through the propagation of computer viruses [Cohen 1987; Denning 1988]. The distribution of control between the ground and the space segment through the communications link could potentially yield viral infection of the spacecraft onboard computer. This is particularly the case when uploading software from the ground to the spacecraft. Host-based and network security such as anti-viral software and firewall security systems are a first line of defence. Network security provides cryptographic and user authentication procedures. There are 4 types of programs that can attack other programs in a computer memory. A worm is a program that performs a legitimate function but also contains hidden code which performs an illegitimate one. A bacterium is a program that replicates itself and uses up processor and memory capacity from the host. A virus is a program that incorporates copies of itself into the code of other legitimate programs usually at the start and takes control performing illegitimate functions when invoked. There are several types of viruses, some 89% of which are parasitic in that they modify the content of executable files. File execution is diverted to execute viral code first prior to normal execution. A shell virus forms a shell around the original program which then becomes a subroutine of the virus. A carrier virus appends its code to the host program. It can then locate other uninfected programs and infect them taking up computational resources. Viruses are usually introduced as Trojan horses which appear innocuous and can comprise of only a few lines of code. The ability of viruses to replicate and infiltrate programs as Trojan horses coupled with the widespread use of sharing and networking in computer systems makes them a considerable threat. Around 85% of known viruses are file infectors, 5% are boot sector viruses which load on booting up and 10% are macroviruses which run as scripts in applications. The 10 most common viruses comprise around 65% of all viral attacks. The activity monitor alerts to activity usually associated with viral activity but not characteristic of legitimate programs. They cannot however locate the source of the infection. Generic virus detection programs can detect unknown viruses but are prone to false alarms. Although it is critical to have low false positives and low false negatives, tolerance to false negatives should be more desirable than tolerance to false positives. Virus scanners search files and memory where executable code can be stored. Scanning programs look for specific viral signatures of known viruses ~16-32 bytes in length (which identify them and mark them from legitimate code) to locate them utilising pattern recognition algorithms. The assembler code is analysed to determine the viral behaviour and its attachment method. A 20 byte signature often forms a family signature which is common to several distinct viruses of the same family. This pattern matching process is similar to biological viral receptors which bind to 8-15 amino acid sequences in viral protein coats. Once detected, the viral code is deleted and the infected file repaired. The difficulty is that computer security is an NP complete problem and is undecidable by both a priori and runtime analysis. To determine if a program P is a virus, it is necessary to determine if it infects other programs and any detection mechanism can be infected. This is undecidable. An evolutionary virus may be constructed by allowing it to add random statements between each line of necessary statements and this evolutionary process is also undecidable. A compiler is in fact a virus as it compiles a new version of itself and so every use of the computer is a potential viral attack. Once a virus is detected it must be removed. If the system is kept running during removal a disinfected program may be reinfected unless the program performing the removal is faster at spreading than the virus to be removed. This usually implies that the system must be shut down or at least a partial denial of services. The use of virtual memory which restricts a program to a limited region of memory will only slow viral propagation rather than stop it. If the virus is a known one, an antibody program can remove it. Such an antibody program resides in the operating system until a virus triggers its activation. The only way to detect an unknown virus is to create an authenticator when the program is installed which checks the sum of bits in the file to ensure that there is no additional unwanted code. Integrity monitors of the authenticator use checksums to detect changes to programs and files and differences are flagged. Activity monitors of the authenticator examine the nature of such modifications using a variety of heuristics. The authenticator is opened with a secret key and contains the proper checksum value. The only infallible way to protect against viral attack is to limit networking and sharing. An immune system is under development for computer systems based on the vertebrate immune system [Kephart 1994]. This system can develop antibodies to previously unkown computer viruses and subsequently remember and respond to them in the future. Computer viruses utilise the host's computational energy to "metabolise". Present viruses do not evolve as only programmer intervention can generate mutations. Such evolving computer programs would have to be large and complex. Such an immune system must recognise known intruders, determine unknown intruders, eliminate those intruders, and selectively self-replicate solutions for effective recognition and response. The immune system provides a fuzzy match to the virus. This fuzzy match to a short sequence of bytes in the virus that comprise its signature must be efficient in time and retain the ability to recognise variants and mutations of the virus. Clonal selection theory of immune system function suggests that the vertebrate body generates a vast range of immune cells and then selects and amplifies the population of those that match pathogens. In computational systems, pathogenic signatures may be generated randomly, and each must be checked against existing signatures. A major difficulty concerns distinguishing harmful code of unknown viruses and legitimate software whose de-activation or erasure could be catastrophic. Viruses tend to seek out programs that are most executed by the operating system to maximise their spread. Decoy programs may be presented to attract viruses by their generalised manipulation by the operating system and thereby be localised and detected. Infected decoys are then analysed by the signature extraction program to develop recognition information for the new virus. The attachment pattern of the virus and the constant executable code sequences across all variants in the system can then be extracted. This information is then held in a database for re-activation if further infection occurs. To limit the spread across networks, one infected machine signals the database information to all connected neighbours in that network. Computer viruses are less rife than formerly feared. Less than 15% of 1500 known viruses have been observed, most only once. The top 10 viruses account for 2/3 of all infections. This is due to localised software and diskette sharing [Kephart et al. 1993].

3.7.2. Key Principles of Immunity

The human immune system is a complicated network of interacting cells and chemical messengers (hormones). It is able not only to produce a rapid and generalised response to an invading pathogen (the "innate response") but also to adapt specific counter-measures to combat it over a longer period of time. For the purposes of constructing AISs it has not proved necessary to mimic a natural immune system in every detail to achieve results; instead, certain key principles have been abstracted from the lifecycle of a particular kind of immune system cell known as a B-lymphocyte or B-cell and applied to computer algorithms, as shown in Figure 88.

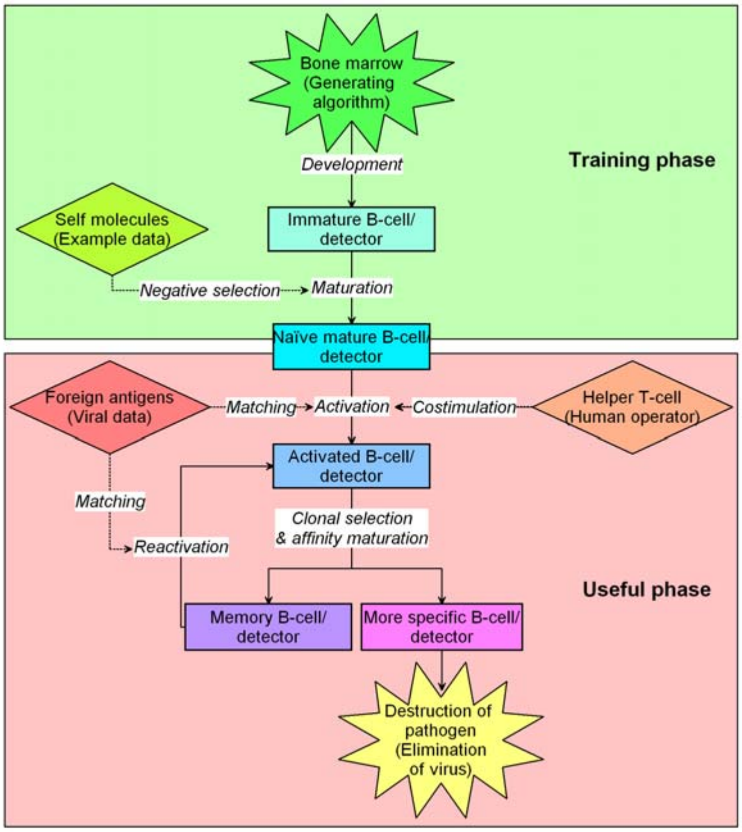

Figure 88 – A diagram showing the stages in the life of a B-cell and their AIS equivalents. The B-cell is produced and undergoes maturation in the training phase, and then acts to protect the organism during the useful phase.

These principles are discussed in the sections below in connection with the protection of data from computer viruses, but many other applications are possible – see section 3.7.4 below for details of some of them. The principles are arranged roughly in order of decreasing importance; only the first two principles, "Antigen-antibody matching" and "Antibody generation and negative selection" (described in sections 3.7.2.1 and 3.7.2.2), are necessary to create a functioning AIS. Nonetheless, implementing some or all of the other principles is likely to enhance a system's capabilities.

Note that the description of the immune system given in this document is greatly simplified. For a more detailed introduction to immunology, see Roitt and Delves (2001). Biological organisms must monitor and protect their internal systems to support metabolism, e.g. immunity. The immune system defends vertebrates against infection. The immune system is comprised of cells which circulate through the lymphatic system. Some immune cells are non-specific – such cells are either phagocytic, enzyme-producing, or chemical producing (eg. histamines). Phagocytes do not synthesise antibodies but concentrate the antigen for adjacent lymphocytes. Others are lymphocytes which manufacture antigen-specific antibodies to de-activate foreign antigens. Lymphocytes circulate in bone marrow, lymph nodes, spleen and lymphatic vessels. The immune response is initiated by antigens which may be complex proteins or large carbohydrates. Injury to tissue causes inflammation – blood accumulates in the infected area impregnated with leukocytes and fibrinogens (white blood cells). There are ~10^{10} lymphocytes in the human adult blood with ~10^6 different clones. Antigens are lymphocytes that recognise foreign tissue and are large protein molecules. Adaptive immune system responses are initiated by interactions of the antigens with specific receptors on the surfaces of lymphoid cells. Each lymphocyte clone shares a unique cell surface receptor that enables them to bind to particular "antigenic determinants" comprising of a portion of a molecule. All lymphocytes develop from pluripotent homopoietic stem cells. They circulate continuously between blood and lymph. The most basic requirement of the immune system is to distinguish between "self" cells and "non-self" foreign cells. The ability of the immune system to distinguish between self and non-self depends on specific paralysis to anti-antigens during foetal development. Anti-antibodies may occur if auto-antigen structure changes, loss of T-suppressor activity, or circulation of antigens shielded from the immune system. Antigen-antibody complexes in circulation may also form micropreciptiate in small blood vessels and cause leucocyte accumulation and inflammation. The second requirement is to eliminate non-self invaders such as bacteria. In addition, the immune system should recognise and eliminate altered "self" tissue such as cancer. The immune system recognises particular antigens through antibodies and immune cell receptors which bind to active portions of antigens (typically 4-6 amino acids). An antibody is a Y-shaped protein with two identical antigen-binding sites at the tip of the Y and complementary binding sites on its Fc region. If an antibody meets an antigen, the two bind together thereby neutralising the antigen. The strength of antigen-antibody interaction depends on the affinity and number of binding sites. Antibodies de-activate viruses and bacterial toxins and kill and ingest invading microorganisms. Lymphocytes are produced during pre-natal development from stem cells of bone marrow and they have the ability to recognise their own proteins as opposed to foreign bodies. Macrophages are phagocytotic white blood cells which engulf invading microbes and release proteins to activate other parts of the immune system. Cytokines are proteins released by immune cells when activated. They include interferon IL-1 and IL-6 and TNF (tumour necrosis factor). Macrophages produce IL-1 and IL-1 stimulates macrophages to destroy micro-organisms. They comprise the innate/natural immunity system of rapid response. The complementary/acquired immune system works in cascade to identify and destroy invaders. There are two classes of lymphocyte – B cells which make antibodies and T cells which are involved in cell-mediated immunity response. Each B cell clone makes antibody molecules with a unique antigen binding site. T cells undergo transformation and mitosis generating a population of specifically reactive anti-antigen cells. B cells differentiate into plasma cells which secrete antibodies in cooperation with T-helper cells. T suppressor cells interact with B cells to inhibit antibody synthesis. Molecules are inserted into the plasma membrane to act as cell surface receptors for the antigen. When the antigen binds to the antibody in the membrane, B cells are activated to multiply and synthesise quantities of the antibody with the same antigen-binding site which is secreted into the blood. Lymphocytes are specialised white blood cells which multiply when they encounter specific antigen proteins associated with foreign organisms. The antigen does not serve as a template for the synthesis of antibodies of the correct fit from lymphocytes. Clonal selection operates by the production of a vast diversity of lymphocytes which are selected by fit to the antigen. Once selected, those lymphocytes are clonally reproduced and perpetuated as antibodies. The body can produce antibodies against a vast variety of potential antigens, i.e. the ability is genetically determined (clonal selection). Large numbers of small populations of lymphocyte cells exist, each capable of producing certain types of antibody. Foreign antigens react with particular lymphocyte cell population causing them to amplify the manufacture of that antibody. In the thymus gland, the classes of lymphocytes which would react with the body's own chemicals are eliminated. Lymphocytes react to antigens by multiplying and transforming into plasma cells. Plasma cells manufacture antibodies from their substantial cytoplasmic and rough ER reserves. The process of mutation and selective proliferation allows the immune system to learn to fabricate antibodies, B cell and T cell receptors which can recognise particular intruders. The immune system can remember the antigen by retaining immune cells with the specific receptors that recognise specific antigens. Antibody or receptor must not bind to self-proteins so early in lymphocyte development, immature B and T immune cells are subjected to a training phase in the thymus where self-recognising receptors are suppressed and eliminated so that the immune system responds only to foreign antigens. The immune system works by clonal selection – most antigens stimulate many different lymphocyte clones. In primary immune response, immune cells with receptors that match a given antigen well are stimulated to reproduce and proliferate by clonal expansion. Strong selection pressure is applied to good recognition immune cells. Lymphocytes are of two types: B lymphocytes secrete defensive antibody proteins that bind to antigens. There are ~100B specific types of B lymphocyte. Each lymphocyte recognises only one antigen and its is DNA mechanisms which provide the diversity. Antibody immunoglobins are large molecules with a characteristic structure – the Ig fold – which serves to recognise non-self molecules. Humans have 5 major types of immunoglobin. All antibodies on a single B cell are of the same type and bind to a specific antigen. If the B cell encounters an antigen, it binds to it and the B cell reproduces to secrete its antibody. Diversity is relieved by genetic processes of both random and inherited components that produce a unique receptor section which is the binding section of the antibody.

Serum antibodies (immunoglobins Ig) are large proteins of two heavy and two light polypeptide chains. Different classes IgG, IgA, IgM, IgD, IgE are distinguished by their heavy chain structure. Each chain comprising repeating units of ~70 amino acids enclosed by a disulphide bond. The structure has evolved to perform two specific roles. Antibodies combine with a specific antigen. There are around 10^5 different specifities of antibody. Two combining sites (V regions) per four-chain molecule are generated by interaction of N-terminal unit of heavy and light chains which show variability of structure in different antibodies. V regions are common to all immunoglobins Ig. Antibodies also have effector functions which involve specific interaction with bodies through C regions. IgG and IgM have sites for complement fixation. Only IgE has a site for attachment to mast cells. IgG and IgM attach to macrophage surface and induce endocytosis. Only IgG attaches to killer cells to initiate antibody dependent cell-mediated cytotoxicity. Antigen stimulation also leads to the production of small lymphocyte cells with specific combining sites (T cells) originating from bone marrow but modified through the thymus. The receptors are similar to V regions of serum immunoglobins. They comprise around 80% of lymphocytes. T-helper cells cooperate with B cells in Ig synthesis, T-suppressor cells inhibit B cells and T-cytoxic cells can kill virally infected target cells and tumour cells. Immunity is also dependent on the reaction of bound antibodies to cause cell lysis of the pathogen or to facilitate phagocytosis of the pathogen.

Each antibody's antigen receptor is formed through interaction between two amino acid chains, one light L and one heavy H. Basic antibody molecule has two pairs of such polypeptide chains and so two antigen receptor sites. Five different classes of antibody (IgA, IgD, IgE, IgG and IgM) each with distinctive H chain (α, δ, ε, γ and μ) which also forms the Fc region to define the biological properties of the class. There are two types of L chain (K or λ but not both) can be associated with any class of H chain. Each immunoglobin L and H chain comprises of variable regions ~110 amino acids residues at its terminal and followed by constant regions. The constant region is the same length in the L chain but 3-4 times longer in the H chain. Each chain consists of repeating and folded domains. The L chain has one variable region V_L and one constant region C_H while the H chain has one variable region V_L and 3-4 constant regions C_H. The amino acid sequence variation in the variable regions are confined to 3 small hypervariable regions which interact together at one end of the molecule to form antigen binding sites. Each anntigen binding site can contact 5-6 sugar residues of the antigen. More than one gene segment codes for each L and H chain. Antibodies are produced from 3 separate gene pools encoding the K, λ and H chains respectively. In each pool, separate gene segments that code for different parts of the variable regions of L and H chains can be brought together by site-specific recombination events during B cell differentiation. Hence, antobody diversity is increased by somatic recombination, by combinatorial joining of light and heavy chains and by somatic mutation. Specificity of the receptor site depends on amino acid sequence of the chain and their arrangements which are specified by gene segments (antibody genes) in the B cell nucleus. There are 3 types of gene segments: V (variable), D (diverse), J (joining). Amino acids of the heavy chain are specified by all 3 segments while the light chain is specified by V and J only. Fourth gene segments C (constant) determines the class of antibody. In humans, V, D J and C segments all lie spaced out on a single chromosome. Segments are located in clusters of 50V, 30D, 6J and 8C elements provide combinatorial receptor diversity. Two gene segments code for V regions of each L chain while 3 gene segments code for V regions of each H chain. To make an antibody, a V gene segment is combined with a J gene segment to produce a V gene for the L chain and a V_H gene is combined with a D and J_H gene segment to produce a V gene for the H chain. Each of the assembled gene products is then co-transcribed with the appropriate C region sequences to produce mRNA molecule that codes for the complete polypeptide chain. By randomly combining inherited gene segments for V_H and V_L regions, vertebrates can make 10^3 different L chains and 10^3 different H chains that can combine to form 10^6 different antibodies. Somatic mutation in the gene segments coding for V regions can increase that variability by 10^2. All B cells make IgM antibodies. Some later switch to making antibodies of other classes with the same antigen binding site, allowing the same antigen binding site to be distributed across antibodies with different biological properties. There are other sources of diversity. When joining of V, D and J segments, at the joining boundary between two segments, several base pairs are removed locally and new bases added randomly, thereby changing amino acid sequence and antigen receptor. This allows ~10^6 different variants of antibody molecule from each cluster, each slightly different across two locations (between V and D, and D and J segments). Junction diversity can generate vast diversity of specificity but it is rapid selection of specific antibody by B cells that allows expansion of antibodies. B cells possess immunological memory to reduce delay in combating subsequent infection. T lymphocytes that recognise and kills cells that bear non-self antigen proteins on their surfaces, and enable B lymphocytes to secrete proteins. T cells carry specialised receptor molecules for recognition and it is T cells that inhibit foreign skin grafts. T cell production follows a similar pattern as both T cell receptors and antigens are specified by similar gene mechanisms. Antibodies and T cell receptors are the primary means for immune response to antigens. Killer T cells may encounter infected host cell and kills it. In doing so, the killer T cell hastens the inevitable demise of the cell but prevents a virus from completing the replication process to infect other cells. T cells also help regulate B cell formation. Host defence systems must have arisen at the same time as life itself.

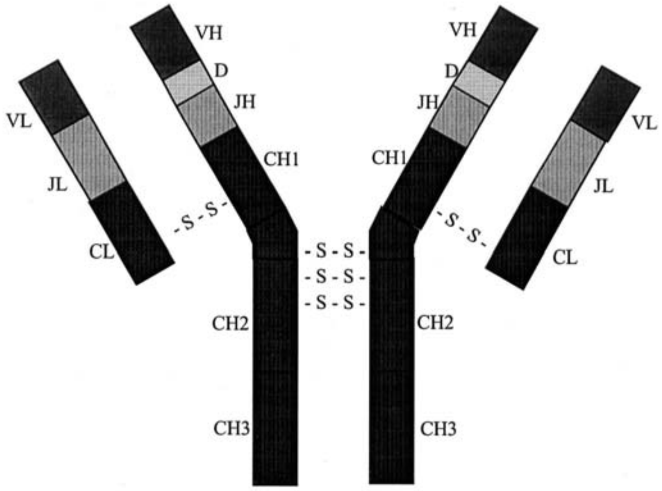

Figure 89 – Typical heavy-light chain bonds (Tyrell 1999)

Antibodies recruit the complementary system of K (killer) cells to fight infection. The complement system acts on its own and in cooperation with antibodies in defending vertebrates against infection. Its activation involves sequential proteolytic cascade. It is composed mainly of inactive blood proteins that are sequentially activated in an aplifying series of reactions either by classical pathway which is triggered by IgG or IgM antibodies binding to antigens as complexes, or by alternative pathways triggered directly by cell envelopes of invading micro-organisms. The most important complement component is the C3 protein which can be activated by proteolytic cleavage and then binds covalently to nearby membranes; micro-organisms with activated C3 (C3b) are readily ingested and destroyed by phagocytosis. C3b also initiates assembly of late complement components which form large transmembrane lytic complexes in the membrane to kill cells. Complement activation also releases a variety of small soluble peptides that attract polymorphonuclear leucocytes and stimulate mast cells to secrete histameric – this results in inflammatory response at sites of complement activation. The complement proteolytic cascade is focussed on membranes of target cells by the fact that several of its components including C3b remains activated for less than 0.1s. There are three functionally distinct subclasses of T cells: (i) cytotoxic T cells which directly kill foreign cells; (ii) helper T cells which help B cells make antibody responses, help other T cells make cell-mediated increase responses, and activate macrophages by secreting lymphokines; (iii) suppressor T cells which can inhibit responses of B cells and other cells. Helper and suppressor T cells are the main regulators of immune responses. They interact with their target lymphocytes by recognising either foreign antigens or receptor idiotypes on the target cell surface. Transplantation reactions are T cell mediated immune responses. Most T cells recognise foreign antigens only when these antigens are associated on cell surfaces with membranes glycoproteins encoded by genes in major histocampatibility complex (MHC). In man, they are called HLA antigens. T cells recognise foreign antigens in association wiIh self MHC molecules. There are two main classes of MHC glycoproteins: (i) class 1 MHC glycoproteins are expressed as all nucleated somatic cells and operate by presenting viral antigens to cytotoxic T cells; (ii) class II MHC glycoproteins recognised in association with foreign antigen by helper T cells and which are expressed on most B cells, some T cells and macrophages, and on specialised antigen-presenting cells. MHC glycoproteins serve as primitive antigen-binding receptors to ensure each antigen type involves the appropriate T cell response.

3.7.2.1. Antigen-antibody matching

An antigen is any substance that provokes a response from the immune system when introduced into an organism – substances as simple as the metal nickel can act as antigens, but most naturally occurring antigens are proteins. The basic principle underlying the immune system's recognition of antigens is the secretion by B-cells of antibodies: protein molecules with a section called the "paratope" that chemically attaches or "binds" to a part of the antigen (the "epitope"). The strength of this binding between the antigen and antibody, or "affinity", depends on the relative shapes of the paratope and epitope – if the paratope fits neatly into the epitope then the affinity between the two molecules is said to be high. One type of antibody is specific to one type of antigen (i.e. a particular chemical compound), though an antigen may have several epitopes and thus be capable of binding to several types of antibody. This is similar to the idea of a bunch of keys (epitopes on an antigen) that each fits into one type of lock (paratopes on antibodies). Due to the specificity of antibodies and the presence of a wide variety of antigens, a similarly wide variety of antibody must be produced – adult humans possess about 25 million types of antibody in their bloodstream. If an antibody successfully binds to an antigen the pathogen represented by the antigen is assumed to be present, though further events may also be required before an immune response is produced (see sections 3.7.2.5 and 3.7.2.6 below).

In AISs, antigen-antibody matching is modelled by comparing test data with sequences of characters (strings) called "detectors", the equivalent of antibodies, and determining their similarity. The symbols making up the detectors can represent anything from individual bits up to commands in computer programs – strings of bits are most popular. If the two strings "match" then a virus must be present and appropriate measures should be taken. Importantly, viruses in either the test data or the detectors will allow matching, and hence the detectors do not need to be themselves protected against viruses.

The likelihood of an exact match, with the two sequences completely identical, decreases with increasing length of the detectors so ways of detecting partial matches are required. One popular method, thought to be similar to the natural mechanism, is the *r*-contiguous matching scheme: two strings match if there are identical symbols at *r* neighbouring positions anywhere in the two strings.

Like all such systems, detection of viruses by an AIS is probabilistic rather than certain – the probability that a virus will be detected can be traded off against the amount of computing power required to implement it. An important value for this trade-off is P_M, the probability that a random detector will match a random piece of data – this is in effect the specificity of the detector and so governs how many detectors are required to give a certain probability of detecting a virus. The value of P_M depends on the particular matching system used. For the *r*-contiguous method, P_M can be estimated by the following if $m^{-r} << 1$ [Percus et al. 1993]:

$$P_M \approx \frac{(l-r)(m-1)}{m^r\,(m+1)}$$

where:

r = *number of neighbouring characters that must be the same*
l = *length of detector/data string*
m = *number of symbols in the alphabet*

Alternatively, P_M can be calculated exactly as the coefficient of x^l in the power series expansion of the generating function $R_p(r, l)$ [Ranang 2002]:

$$R_p(r,l) = \frac{p^r x^r (1-px)}{(1-x)(1-x+qp^r x^{r+1})}$$

where:

p = *probability that two random symbols chosen from the alphabet are the same* = m^{-1}
q = *1* – p
x *is the dummy variable required in a generating function*

3.7.2.2. Antibody generation and negative selection

Since the human genome contains on the order of 100,000 genes, it cannot include separate plans for all the different antibodies used by the immune system. Instead, the genes containing the blueprints for antibodies have alternative options for each part of the antibody. During the development of a B-cell, one option is chosen at random for each part in a "mix-and-match" fashion to give unique DNA sequence out of a large number of possible permutations. After decoding by B-cells, this gives rise to the required diversity of antibodies which between them are able to detect almost any possible substance. This process is shown in Figure 90.

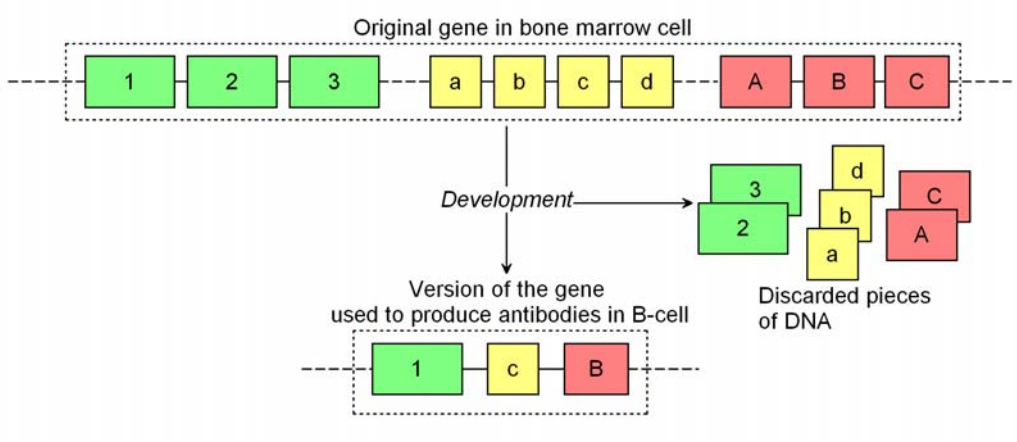

Figure 90 – A diagram showing the modification of an original gene, containing alternative versions of several sections of DNA, into one that can code for a functioning antibody

By chance, some of the B-cells thus generated will be specific to molecules associated with the organism itself which would lead to the immune system attacking the body's own tissue. To avoid this happening, newly-produced or "immature" B-cells go through a period of "maturation" or "tolerisation" during which they are tested against a variety of "self" tissues; if the antibodies produced by a B-cell bind to a molecule known to belong to the host organism, the B-cell is destroyed. At the end of this "negative selection" process, only B-cells whose antibodies react with non-self molecules remain; they are described as "naïve mature" B-cells.

For AISs, the simplest way to produce detectors, and that most analogous to the way the immune system is thought to operate, is the "exhaustive" algorithm: random number generators produce a large number of possible detectors, which are then compared against a training set of data known to be virus-free –

those detectors that match are eliminated [Forrest et al. 1994]. The disadvantage of this algorithm is that very many detectors need to be generated and screened to find a sufficient number that do not match self data.

Alternative algorithms that are more efficient at finding appropriate detectors have been found, but only for systems using the *r*-contiguous matching criterion [D'haeseleer et al. 1996; Ayara et al. 2002]. In other cases where the number of non-self states is small the detectors have been hand-picked to match them.

The number of detectors required is dependent on the desired probability that the system will detect a virus, since the more detectors there are the more likely it is that the system will notice viral data. The number of detectors required for a certain probability of detection P_d is given by:

$$N_R \approx \frac{-\ln P_f}{P_M}$$

where:

N_R = *number of detectors required*
P_f = *allowable probability of detection failure = 1 –* P_d
P_M = *probability that a random detector will match a random string, as described in section 3.7.2.1 above*

3.7.2.3. Clonal selection and affinity maturation

In the natural immune system, antibodies function both to detect antigens and as part of the defence mechanism against the corresponding pathogens. To fulfil this latter function, however, requires much higher concentrations of the antibodies than the former – and since most types of antibodies will never come into contact with their appropriate antigen, it would be a waste of the body's resources to produce all types at a high concentration. Instead, all types of antibodies are initially produced in small numbers, but if a particular type of antibody discovers an antigen to which it binds then B-cells that produce that type of antibody are stimulated to increase in number (forming a "clone", a group of cells identical in genetic make-up) until they can secrete enough antibodies to defend the organism. This is termed "clonal expansion", and the whole mechanism of producing large numbers of only those B-cells whose antibodies react with pathogens is referred to as "clonal selection".

In addition, during the replication of B-cells in clonal expansion the copying of genes containing the plans for the antibody is subject to many more errors than usual, around 10^{-3} errors per base pair per generation rather than 10^{-9}. This phenomenon is known as "somatic hypermutation". These errors, or "mutations", are likely to have some effect on the structure of the antibody produced by the gene. If they make it bind to its antigen less strongly than before, the strain of cells carrying those particular mutations dies out; however, if they make the paratope fit the epitope more closely then the modified strain of B-cells will reproduce more rapidly and come to dominate over the original strain, so the majority of the antibodies produced will be of the improved type. Thus a type of antibody with an initially low affinity for an antigen can be optimised to bind more strongly through the process of "affinity maturation".

In AISs, reproduction of detectors in an analogue of clonal expansion is not usually required. This is because the method for dealing with a virus in a computer is separate from the detection mechanism and so a single activated detector is sufficient to stimulate anti-virus systems, in contrast to the natural immune system where the same cells fulfil both detection and defence functions. However, if somatic hypermutation is implemented in an AIS, it is necessary to create copies of an activated detector in order to mutate them – the usual procedure is to produce a number of copies that is proportional to a detector's affinity to the viral data and mutate each of these copies by an amount inversely proportional to its affinity through randomly changing a few of the symbols that make it up [Ranang 2002].

The term "somatic hypermutation" has also been used to describe a slightly different algorithm in Ayara et al. (2002), where during the negative selection process detectors that matched the "self" data underwent guided mutation until they no longer matched and so could be used. This is not strictly the correct interpretation of the term from a biological point of view, but it is still a useful principle to implement in an AIS.

3.7.2.4. Immune memory

When a B-cell reproduces itself after being activated, some of the resulting cells do not immediately take part in the immune response. Instead, they remain alert in the bloodstream as "memory cells", ready to react to their matching antigen if it reappears in the body. The "secondary response" to an antigen (the second time it appears) is therefore both more forceful and more quickly produced than the primary response. This explains why humans can become "immune" to certain diseases (such as chicken pox and measles), and why vaccination is effective: in each case the initial exposure to the antigen causes the production of memory cells against it, so if the pathogen is encountered again in later life infection is rapidly fended off before symptoms of the disease become noticeable.

AISs implement this behaviour in a straightforward manner, by compiling a list of all detectors that are known to have been activated in the past (often called "memory detectors") separate from the list of naïve detectors. Depending on the system used, the memory detectors may be checked for matching more frequently than the naïve ones, or more importance may be attached to their matching when it occurs (see section 3.7.2.5 below).

3.7.2.5. Costimulation

It is good to have a very sensitive immune system, but safeguards are needed as well. If an immune response could be triggered by an single antibody matching an antigen, the potential for erroneous responses would be unacceptably high – such "false positive" responses would likely involve reaction against the organism's own tissue (an "autoimmune" response) and could be very damaging. This problem is avoided in two ways, the first of which is the requirement for "costimulation" before a response is launched (the second is described in section 3.7.2.6 below). The costimulation signal is an independent indication that an antigen is non-self, from a system of antibody-like receptors attached to a different type of immune cell ("helper T-cells"). Thus, even if a B-cell whose antibodies react with self molecules has slipped through the negative selection net and become a naïve mature B-cell, it still cannot trigger an autoimmune response by itself; indeed, if a B-cell's antibodies bind to antigens but it does not receive the costimulation signal, it is destroyed. However, once an antibody has been proven to be useful (i.e. the B-cells that secrete it have become memory cells), no costimulation is needed to provoke a response.

The usual way in which costimulation is implemented in an AIS involves sending a message that a detection has occurred to a human operator who takes the place of the helper T-cell, and waiting for their confirmation that the detection represents a real virus before the anti-virus software is triggered. If this confirmation is not received within a certain time, the detector is assumed to have matched with self data and is destroyed [Kim & Bentley 2002]. Once a detector has become a memory detector it may be able to activate anti-virus countermeasures without costimulation.

3.7.2.6. Thresholds and inflammation

The second method for reducing the number of false positives relies on the B-cells having an activation threshold, which prevents an immune response unless a certain minimum concentration of the antigen is present. This threshold can be modified by a number of other chemicals; in particular it appears that some of the chemicals that stimulate inflammation, produced by damaged body cells and activated immune cells to help fight off an infection by recruiting more immune cells to a particular area, may also decrease the threshold. This would mean that the response of B-cells to antigens is more rapid than usual in a region where damage is known to have occurred, allowing early detection and destruction of the other pathogens that are probably present.

In AISs, thresholds are frequently used to reduce the number of false positives, especially if humans are required to provide costimulation (see section 3.7.2.5 above). They have been modelled by allocating each detector a “match counter”; this counter is incremented every time the detector matches a sample string, but the count also gradually decreases over time. A response is triggered if the count rises above a certain threshold value – i.e. if there are more than a certain number of matches in a given time period [Hofmeyr & Forrest 2000]. Like in natural immune systems, in some AISs the threshold value can be reduced by “inflammation”: the whole AIS has a “sensitivity level” variable which is increased every time a match is found by a detector that had not been successful recently (i.e. whenever a detector’s match count goes from 0 to 1) and, again, gradually decreases. The higher the sensitivity level, the lower the threshold levels of the detectors. This guards against the possibility of a virus stimulating multiple detectors at a level that would be too low to cause a response from any one detector in isolation; by allowing a match by one detector to count as a partial match by others, such a virus can still be uncovered.

3.7.2.7. Immune networks and idiotopes

The concept of immune networks, proposed by Jerne (1974) and later experimentally verified, suggests that antibodies can also act as antigens as they possess an epitope (a region to which other antibodies can bind) in addition to, and different from, their paratope. Such an epitope is called an “idiotope”, and the collection of idiotopes on an antibody is called an “idiotype”. Thus, an immune response and the consequent large-scale production of antibodies against a foreign antigen may stimulate the production of anti-antibodies which reduce the concentration of the original type, as a negative feedback mechanism. This is in contrast to the traditional view that antibodies recognised only foreign molecules and different strains of B-cells were entirely independent of each other. Such networks of antibody-antibody interaction could allow sophisticated control of the immune response to a pthogen, as well as playing a part in immune memory – however, their role in natural immune systems is disputed by many immunologists.

In some immune system models incorporating the ideas of immune networks, all types of antibodies are present in equilibrium with each other when no pathogens are present; the concentration of an antibody is enhanced if it matches the idiotopes of others and diminished by others binding to its idiotope. In this interpretation, the epitopes of foreign antigens are identical to idiotopes of some antibodies in the network (the antibodies are described as the “internal images” of the antigens) and so a pathogen invasion corresponds to an apparent increase in the level of an antibody – an immune response is seen merely as the attempt of the network to return to its initial equilibrium state.

Although immune networks are still not clearly understood in humans and other animals, versions have been included in AISs where they seem to have been successful. The immune system possesses the properties of learning, memory and pattern recognition without preprogramming through an evolutionary mechanism, and strongly resembles classifier systems [Farmer et al 1986]. Such an approach offers considerable promise for combating unauthorised infection of computer systems by viruses – in 1986, a satellite TV broadcast in the USA was hacked into by an unauthorised hacker "Captain Midnight". The immune system can distinguish between foreign molecules (antigens) and molecules that comprise itself. Antibody molecules – Y-shaped molecules – recognise and tag specific foreign invaders for removal by phagocytic lymphocytes. Identification of the antigen is based on the antibody combining region (paratope) of the antibody molecule which is a groove or lock geometry – if this matches the shape of the foreign antigen, the antigen attaches itself at the matching portion of the antigen (epitope) which is the complementary key. Such pattern matching can be simulated computationally by matching complementary segments of binary strings, the degree of match determining the strength of the bond. The antibody combining region (paratope) of the antibody molecule is a groove or lock geometry – if this matches the shape of the foreign antigen, the antigen attaches itself at the matching portion of the antigen (epitope) which is the complementary key. A paratope may function as a epitope, and an antibody/antigen may possess several different epitopes. Around 10^5 antibodies of the same type with identical paratopes reside attached to the surface of each B-lymphocyte. If an antigen epitope is detected, the lymphocyte clones itself to amplify its population of lymphocytes with the same antibodies. There are $\sim 10^7$-10^8 different antibody types in the mammalian body. The antibodies are generated randomly within bone marrow $\sim 10^6$ per second. To ensure that self-generated proteins are not tagged, self-destructive antibodies are regulated by special antibodies generated and learned during development. In a computational scheme, the degree of proliferation is determined by the strength of the match between the antibody and antigen. The construction of antibodies is based on a fixed set of DNA building blocks which may be combined in different ways to produce a diverse set of antibodies, each capable of locking a class of antigens. Diversity is achieved by reshuffling the protein segments of the antibody corresponding to the fixed set of DNA building blocks. Antibody molecules are composed of two polypeptide chains – a heavy chain and a light chain. Any of the 10^4 light chains may combine with any of the 10^4 heavy chains to generate up to 10^8 different antibodies. The heavy chain is coded by four separate genes (V, D, J, C genes). The light chain is coded by three separate genes (V, J, C genes). Light and heavy chains are assembled by first selecting one gene from the V-gene library then sequentially selecting from each of the other gene libraries randomly. Thus is a vast combinatorial increase in possible antibody types. Furthermore, antibody genes undergo high mutation rates increasing their diversity further. In one computational scheme based on the immune network system, the antigen correlates to sensory input which stimulate the selection of specific antibodies (correlating to tasks or processes to be performed). For simplified computational modelling purposes, heavy/light chains may be treated a binary variables 0 and 1 with each antibody (and antigen) comprising one epitope. An antibody is thus represented as a pair of strings (p, e) of length l_p and l_e respectively with a matching threshold of s below which no reaction occurs. Rate of change of concentrations of antibodies is modelled by [Farmer et al 1986]:

$$\dot{x}_i = c\left(\sum_{j=1}^{N} m_{ji}x_i x_j - k_1\sum_{j=1}^{N} m_{ij}x_i x_j + \sum_{j=1}^{n} m_{ji}x_i y_j\right) - k_2 x_i$$

where:

x_i = *concentration of* ith *detector*
N = *number of detectors*
n = *number of antigen types*
y_j = *concentration of* jth *antigen*
c = *rate constant*
k_1 = *constant to give asymmetry between recognising and being recognised*
k_2 = *decay rate constant*
m_{ij} = *matching parameter of ith paratope with jth epitope.*

An antibody is thus represented as a pair of strings (p, e) of length l_p and l_e respectively with a matching threshold of s below which no reaction occurs then matching parameter is given by [Farmer et al 1986]:

$$m_{ij} = \sum_{k=-rng}^{rng} G\left(\sum_{n=1}^{l} e_i(n+k) \wedge p_j(n) - s + 1 \right)$$

where:

$e_i(n)$ = n*th character of* i*th epitope string*
$p_j(n)$ = n*th character of* j*th paratope string*
G(x) = *threshold function* = x *for* x > *0*
= *0 otherwise*
k = *index for shifting epitope over paratope*
n = *position in paratope*
^ = *equality operator = 1 if arguments equal*
= *0 otherwise*
l = *number of symbols in epitope or paratope, whichever is the shortest*
s = *threshold for matching* < l
rng = *range over which epitope is shifted* = l – s

This encapsulates the stimulation of antibody type i by antigen type j, the suppression of antigen i by antibody j, the consequent change of antigen concentration y_j, and decay of antibody/antigen cells without interaction. New antibodies can be created by applying genetic operators such as cross-over by selecting two antobody types and interchanging their strings. The antibodies with the highest concentration are selected and the corresponding behaviour selected for action, the others being suppressed. The connectivity between antibodies may be learned by reinforcement learning.

Immune networks are suited to different types of problems from more conventional AISs –they give an analogue output, the concentrations of a set of detectors, whereas the methods described above return a binary classification of a piece of data as self or non-self. Frequently immune network systems are used for discovering clusters of data and finding near neighbours in terms of similarity: the target item of data is encoded as an antigen and the others as detectors, and similar pieces of data are represented by the detectors that have a high concentration once the network has stabilised.

If the paratope and epitope of a detector are identical, network interactions in an immune network can to reduce redundancy by removing detectors that are too similar to others – this has been used for data classification in the system called aiNet [de Castro & Von Zuben 2001].

3.7.2.8. Other principles

Protected maturation: The thymus gland (situated in the chest) and the marrow of large bones play an important part in the natural immune system, as several types of immune cells finish their maturation in one of these locations. Negative selection takes place during their maturation, so it is important that non-self material is excluded from these locations lest the immune system grow tolerant of it. This idea has been adapted to AISs by training detectors against known uncorrupted data, or, if this is not obtainable, test data is checked first against the memory and mature detectors – if a match occurs then the item of data is not used for the negative selection of immature detectors [Ranang 2002].

Finite lifetimes: A real immune system is dynamic in composition, with the destruction of naïve mature immune cells that have not detected a pathogen after a finite time and the continual production of new immature immune cells. If the system had been static, it would become tolerant of a foreign antigen that had happened to be present while the immune cells were maturing, while unnecessary cells that never matched detectors would remain a drain on resources; its dynamic nature avoids both of these problems. Thus finite detector lifetimes coupled with the continuous generation of new detectors are incorporated into many AISs.

Decentralisation: A major advantage of the natural immune system is its decentralisation – immune cells are manufactured in the marrow of a number of large bones in the body, finish their development either there or in the thymus gland and circulate in the bloodstream and lymphatic system once mature. This means that damage to any one area of the body cannot destroy the immune system, making it extremely robust. The same principle can be applied to AISs if they are implemented on a network of computers: instead of one server running the AIS program, each computer can have its own set of detectors and its own sensitivity level, making each computer the equivalent of a lymph node and preventing a successful viral infection of a single node from disabling the whole network's anti-viral systems. This approach can be further enhanced by exchanging activated detectors between nodes so that the whole network becomes immune as a result of one node's experience.

Danger theory: There are several observations that the self/non-self recognition hypothesis was unable to explain, including the fact that there is no immune reaction to the foreign bacteria that inhabit the digestive system or to food, but that there is a reaction to cancerous self cells which should be outwardly identical to other self cells. The explanation suggested by Matzinger (1994) is that a B-cell effectively requires a third costimulation signal, either specific molecules that are released from diseased and dying cells or the lack of signals that confirms a cell's well-being, in order to stimulate an immune response after binding of its antibody to an antigen. Such a signal represents "danger", not merely "non-self", and enables B-cells to ignore harmless non-self while responding to harmful self. In an AIS implementation, danger theory methods would involve combining or replacing the conventional human costimulation signal with a confirmation-of-danger signal based on evidence of a viral infection (such as abnormal memory or disc usage, abrupt process termination or unusual patterns of system activity). To date, few implementations have been constructed.

MHC proteins and permutation masks: Almost all cells in the body produce a set of proteins on their surface called the Major Histocompatibility Complex. MHC proteins are of two main types, both of which function by presenting small pieces of protein ("peptides") to cells of the immune system – the difference lies in the origin of the peptides, which in the case of MHC class I proteins are samples of what the cell is currently manufacturing (which may be foreign antigens if the cell has been infected by a virus) whereas MHC class II proteins present samples of proteins the cell has taken in from outside (which may be from viruses, bacteria etc.). There are many different varieties of these proteins (out of which each individual only produces a small number) and each binds slightly differently to a given peptide, presenting it in a slightly different way; there is evidence that this diversity of presentation contributes to the ability of an organism to resist diseases, and the population-wide diversity means that it would be very unlikely for the whole population to be susceptible to a single pathogen. Although there is no need for specific mechanisms to present test data to detectors in an AIS, as unlike B-cells the detectors can read the insides of files without destroying them, diversity can be useful in an AIS that is spread over a number of nodes by reducing the prevalence of "holes" (non-self data strings that cannot be detected by any of the detectors). Diversity has been achieved by several means [Hofmeyr & Forrest 2000], including reordering the characters in a test data string with a "permutation mask" that was different for each node (e.g. under the permutation mask 1-6-2-5-8-3-7-4, an initial data string of "abcdefgh" would become "afbehcgd").

3.7.3. Comparison with Other Biomimetic Computing Paradigms

3.7.3.1. Genetic Algorithms

AISs are very similar to genetic algorithms (GAs) in a number of ways, summarized in Table 14.

Table 14 – Comparison between AISs and GAs

Artificial immune systems	Genetic algorithms
Large initial population of randomly-generated detectors	Large initial population of randomly-generated individuals
Best detectors are modified by somatic hypermutation	Best individuals are modified by mutation and recombination
Best detector is selected on the basis of matching a given template	Best individual is selected on the basis of satisfying given criteria
Worst solutions are retained in case they match an as-yet unseen antigen	Worst solutions are discarded and replaced
Continuous production and testing of detectors	Production and testing of detectors occurs in discrete generations

The most important difference between the two methods is the fact that an AIS can only produce detectors for a complete template data set. GAs, on the other hand, can evolve solutions to a problem without a template to follow – all that is required is a way of evaluating the individuals in the population against criteria for a successful solution.

The similarities between the two approaches have been widely recognized, and Nikolaev et al. (1999) made a direct comparison of the solutions to a couple of machine learning problems found by the two methods.

3.7.3.2. Artificial Neural Networks

Analogies can also be made between immune networks and artificial neural networks (ANNs) as detailed in Vertosick and Kelly (1991). These are summarized in Table 15.

Table 15 – Comparison between AISs and ANNs

Immune networks	Artificial neural networks
Consists of many interacting detectors	Consists of many interacting neurons
One detector can affect many others	One neuron can affect many others
Input is an antigen or set of antigens	Input is a pattern of neural activity
Output is a set of concentrations of detectors	Output is a pattern of neural activity
Network structure stored as interaction strengths	Network structure stored as synaptic weightings
Interaction is mostly stimulatory to one detector and inhibitory to the other	Interaction can be stimulatory or inhibitory in nature, but only affects the postsynaptic neuron
Patterns are recognized and memories are	Patterns are recognized and memories are

stored by individual detectors	stored by groups of neurons or the whole network

Several pieces of work have exploited this similarity. For example, Hoffmann (1986) proposed an “unorthodox neural network” using ideas from immune network theory, and Abbatista et al. (1996) used an immune network as an associative memory that could store information and retrieve it if exposed to a similar stimulus. An excellent review of recent research in the field is given by de Castro et al. (2003), which also details their implementation of some ideas from AIS, specifically the use of clonal selection and affinity maturation mechanisms for neuronal replication and pruning, in a Boolean competitive neural network that proved able to solve several benchmark problems.

3.7.4. Application Areas

This section gives brief details of the many and varied applications of AISs to date – see de Castro and Von Zuben (2000a) or Dasgupta (1999) for more information.

3.7.4.1. Computer security

The most important applications of artificial immune systems to date have been in the field of computer security, following the example of the natural immune system’s biological role. Three separate kinds of computer security systems have been addressed: virus detection [Khephart 1994] and intrusion detection on a single computer [Forrest et al. 1996], and network intrusion detection (NID) distributed between several computers [Hofmeyr & Forrest 2000]. In the first case, the data given to the AIS consisted of sections of files stored on the computer; in the second case, the data represented the programs running on that computer at a certain time; and in the third, the items monitored were packets of data sent between computers on the network (which contain information about their origin and destination, as well as a message). An excellent review of this area is given in Ranang (2000).

3.7.4.2. Robotics

Watanabe et al. (1999) conducted a series of experiments with a rubbish-collecting robot called “Immunoid” that used immune network methods to decide between alternative courses of action. An action and a set of conditions under which it would take place were used as the paratope of an antibody, and an identification number as the idiotope; the action represented by the most concentrated antibody was executed. Jun et al. (1999) also used immune network methods to decide upon appropriate patterns of collective behaviour and transmit them amongst a group of mobile robots.

3.7.4.3. Immunotronics

A new paradigm in electronics is “immunotronics”, the use of AISs for fault detection. Detectors that have been trained against information on the proper operation of a circuit are compared with the actual input and output signals, and if a match occurs then a fault in the circuit is presumed. A system of this kind was initially implemented by Bradley and Tyrrell (2000): they used a field-programmable gate array (FPGA) which contained both a simple circuit to be protected and an AIS designed to detect improper state transitions and signal a fault. The same team further developed the work and in Canham and Tyrrell (2002) reported an FPGA-based robot controller that could not only detect faults in its operation by AIS methods but also correct them using techniques inspired by the development of embryos, combining immunotronics and "embryonics".

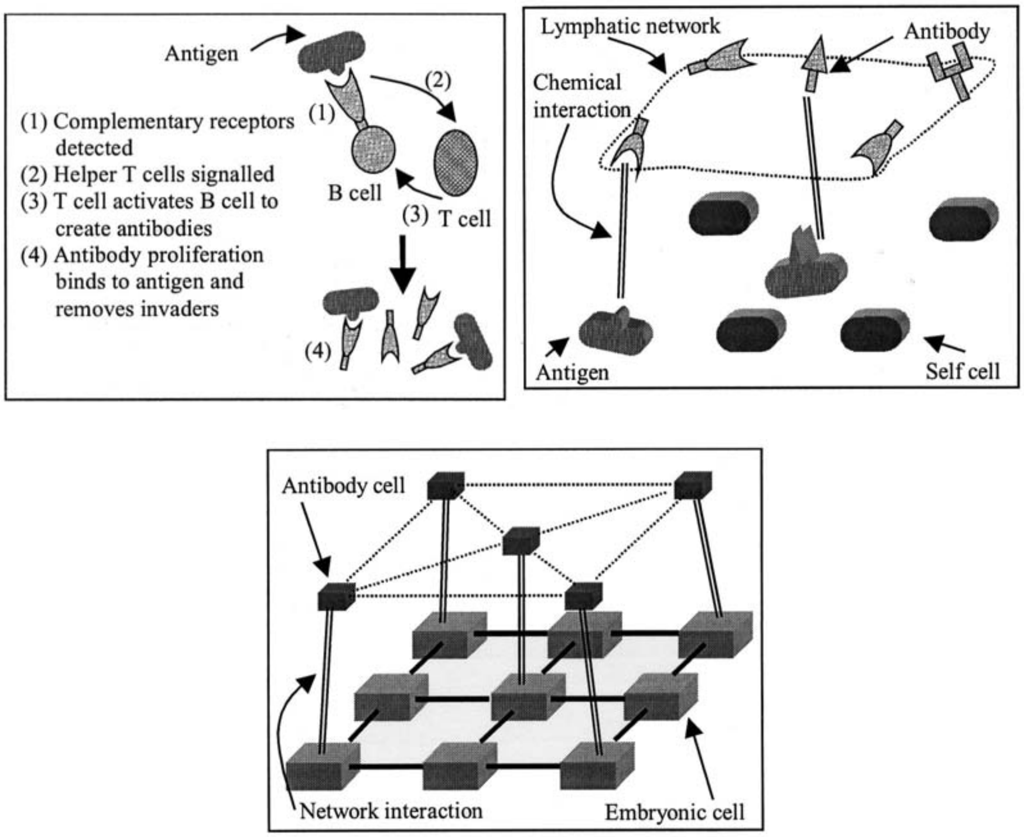

Figure 91 – a) antibody-mediated immunity; b) lymphatic interactions with invading antigens; c) immune-embryonic interactions (Bradley et al 2001)

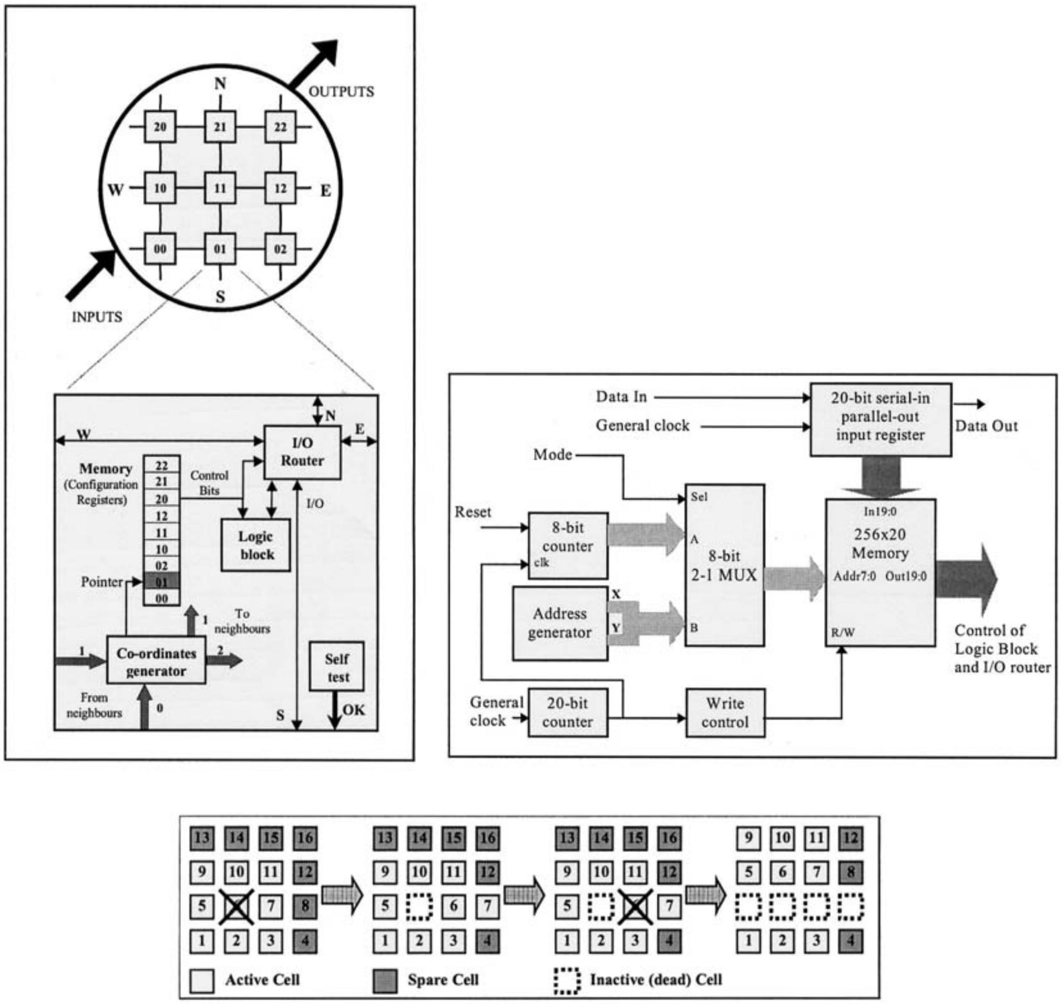

Figure 92 – a) Basic components of an embryonic cell (Ortega & Tyrell 1999); b) Memory system located in each cell (Ortega-Sanchez & Tyrell 1997); c) Fault tolerance by cell elimination (Ortega & Tyrell 1999).

3.7.4.4. Web server management

Suzuki and Yamamoto (2001) used an immune network system called iNet to optimize the behaviour of a web server program to the demands placed on it (in terms of number of files requested, file size etc) by changing a number of different system policies. The different detectors represented different policies; they had a chance of being selected for implementation dependent on their concentration, which was affected by both the conditions of the external demand and the concentrations of other detectors. The system was able to learn and adjust the network interaction strengths by analyzing the resulting performance of the web server.

3.7.4.5. Pattern recognition

The use of AISs for a generalized pattern recognition system has been reported by the Hunt and Cook (1996), who used an immune network system, and others. Work has also been carried out in the following specific areas:

- **Image recognition:** AISs incorporating immune networks have been used for extracting features from aerial photographs [McCoy & Devarajan 1997]: detectors were generated by negative selection to produce a population that responded only to areas in aerial photographs that were not roads. AISs have also been used for image inspection [Aisu & Muzutani 1996], classification of stained wood by colour [Sathyanath & Sahin 2001] and, in a very self-referential application, analysis of images of biochemical tests to determine the reaction of natural antibodies with antigens [Tarakanov et al. 2002].
- **Chemical spectrum recognition:** Dasgupta et al. (1999) applied an AIS to the task of identifying chemical compounds from spectra. They represented each spectrum by a binary string according to the presence or absence of a peak in certain frequency ranges, and then used standard AIS techniques to produce detectors specific to each compound which were added to a library and could be used to identify the compound in future.

3.7.4.6. Anomaly detection

AISs are ideally suited to the detection of anomalies in time-varying data. Negative selection of detectors takes place on data that are known to be free of anomalies, and if a match later occurs then it is likely that the current variables being monitored are out of the expected range. Previous applications have included tool breakage detection in machining [Dasgupta & Forrest 1995] and diagnosis of faults in industrial processes [Ishida 1993]. In particular, immunological analogues have been proposed for robust fault tolerance in computing systems. This immune system analogy has provided the basis for developing fault-tolerant systems of processor arrays [Tyrell 1998; Ortega & Tyrell 1999].

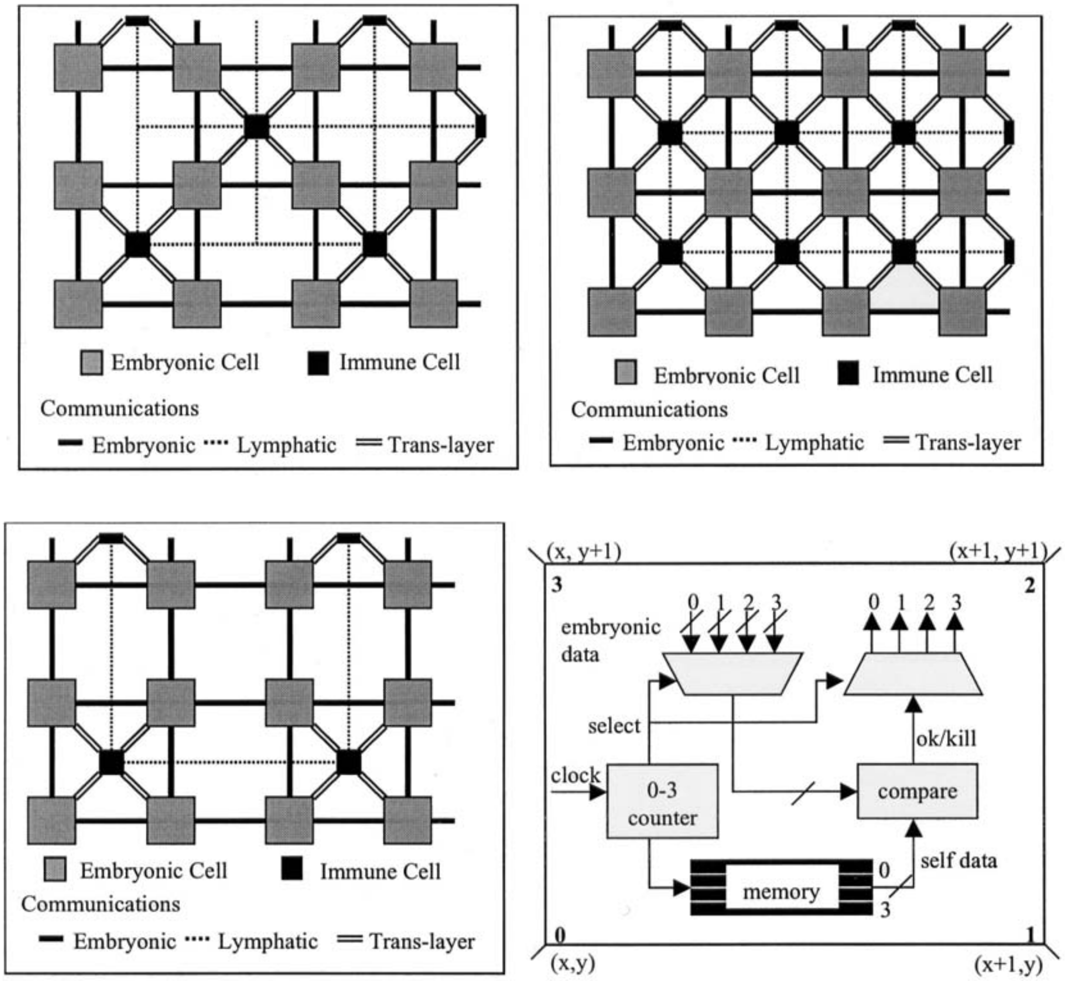

Figure 93 – Antibody cell monitors 4 closest neighbours: a) no replication of antibody cells; b) each embryonic cell is monitored by two antibody cells; c) each embryonic cell is monitored by 4 surrounding antibody cells; d) Architecture of the antibody cell – comparison logic (Bradley et al. 2001).

3.7.4.7. Scheduling

Hart et al. (1999) used an AIS to generate schedules of jobs to be carried out in a fictional machine shop. The first stage of the process used AIS pattern recognition techniques to compile a list of sequences of jobs that were frequently carried out by the machine shop; a second stage then used these sequences to compile new schedules.

3.7.4.8. Data mining

Hunt & Fellows (1996) showed that the immune system had similar capabilities to other case-based reasoning systems (artificial intelligence systems that use memory of their past experience to guide future actions). The antigen-antibody matching principle allowed the grouping of a new item of data with others that it resembled, using an immune network.

3.7.4.9. Multi-agent systems

Sathyanath and Sahin (2002) defined a general framework for applying AIS techniques, specifically the idea of an immune network, to the co-ordination of multi-agent systems. Antigens and B-cells are represented as agents of two opposing types (non-self agents, "NAGs", and self agents, "SAGs" respectively), each possessing a string of information representing their epitope or paratope. The SAGs are able to communicate with each other and decide amongst alternative actions using an application-specific affinity function. Two simple simulated applications were considered: barcode character recognition, and landmine destruction by mobile robots where a group of 4 was required to defuse a mine (information about mine positions was readily available). Dilger (1997) proposed the use of an immune network-type AIS for detection of and response to intruders or other anomalous events (e.g. thunderstorms) in intelligent buildings, where each sensor or actuator was an agent in the system with certain capabilities. Ishida (1996) devised an agent-based AIS algorithm that could be applied to areas such as noise neutralization, group task achievement, model adjustment and decision-making.

3.7.4.10. Learning

The negative selection idea incorporated in almost all AISs can be regarded as a form of reinforcement learning. Potter and De Jong (1998) gave a set of training data, whose contents had been labelled if they were examples of a particular concept, to an AIS and a symbolic inductive learning system called AQ. They compared the population of detectors with the symbolic rules abstracted by AQ, and concluded that the AIS represented more concise description of the concepts with similar predictive accuracy.

3.7.4.11. Recommenders

Computer programs that give users recommendations for films they might enjoy watching [Morrison & Aickelin, 2002] and web sites they might find interesting [Cayzer & Aickelin, 2002], based on their current preferences, have been constructed using immune network principles. AIS methods were used to find a group of users whose preferences were similar to those of the target user, by encoding the target user as an antigen and all other users as antibodies and iterating their concentrations until they stabilised. The items preferred by the group thus discovered represented the recommendations. In the case of the film recommender idiotypic interactions encouraged diversity in the group of people finally chosen, but in the latter it proved unnecessary as there were many more possibilities for websites (even when sorted into categories) than films.

3.7.4.12. Optimisation

AIS methods have been used analogously to genetic algorithm methods (see section 3.7.3.1) with which they share many characteristics, to optimise solutions to problems. These have ranged from the design of a 10-bar truss [Hajela & Yoo 1999] to the n-agent travelling salesman problem [Toma et al. 1999]. More generalized optimisation algorithms have been produced by Fukuda et al. (1999) and de Castro and Von Zuben (2000b).

3.7.4.13. Control

At least two generic real-time control systems have been developed using AIS methods. KrishnaKumar and Neidhoefer (1999) produced an "Immunized Computing System" for autonomous aircraft control which included an AIS to adjust the structure of the control system (made up of fuzzy logic elements connected in a neural network-type arrangement) in response to disturbances. In a different approach, Ootsuki and Sekiguchi (1999) developed an AIS-based method for designing the state transition sequences of state machines associated with industrial process equipment.

3.7.4.14. Inductive problem solving

Slavov and Nikolaev (1998) developed an immune network model to solve the finite-state automaton induction problem, of constructing a finite-state automaton that can replicate a given pattern of behaviour. Detectors represented random finite state machines; their concentration was increased according to how well they match the target behaviour sequence and modified as appropriate by idiotypic interactions with the other detectors.

3.8. BIOLOGICAL BEHAVIOUR CONTROL & NAVIGATION

The robotic system provides the most demanding test of AI techniques since it requires the intelligent control system to interact with the real world via a physical robot body. Sensory input and motor output are not analogues of the standard peripheral read and write commands in a computer programming language. The physical embodiment of the information processor (the body) provides a buffer to the external environment [Fritz et al 1989]. Dreyfus (1967) suggested that bodies are a necessary pre-requisite for intelligence and that disembodied computing systems will never exhibit truly intelligent behaviour. The body provides a fundamental system of reference for cognition. All animals that are mobile require general-purpose perceptual and behavioural repertoires sufficient to deal with a variety of different environments in order to survive compared with animals such as molluscs that remain fixed and so are exposed to a limited repertoire of problems and conditions. Behaviour depends on the environment in which the agent is embedded by virtue of evolution, e.g. foraging strategies for food require complex navigational skills. Control is concerned with the generation of desired behaviours in robotic agents – we consider control at different hierarchical levels. Indeed, control and navigation are fundamental to animal survival. Animal behaviour is adaptive in that it adjusts to the changing external environment. Simple animals like insects are capable of autonomously adapting a limited repertoire of basic behaviours to variations in the noisy complex environment of the real world. Animals possess limited motivations which determine their behaviour – hunger, libido and fear. Such motivations are combined with sensory data to generate dynamic goals and control perceptual attention. Hence, the environment does not produce a given behaviour, but merely triggers it but complex behaviour is still a reflection of the complexity of the environment. Motivation for motor action usually resides in the physiological state of the agent often correlated with deprivation of a resource, e.g. food. Insects are capable of adapting a limited behavioural suite to the uncertainties of the real world. Animal behaviour is adaptive in that it is continually adjusted to meet the changing internal and external conditions. Complex behaviour is generated by a series of sequential behaviours whereby the output of one behaviour triggers the next one.

Simple cyclic muscular movement, particularly those involving the alternation of extensor and flexor muscles in locomotion, appear to involve central neural pattern generators [Delcomyn 1980]. These automatically drive the muscles used in the behaviour pattern in the proper sequence with the proper temporal relationships without the need for sensory feedback – the completion of each movement provides the trigger for the next segment. These behaviours are innate and passed from generation to generation genetically. Such innate reflexes and behaviour pattern generators require no knowledge of causation, space, time, objects or their relations. The innate response has the benefit of immediacy in that no time for learning is required. Innate behaviours are most important for animals with small body size and short lifespans such as insects where there is too little time for learning and too little volume for the large brain required for learning.

The simplest form of behaviour is the reflex which is characterised by a rigid relationship between stimulus and response. A rapid stereotyped response is triggered by a particular type of environmental stimuli and it is usually protective. They allow the animal to rapidly adjust its behaviour to sudden environmental changes. The intensity and duration of the response is governed by the intensity and duration of the stimulus. The stimuli required to activate a complex behavioural pattern is usually only a very small subset of the total amount of sensory information available (sign stimuli). These sign stimuli include behaviours which are absolutely critical to survival and typically animals respond to other objects which are no threat but which still possess the sign stimuli. Furthermore the sensory organs are specialised to sense certain environmental conditions, eg. the frog detects the movement of small dark objects at close range characteristic of food. Indeed, frogs can detect movement but not static objects – static objects are not of biological significance to the frog. Generally animals are capable of abstracting space representations for foraging and navigation and time representations with respect to their internal circadian rhythms. There is the question as to what is the minimal requirement for cognition – Primio et al (2001) take the view that basic cognitive functions of associative learning and memory are properties of all cellular creatures, even those lacking nervous systems such as single-celled bacteria. They identify the most fundamental cognitive abilities to be distinguishing of the environment from self, multisensory identification of elements in the environment, to adapt to changes in the environment, and to decide between alternative courses of action (thereby exhibiting a degree of purposefulness). Chemotaxis allows bacteria to avoid toxic environments and move towards attractive environments for food. Bacteria swim within chemical gradients detecting differences in stimulus intensity and integrating these over time. Simple memory is based on the number of sensors stimulated. Their search strategy is simple:

```
LOOP: REPEAT twice;
        linear movement for 1s; !counter-clockwise coordinated flagella rotation
        IF conditions are improving !increasing attractant or decreasing repellant
              THEN do nothing
              ELSE tumble randomly for 0.1s !clockwise uncoordinated flagella rotation
        END REPEAT;
        tumble randomly;
END LOOP
```

This represents a minimal, probabilistic search algorithm. Exposure to attractant stimulates the glutamate receptors of its chemosensors, thereby switching them off – this ensures that the flagella function in a coordinated fashion by default for linear movement. The chief requirements to enable such behaviours are the existence of a variety of internal and external sensors and the existence of motility provided by rotating flagella. In addition, bacteria posses the ability to modify their actions to stimuli and alter their behaviour to the same stimuli over time, i.e. possess both simple learning and simple memory (one-shot memory rather than stack memory). Bacteria utilise switchable sensors whereby prolonged exposure induces methylation of glutamate receptors of its chemosensors, thereby gradually switching them on – this activates the flagella to operate in uncoordinated fashion for tumbling as if attractant has been lost. The glumate receptors are de-methylated over time switching the flagellum from on to off to return to is default algorithm. This suggests that even simple organisms must possess indirect coupling between sensors and actuators through memory as the core of cognition (in this case habituation). Development beyond this core cognition introduces the capacity for anticipatory behaviour based on predictions which requires an internal model of both the environment and itself. Animals developed the ability to move through muscular exertion controlled by a signal processing nervous system. To relate behaviours to the outside world, sensory organs evolved. One of the most primitive creatures is the sea anemone which includes a system of sense organs linked directly to a motor system. The nervous system evolved subsequently as an intermediate system.

Much animal behaviour is the result of innate responses to certain environmental stimuli which allow animals to respond to different situations. Many types of behaviour are specific to certain to certain species. Taxes are orientation responses toward or way from certain environmental agents such as light or gravity and is exhibited by plants. Fixed action patterns are more complex and are extended responses to triggering stimuli. The fixed action pattern involves a temporal sequence of component behaviours which if triggered, run to completion, e.g. escape response and courtship rituals. There appears to exist in the central nervous system an innate releasing mechanism whereby a specific response is induce by certain specific environmental stimuli. The stimulus acts as a releasing mechanism to the fixed action pattern response. Innate behaviours are automatic and data-driven in that they are evoked by specific stimuli in the environment which releases specific action behaviours. A rigid relation exists between the stimulus and response providing reflexive behaviour with the benefit of immediacy. The fundamental unit of behaviour is thus the reflex arc which is characterised by a rigid pattern of sensorimotor coordination. Behaviours are well-defined control laws that are goal-oriented. The required input stimulus to activate reflexive responses is much less than the total sensory information available. This minimal sign stimulus is guaranteed to evoke response and is mainly due to specialised stimulus recognition hardware which increases the speed of response. Such sets of innate behaviours tend to be species-specific and are manifest as primary drives such as hunger (energy resupply), mating (self-replication), and fear (self-preservation). Such drives exhibit the characteristic that they are satiable. The threshold for triggering fixed action patterns varies with internal states such that the brain actively selects the sensory information to direct the animal's response. This is motivated behaviour which is selected according to its consequences and the motivational potential affects possible behaviours, e.g. hunger. Motivated behaviours are characterised by sequences of component behaviours, goal directedness, spontaneity in the absence of stimuli, responsiveness changes according to the level of arousal or satiation, and persistence beyond the initial stimuli. In this way an animal can coordinate its many possible actions into coherent behaviour directed towards long-term survival. It appears that the behavioural repertoire is a product of evolution and has embodied a form of long-term learning. Complex behavioural patterns are decomposed into more simple behaviours in a structured organisation. Some behaviours usually take precedence over others such as feeding over exploration and some may be mutually contradictory if they share the same motor apparatus. This organisation may be hierarchical or non-hierarchical. Animal behaviours as instinct centres are hierarchically organised with negative feedback connections between them – the Tinbergen model [Baerends 1976]. The three relevant basic centres are locomotion, food-seeking and pain aversion to generate feeding and fleeing behaviours under the appropriate conditions. Each instinct centre is decomposed into finer-grained behaviours at lower levels. Simple animals such as insects behave through standard responses to stimuli. More complex vertebrates augment these responses with a world model which they use to select sensory information to plan courses of action to achieve its goals. The addition of learning provides the ability to adaptively use the world model to predict and compare predicted and actual outcomes of courses of action to update the knowledge structures for further use.

The bottom-up approach to intelligent behaviour reduces the requirement for explicit representation of the external environment on the basis that the physical world is its own best model. Intelligence is seen as the result of the interactions of sets of non-linear processes. The central core notion in such approaches is that higher-level behaviours of greater complexity will emerge through the interaction of the component simpler behaviours and its environment without the need for explicit symbolic representation of those higher-level behaviours. The outside world is only partially observable and undergoes continuous change. Particularly important is the ability to cope with unexpected events which cannot be anticipated but which must be dealt with by concurrent contingency routines. The animat approach has become dominant in robotics in which autonomy, survivability and robustness are the key characteristics inspired by the capabilities of insects despite their limited neural resources in which ecological niche is an important design consideration [Wilson 1991]. This is the Animat approach to robotics which simulates biological approaches by focussing on complete systems and on the tight coupling interaction between the robotic agent and the environment in which it is situated (situated robotics). Reflexive/instinctive control is part of the behaviour control repertoire of all animals – these are the innate behaviours which are often simple and rapid being associated with self-preservation and are exhibited by Braitenburg vehicles. Some of these behaviours are elemental and include central pattern generators which provide the basis for rhythmic movement of muscles. Braitenburg vehicles indicate that simple mechanisms can yield complex behavioural manifestation in the appropriate environments.

Behaviour control methods based on insect behaviours are an approach which have been highly successful [Brooks 1986]. Planner-based control strategies employ a centralised world model generated updated from sensor fusion data, while reactive control strategies employ reactive condition-action rules linking sensors and actuators directly without sensor fusion into a central world model representation. The generation and maintenance of a single monolithic world model representation is computationally intensive. Although planners appear to be more amenable to analysis and proof of correctness than emergent dynamics, this is in fact not so as any interaction with the real world yields unpredictable conditions. An alternative approach to path generation is behavioural decomposition – a bottom-up approach based on the notion that the world is its best own model. The behaviour control method is an assertion of the Physical Grounding hypothesis which in turn rejects the Physical Symbol hypothesis. This is rooted in the biological paradigm that intelligence evolved through interactions with the real world in the bid for survival. Intelligence must react to the exigencies of the environment and generate robust behaviours to survive. Such reactive-level behaviour-based control was inspired by the singular lack of success of AI in developing real-time intelligent behaviour. The behaviour approach seeks to build intelligent control in robotics by developing modules in which each is capable of task competences and which performs its own sensing, behaviour-generation and actuation. It is a parallel-distributed approach to robotics. There are several tenets to based-based control:

- ➢ situatedness – robots are agents in the real world, so the world is its own best model;
- ➢ embodiment – robots have bodies with actuators and sensors with which they interact with the real world directly;
- ➢ intelligence – intelligence is a property of the robot's behaviour in surviving in the real world;
- ➢ emergence – intelligent behaviour emerges from the robot's interactions with the real world.

Brook's (1990) proposed a "nouvelle" AI paradigm called the Physical Grounding Hypothesis such that autonomy requires representation to be grounded in the physical world – the incompleteness of simulated world models can miss important parameters. The upshot of this is to couple perception and action directly without an intervening world model and to minimise internal symbol processing. This denies the linear sense-think-act cycle philosophy of subdividing intelligent architectures into perceptual-central-motor subsystems. A single object can generate large variations in sensory signals. This perceptual aliasing problem is inherent in attempting to map highly variable sensory stimuli patterns onto the same internal symbolic representation. Perceptual classification is viewed as part of the sensory-motor coordination process. The robot agent can move around so sensory patterns will vary according to distance from the object, the viewing direction and illumination. Hence, only a small segment of the world is filtered by the perception process for analysis, and this perception process determines behaviour. Only a small part of the environment is relevant for behaviour. What is perceived is determined by the sense organs, the transduction process, and the process of learning to discriminate important data. Autonomous robots interact with the external environment and their behaviour is determined by the stimuli they receive as input. Physical embodiment is a necessary component of such interaction and it is through such coupling with the world that effective behaviour emerges. There is no central global control. Global communication arises through local interactions with the environment. Situated robotics provides existence proofs for the emergence of structured behaviours (symbols) through the interaction of primitive behaviours in the real world. Reflexive behaviour is the basis of this approach (physical grounding hypothesis) [Brooks 1990]. This approach couples robotic sensors directly to behaviours without intervening symbol manipulation of internal models [Brooks 1990]. Such reflexive behaviour is characterised by simple stimulus-response behaviour. This is situated action which stresses the dynamics of interaction of the agent and its environment. Perceptual classification is viewed as part of the sensory-motor coordination process. The robot agent can move around so sensory patterns will vary according to distance from the object, the viewing direction and illumination. Hence, only a small segment of the world is filtered by the perception process for analysis, and this perception process determines behaviour. Autonomous robots interact with the external environment and their behaviour is determined by the stimuli they receive as input. The effectiveness of the embedding of such a robot in the real world depends on both the bandwidth of the sensors and the performance of the actuators. Physical embodiment is a necessary component of such interaction and it is through such coupling with the world that effective behaviour emerges.

Anderson & Donath (1990) used multiple components of reflexive behaviours which operated concurrently and independently. They found that high-level global behavioural patterns emerge which are not ascribable to individual simpler components at lower levels. Control structures emerged from the distributed loosely coupled behavioural modules through parallel local interactions of a non-linear nature. This organisation gave rise to more complex and irreducible behaviours. They used avoidance behaviours to cause motion away from a stimulus and attractive behaviours to cause motion towards a stimulus. Each behaviour was mapped onto a potential field where repulsive forces protected objects and attractive forces guided the robot towards specific locations. Generalised wandering behaviour was an exploratory mode of behaviour. Feather (1967) suggested that intelligent behaviour is information seeking and that this provides the basis for curiosity as modelled by such wandering behaviour. It is through exploration that cognitive maps of the environment are constructed and stored as the basis for navigation in complex environments. There is little doubt that as well as propositional forms of representation, humans as animals utilise pictorial forms of representation through cognitive maps of spatial relations. Such spatio-temporal maps provide data on their location, the availability of food and changes of the environment over time. On the basis of this information, predictive monitoring of the environment and the derivation of behaviours that are likely to achieve success are possible.

Passive avoidance behaviour halted movement similar to the "freeze" behaviour in many animals. Active avoidance behaviour avoided approaching objects similar to flight behaviour in birds. These two avoidance behaviours effectively model predator-prey behaviour. Location attractive behaviour was regarded as a component of migratory habits of some animals. Object attraction in the direction of an object simulates food seeking or mate seeking behaviour. Open space attraction for unobstructed travel is also characteristic of animals which use range data for this purpose. Mataric (1995) suggested that a basic set of behaviours comprises of: wandering, avoidance, following, aggregation, dispersion and homing for individual and groups of mobile robots. The behaviours avoidance (of obstacles) and following (targets) imply the ability to discriminate. Each behaviour is selected by the environmental conditions. The behaviours may be combined to generate higher level behaviours. Wandering and avoidance generate safe wandering. Wandering, avoidance and aggregation combined produce flocking. The addition of homing produces flocking towards a specific location. General foraging can be accomplished by combining different behaviours triggered by different conditions, e.g. temporal switching between avoidance, dispersion, following, homing and wandering. Hence, group behaviours can be generated by local environmental conditions. Additional basic behaviours for robotic arms could include grasping and dropping.

The behaviour control approach is based on the implementation of a number of discrete parallel basic "behaviours" (typically implemented as finite state machines) such as "wander", "avoid obstacles", etc which must be mediated [Mataric 1995, Arkin 1998]. The control system is built bottom up incrementally to reflect the evolutionary process. Mataric suggested that a basic set of behaviours comprises of: wandering, avoidance, following, aggregation, dispersion and homing for individual and groups of mobile robots. The behaviours avoidance (of obstacles) and following (targets) imply the ability to discriminate. Each behaviour is selected by the environmental conditions. The behaviours may be combined to generate higher-level behaviours. Behaviour control methods were adopted on Sojourner [Gat et al 1994, Miller & Varsi 1993].

Emergent behaviours such as wandering and navigation arise through coupling between each behaviour component. This coupling may be temporal or spatial. Anderson & Donath (1990) held each of the primitive reflexive behaviours to be spatially ordered. The difficulty with spatially ordered behaviours is that they can cause conflict resulting in erratic behaviour. The total number of unique behaviours emergent from different primitive behaviours may be assumed to be a binomial function in that i possible combinations of n items taken k at a time is given by:

$$i = \sum_{k=0}^{n} \tfrac{n!}{k!(n-k)!}$$

Alternative orderings include hierarchical coupling of primitive behaviours by adding temporal ordering. Alternatively, behaviours may be weighted according to a learning algorithm. If temporal ordering is included through hierarchical coupling, the number of potential emergent behaviours will be increased. The behavioural modules are usually considered to be hierarchically organised such that each instinct is decomposed into finer grained behaviours. Instincts at the same hierarchical level compete with each other. Thresholding behaviours eliminates the possibility of chaotic behavioural responses. Spatially distributed ordering introduces the potential for cyclic behaviour due to the deterministic nature of mobile response and the inability to respond to events over several time periods. This is a consequence of a lack of memory which limits response flexibility when a reaction to an input may be dependent on previous inputs. Memory provides the capability of reacting to events over a number of time intervals and so alter behaviour on the basis of previous behaviour.

Brooks (1986, 1989) introduced a subsumption architecture of hierarchical behaviour based units to provide robust and reflexive robot performance with a precedence organisation to avoid multiple activations of mutually exclusive behaviours:

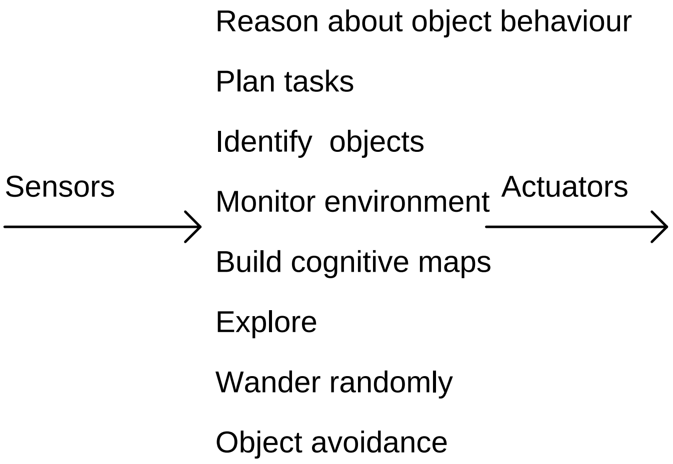

Figure 94 – Subsumption architecture task decomposition (Brooks 1986, 1989)

The classical control sequence is complex and requires extensive computation:

> `SENSORS → PERCEPTION → WORLD MODELLING → PLANNING → TASK EXECUTION → MOTOR CONTROL → ACTUATION`

Behavioural control involves dividing up the control system into modular task-based behaviours which are built up from basic locomotion to increasing levels of competence such as obstacle avoidance in incremental fashion. Each task-based behaviour module encapsulates its own perception and actuation functions – each behaviour module requires only task-relevant data to function. The behaviour-based control sequence is based on behavioural competences as complete modules: The problem is decomposed vertically into a number of levels of task competence based on behavioural units rather than functional modules. The behaviours as functional units comprise of well-define control laws which are goal-oriented. Each task module is connected directly to the outside world via sensors and actuators and operate in parallel. Each behaviour represents an active processing sequence from sensor input to actuator output. Sensor fusion is effectively replaced by command fusion. To eliminate conflicts, the subsumption architecture is adopted so that higher hierarchical level behaviours can suppress lower level behaviours so only one layer controls the robot at any one time. The subsumption architecture is a common example of behaviour-based control whereby task-based behaviours operate asynchronously. The control system is built bottom up incrementally and each level of competence includes the lower level as a subset. Simple tasks are solved first with new layers of competence added for greater complexity. This is an incremental approach. Individual layers can work on individual goals concurrently. Each level is implemented and debugged from the bottom up analogous to the evolutionary process of adding layers of greater behavioural complexity to animal brains to provide greater capabilities. Each level specifies a behavioural pattern which directly links perception to action and are summarised in Table 16.

Table 16 – Levels of Behaviour (Brooks 1990)

Level	**Module**	**Effect**
Level 0	Collide, Runaway	remains stationary until a moving obstacle approaches it
Level 1	Wander Randomly Object Avoidance	generate new headings periodically accepts a force vector input from level 0 and suppresses

		output from Runaway module
Level 2	Explore	finds a corridor to a specified goal at a distance and moves towards it (likened to exploratory behaviour for the provision of information)
Level 3	Build Cognitive Maps	plan routes between landmarks
Level 4	Monitory Environment	monitors dynamic changes in the environment
Level 5	Identify Objects	identifies objects in the environment and reasons about tasks to be performed on them
Level 6	Plan Tasks	formulates plans to change the state of the world as required
Level 7	Reason About Object Behaviour.	reasons about object behaviour in the world and modifies plans accordingly

Each level runs on its own processor and each processor represent a finite state machine running completely asynchronously with no shared global memory or central control. Each module represents a layer directly connecting perception to action operating asynchronously in a data driven strategy. Each module abstracts only those aspects of the environment which are relevant to their operation (egocentric invariants). Each module has input and output lines for message passing. The output line from one module is connected to input lines of one or more other modules. Each level examines data from the lower level. It is permitted to inject data (advice) into the lower level to suppress the normal data flow output from the lower level by actively interfering and issuing replacement data. This generates conflict resolution. There is no overall representation of the world. Complexity of behaviour is a reflection of the environment rather than of the agent and the state of the world completely determines the actions of the autonomous agent [Brooks 1990]. As primitive behaviours operate concurrently, complex higher level but unplanned behavioural patterns emerge, attributable to none of the individual simple behaviours alone imparting reactive behaviour to dynamic changes in the environment.

Subsumption has been implemented as networks of message-passing augmented FSMs (finite state machines) with messages typically 8-bit numbers passed over local combinatorial circuits to generate outputs. As each module requires only limited sensory data, the control system can react rapidly without delays incurred by sensor fusion and world modelling. Originally, each behaviour was implemented as an augmented finite state machine (FSM) of registers and timers with output wires providing the means for message-passing. Some behaviours are connected directly to sensors and/or actuators or to other behaviours to suppress or not suppress each other, i.e. a higher level behaviour can inhibit a lower level behaviour for a period of time. There is some resemblance to a neural network so a variation of this scheme would allow the implementation of reinforcement learning. The 1 kg 35 cm long six-legged robot, Genghis, was developed around the subsumption architecture without central control for operation over rough terrain [Brooks & Flynn 1989]. Each leg was attached at the shoulder joint with 2 DOF, each driven by a model airplane position-controllable servo-motor. It had 4 onboard 8-bit processors (three for motor/sensor processing and one to coordinate the subsumption architecture), 12 actuators with force feedback, 6 forward-looking pyroelectric infrared sensors, 2 collision-sensitive whiskers, and pitch/roll inclinometers and was built in only 12 weeks by two people. It comprised of 57 AFSMs to control it – all AFSMs had sensor values as inputs and actuator commands as outputs. They were arbitrated by a fixed network involving no communication between AFSMs. As soon as a leg was raised, it automatically swings forward and then down. The act of swinging caused all the other legs to move backwards so that the body moved forward. The process cycles through each leg in turn. The control system was implemented in 8 modular task layers – Stand-up, Simple-walk, Force-balancing, Leg-lifting, Whiskers, Pitch-stabilisation, Prowling and Steered-prowling. Of the 57 AFSMs, 48 were organised as 6 complete copies of an 8-machine control system foe each leg, 2 for local behaviours connecting whiskers to the front legs, 2 associated with inhibitions of balance behaviours in the front and back leg pairs, 5 for "central" control (2 for walking, 1 for steering and 2 for tracking of objects). Genghis demonstrated that force control was essential for ensuring reliable performance in robotics [Ellery 2000].

Apart from the subsumption architecture which is essentially a fixed control hierarchy (it suffers from inadequate conflict resolution as it employs an all-or-nothing strategy rather than compromise), a number of alternative, more flexible approaches can be adopted – spreading activation throughout the behaviour network or neural learning networks. Rosenblatt & Payton (1989) introduced a variation on the subsumption architecture which adopted a finer-grained task-oriented behavioural suite organised in a connectionist-type, highly structured architecture. The behaviours are assigned into two layers – an obstacle-avoidance layer and a follow-path layer. Behaviours are taken as atomic functional elements representing a specific concept – each unit receives a multiple-weighted input from other units and computes a single activation level output. Each link is weighted according to its priority. Each ubit computes a weighted sum of votes received from each active behaviour 0<w<1. The sum is normalised by dividing by the sum of weights. The units with the highest activation are selected for command execution. The system implements a gradient field to avoid obstacles and pursue goals with commands being based on compromise arbitration. Rosenblatt (1997) introduced a set of behaviours that compute the utility of world states which were combined to produce maximum expected utility of action a:

$$U(a) = \sum_{c} U_c . p(c \mid a, e)$$

where p(c|a,e) is the probability the consequence c will occur given the observed evidence e and adoption of action a. This allows optimal action selection through the use of Bayes theorem to reason about uncertainty.

Complex behaviours can be generated by the interaction of simple behaviours without the need for world modelling. However, some world map construction such as relational maps is necessary as the behaviour of the robot is only reactive. Purely reactive methods offer minimal computational overheads but they are not capable of sophisticated behaviour. Hybrid architectures offer a compromise by employing reactive control at lower hierarchical levels and planning at higher hierarchical levels of control. Payton et al (1990) stressed the need for the integration of high level planning with lower level reactive behaviours by suggesting a hierarchical approach as a development from Rosenblatt & Payton (1989). The lowest level was essentially behaviour-based with the higher level system being a map-based route planner. The plan is regarded as a resource for use in decision-making – one of these internalised plans was the gradient field imposed on a cellular map. The gradient field was subjected to the A* search algorithm to calculate the route cost. The gradient field was used as "advice" rather than as a unified representation. Other such resources could have included knowledge-based expert systems.

Balkenius (1994) has suggested the hierarchical integration of sets of structured neural network modules to incorporate behaviour control reflexivity with neural learning and a central controller. Such a central controller may emerge as the result of interactions between behavioural modules. The first layer is a subsumption-based reactive component, the second layer adds reinforcement learning and the third layer adds planning capability. The distributed set of neural network modules learn different situation-action tasks which interact forming a cognitive map of the environment. Reinforcement learning (based on ART2) generates consummatory behaviour (positive reinforcement) and avoidance behaviour (negative reinforcement). The central controller is essential to resolve conflicts between different behaviours in complex systems like intelligent robots – they provide motivational biases.

Thrun and his team have proposed a method for a mobile to accurately map its environment. They have developed a system where a human operator tele-operates the robot through the environment and selects a number of significant places of interest through a pushbutton command. Specifically this paper addresses situations where landmarks might be entirely indistinguishable, and where the accumulated odometric error might be enormous.

This navigation and localization idea is based on the Baum-Welch principle, which alternates an E-step and an M-step. The E-step (expectation step) the known map is kept held constant and the probability distributions are calculated for past and current robot locations. In the M-step (maximization or modification step), the most likely map is computed based on the estimation result of the E-step. This alternation allows the robot to simultaneously improve its localization and its map [Thrun 1998].

The robot using this principle must have the ability to observe landmarks and has the ability to estimate type, relative angle, and approximate distance of these landmarks, such as Choset's "meet points" [Choset 1996].

Genetic programming techniques have similarly evolved the subsumption architecture [Koza 1993]. Other techniques such as the potential field schema approach provide the means for fusing behaviours simultaneously and directly [Arkin 1987, 1990] – the potential field also provides a form of cognitive map and represents a particularly powerful representation schema for servo-control and high level navigation by combining behaviours at each level in a sensor fusion-type process independent of the sensory modality [Giszter 2002]. Potential fields are also candidates for animal foraging search strategies [Gelenbe et al 1997]. The potential field approach to local path planning and reactive behaviours generates a direction-based vector based on the negative gradient of the potential field, $F = -\nabla U$. Obstacles exert repulsive forces $F_o = \sum F_{oi}$ while the goal exerts an attractive force F_g. Generally, a minimum distance is defined between the agent and the obstacle. The resultant force $F=F_o + F_g$ determines the motion of the robot. In a cluttered environment, the robot may become trapped in local minima when attractive and repulsive forces become equally balanced. This may be avoided by using a grid solver that determines a potential field which satisfies the Laplace equation. A generalised potential field is a function of both position and velocity of the robot. This gives a smoother path. Forces are treated as real forces so that they can be integrated directly into the servo-control system. A polygonal robot M at position x between start position x_s and goal position x_g in a 2D workspace in the presence of polygonal obstacles O may be modelled through the basic potential field algorithm:

```
$x \leftarrow x_s$
repeat
       $F_{repulse} \leftarrow (0\ 0)$
       for each $o \in O$ do
                 $\Delta_m \leftarrow \Delta(M,o)$
                 $F_{repulse} \leftarrow F_{repulse} + \hat{\Delta}_m / |\Delta_m|^2$
        endfor
        $F_{attract} \leftarrow \widehat{x_g - x}$
        $F_{result} \leftarrow F_{attract} + aF_{repulse}$.
        $x \leftarrow x + sF_{result}$.
until $(x \approx x_g)$ or $(|F_{result}| \approx 0)$
```

The constant a controls the degree of repulsion typically ε^2 (where ε=minimal clearance between robot and any obstacle) while Δ_m defines the minimum distance from obstacles, and s defines the distance of movement per iteration. A typical mean obstacle proximity metric may be defined as:

$$l = \int_0^l (1/l_0)^2 .ds \Big/ n$$

where:

$$l_0 = \min[V_0 \sqrt{(x_v - x_0)^2 + (y_v - y_0)^2} \text{ =\textit{distance to closest obstacle}}$$

It is possible to add a rotational degree-of-freedom by computing the moment on the robot due to the potential field:

```
x←x_s
θ←θ_s
repeat
        F_repulse ← (0 0)
        moment ← 0
        for each p ∈ P do
                for each o∈O do
                        Δ_m ← Δ(c+p,o)
                        F_p ← $\hat{\Delta}_m / |\Delta_m|^2$
                        F_repulse ← F_repulse + (p x F_p).k
                endfor
        endfor
        F_attract ← $\widehat{x_g - x}$
        F_result ← F_attract + aF_repulse.
        x ← x + sF_result.
        θ ← θ + b * moment
until (x≈x_g) or (|F_result|≈0)
```

The point c denotes the centre of mass of the robot, P defines points on the boundary of the robot M relative to c. The constant b controls the degree of rotation determined by the moment of inertia of the robot. The motion of the vehicle is given by:

$$\dot{x} = v_r \cos\theta$$

$$\dot{y} = v_r \cos\theta$$

$$\dot{\theta} = \tfrac{v_r}{L} \tan\phi$$

where:

$\phi = k\ x\ m$ = *steering angle*
m = *moment* = $m + (p\ x\ F_p).k$

Now, (x, y, θ) is the location of the midpoint of the rear axle.

$$-\dot{x}\sin\theta + \dot{y}\cos\theta = 0$$

If tanϕ is constant, the equation is integrable. The resultant force is divided into two components, one in the direction of vehicle heading to control the speed and one perpendicular to it to control the steering angle.

$$\theta_n = \theta_{n-1} + \tfrac{v_r t(\tan\phi)}{L} = \theta_{n-1} \tfrac{\delta(\tan\phi)}{L} \text{ where } \delta = v_r t$$

$$x_n = x_{n-1} + (L/\tan\phi)(\sin\theta_n - \sin\theta_{n-1})$$

$$y_n = y_{n-1} - (L/\tan\phi)(\cos\theta_n - \cos\theta_{n-1})$$

The most significant problem with the potential field is that local minima occur where the solution becomes trapped. It is then necessary to impose a global approach to overcome local minimum, or alternatively to add a probabilistic Monte Carlo element, or techniques such as simulated annealing.

Behaviour-based approaches to reflexive control and obstacle avoidance utilise algorithms based on biological analogies – indeed, Arkin (1987, 1990) has implemented a form of behaviour control within the framework of schema theory from psychology – the Piaget theory of cognitive development is based on schemas as cognitive structures. Arkin (1987) used a form of behaviour based control whereby the primitive behaviours were implemented as multiple concurrent and independent motor schemas capable of communicating and coping with conflicting data – this AuRA (Autonomous Robot Architecture) adopts motor schema which allow encapsulation of the perception-action process within a framework that is compatible with top-down expectation/attention and bottom-up reactive behaviour mechanisms. Schema theory provides the basis for the integration of action, perception and cognition with compatibility with potential field approaches. The motor schemas act as a dynamic network of active behavioural networks which change as perceptual sensing of the environment changes. Reactive navigation of a mobile robot emerged through the instantiation of these motor schemas. The schemas were generic specifications of appropriate patterns of behavioural actions codified in an organised data structure. The primitive schema behaviours may be combined to yield more complex behaviours. Schemas are invoked by sensor data and they provide expectations for the appropriate motor actions. Each schema outputs a potential force field vector:

Avoid obstacle schema:

$$O = \left(\begin{array}{ll} 0 & d > S \\ \frac{S-d}{S-r} \times G & \text{for } r < d \le S \\ \infty & d \le r \end{array} \right.$$

where:

S=sphere of influence of obstacle
r=radius of obstacle
G=gain
d=distance of robot from obstacle centre

Path following schema:

$$V = \begin{cases} P & \text{for } d > w/2 \\ \frac{d}{w/2} \times G & d \leq w/2 \end{cases}$$

where:

w=path width
P=off-path gain
G=on-path gain
d=distance of robot from path centre

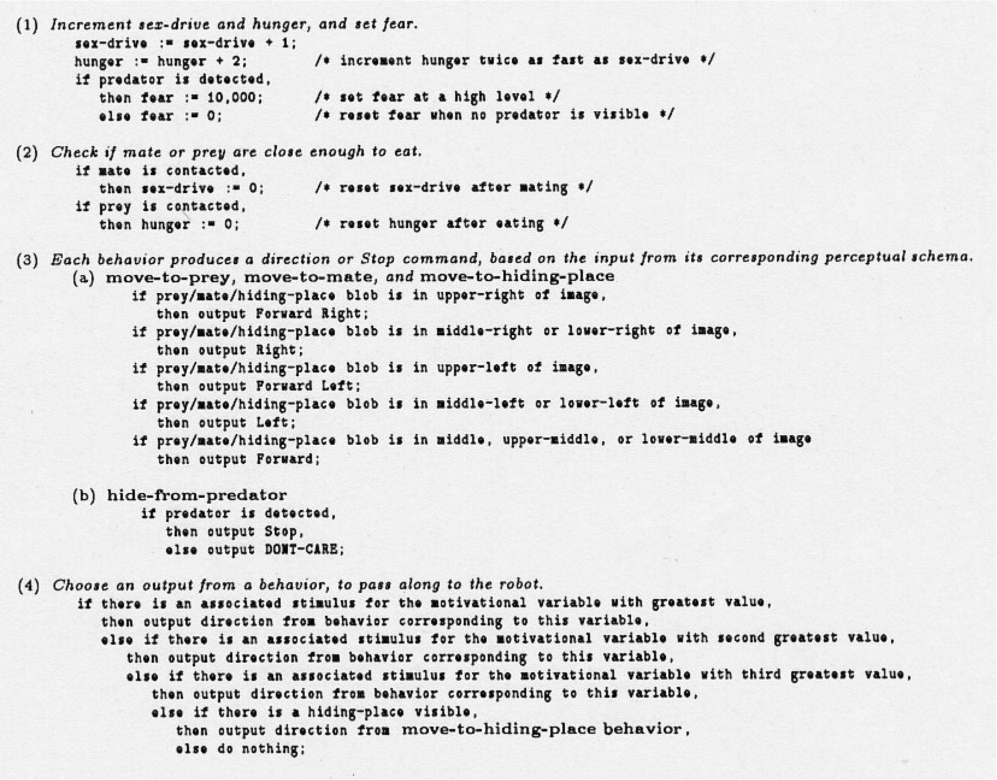

```
(1) Increment sex-drive and hunger, and set fear.
      sex-drive := sex-drive + 1;
      hunger := hunger + 2;        /* increment hunger twice as fast as sex-drive */
      if predator is detected,
         then fear := 10,000;      /* set fear at a high level */
         else fear := 0;           /* reset fear when no predator is visible */

(2) Check if mate or prey are close enough to eat.
      if mate is contacted,
         then sex-drive := 0;      /* reset sex-drive after mating */
      if prey is contacted,
         then hunger := 0;         /* reset hunger after eating */

(3) Each behavior produces a direction or Stop command, based on the input from its corresponding perceptual schema.
      (a) move-to-prey, move-to-mate, and move-to-hiding-place
             if prey/mate/hiding-place blob is in upper-right of image,
                then output Forward Right;
             if prey/mate/hiding-place blob is in middle-right or lower-right of image,
                then output Right;
             if prey/mate/hiding-place blob is in upper-left of image,
                then output Forward Left;
             if prey/mate/hiding-place blob is in middle-left or lower-left of image,
                then output Left;
             if prey/mate/hiding-place blob is in middle, upper-middle, or lower-middle of image
                then output Forward;

      (b) hide-from-predator
              if predator is detected,
                 then output Stop,
                 else output DONT-CARE;

(4) Choose an output from a behavior, to pass along to the robot.
      if there is an associated stimulus for the motivational variable with greatest value,
         then output direction from behavior corresponding to this variable,
         else if there is an associated stimulus for the motivational variable with second greatest value,
            then output direction from behavior corresponding to this variable,
            else if there is an associated stimulus for the motivational variable with third greatest value,
               then output direction from behavior corresponding to this variable,
               else if there is a hiding-place visible,
                  then output direction from move-to-hiding-place behavior,
                  else do nothing;
```

Figure 95 – Schema Algorithm (Arkin 2000)

Each schema had an activation level which acted as a threshold for instantiation and this activation level may be controlled by altering the gains. Each schema is independent and they are activated in parallel without layering and their potential field outputs are summed. Summing requires communication to resolve conflicting forces and this may be accomplished through a blackboard. This essentially provided a world model/cognitive map. Local minima may be overcome using a background noise component to produce a low magnitude random force vector to remove the robot from undesirable equilibrium points. The noise schema in conjunction with the Avoid-Obstacles schema generates exploratory behaviour.

These motor schemas drive the robot to interact with its environment to satisfy the goals generated from a planning system.

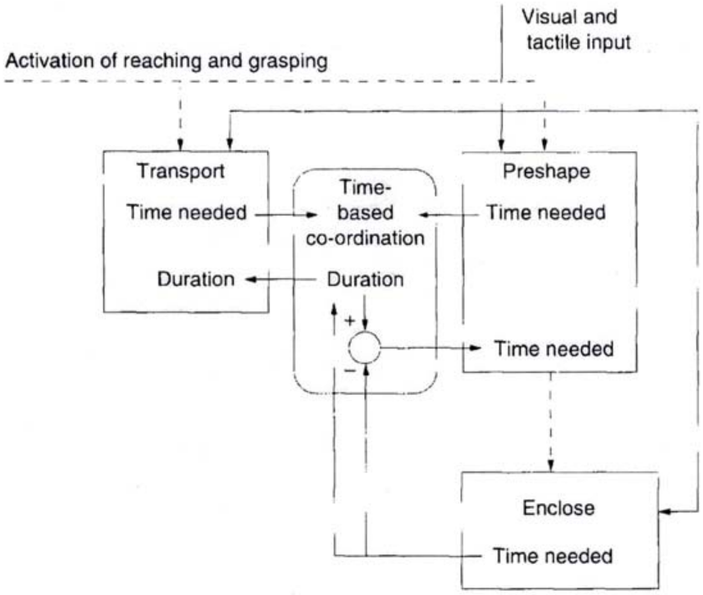

Figure 96 – Motor schema model with timing coordination for reaching and grasping – overall schema combines three motor schemas and a coordinating schema for arm (gross motion) and hand (preshape and enclose) which defines time trajectories (Jeannerod et al 1995)

The behaviour control technique requires pre-programming of the appropriate behaviours and such behaviours are generally of limited sophistication suggesting a lack of scalability. However, schema theory allows integration of both learning methods and cognitive map construction within the architecture. To ignore representation altogether is to limit autonomous behaviour to current sensory data without any consideration of past history – this historical component to cognition is provided by learning and memory. The lack of "memory" in the behaviour-based methodology limits the sophistication of such techniques but are sufficient to mimic many aspects of insect behaviour. Learning provides the means to escape this limitation. There does appear to be some resemblances between motor schema as encapsulated behaviour units and "frames" from cognitive psychology which encapsulate stereotypical cognitive processes (see later).

The navigation behaviours that have been investigated thus far have been extremely simplistic which rather contradicts the importance of environmental interaction in animat-based robotics. There is no reason why these techniques cannot be applied to more complex environments representative of extraterrestrial environments.

Biological navigation exploits cognitive maps stored in the hippocampus – indeed, hippocampal cells in animals store landmarks. Working memory acts as a buffer to hold the current position relative to these landmarks. Several such biomimetic approaches are described below:

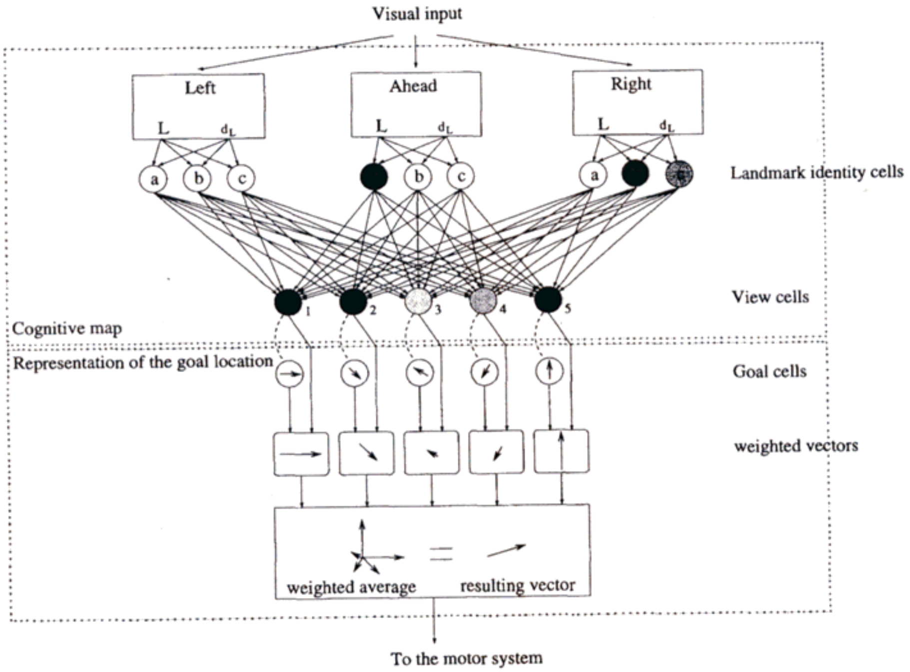

Figure 97 – Zipser's biological navigation system – Each landmark is seen left, ahead or right relative to head direction for which landmark identity cells fire proportionally according to angular distance. Combination of landmark identity cells activates view cells. Each view cell is associated with direction to the goal and corresponding vector is represented by goal cell (Meyer 1997).

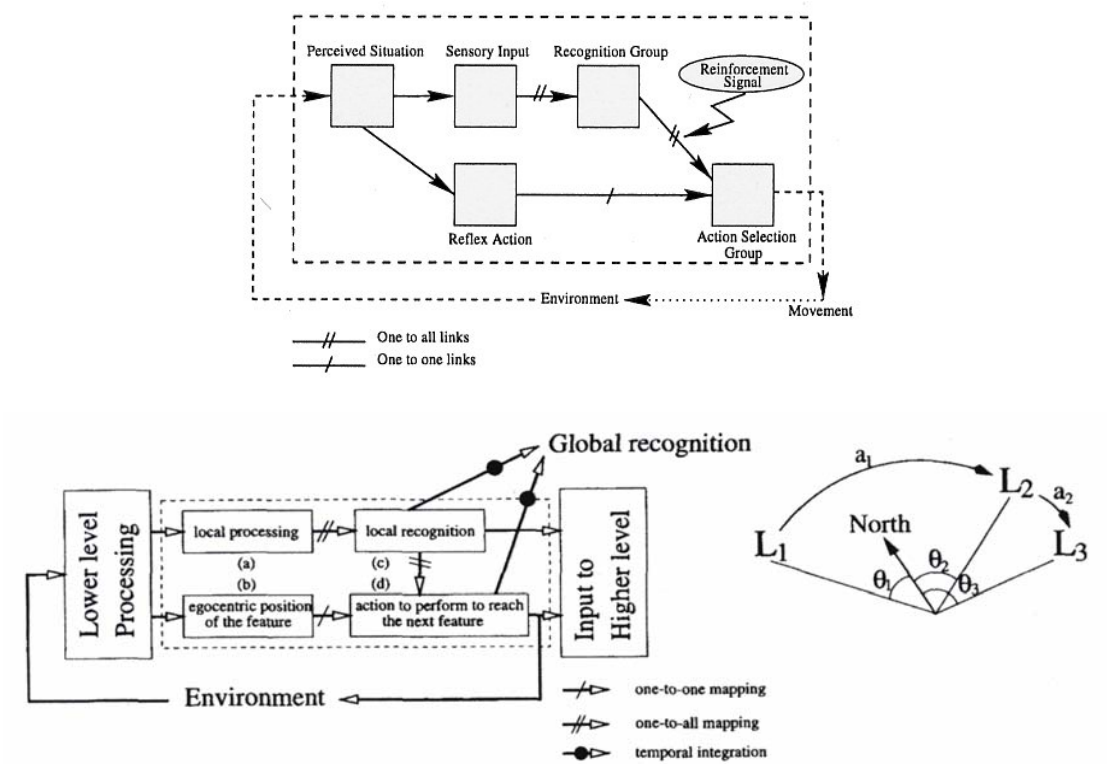

Figure 98 – PerAc navigation module – Each neural net module recognises specific features and landmarks extracted through vision processing. Angular position of these features are recorded and ocular or locomotor actions are generated to orient from the current state to the next feature (Gaussier et al 2000).

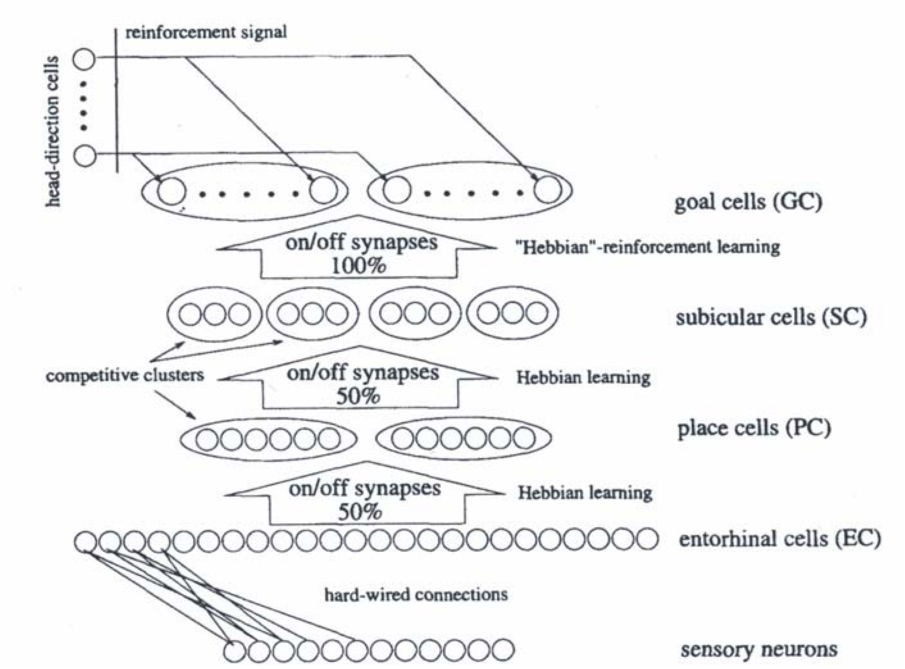

Figure 99 – Burgess' feedforward network based on the hippocampal model – Place cells are clustered to model head direction within goal cells (Trullier et al. 1997)

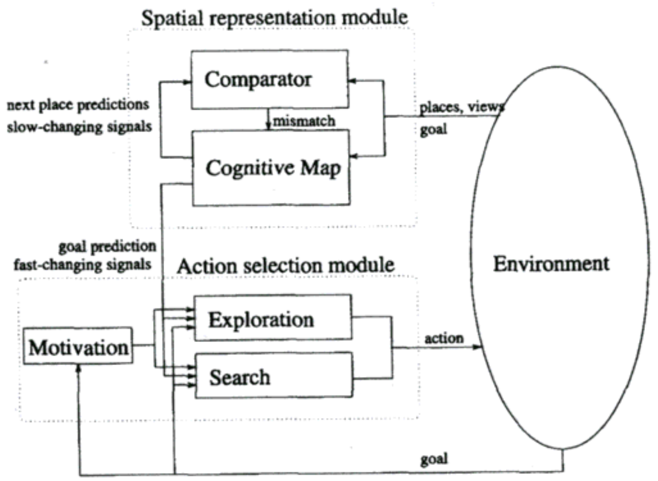

Figure 100 – Schmajuk & Thieme's modular navigation model – the first module encodes topological relationships between places by comparing them to internal predictions based on view. Second module uses predictions to select movements to goal (Trullier et al 1997)

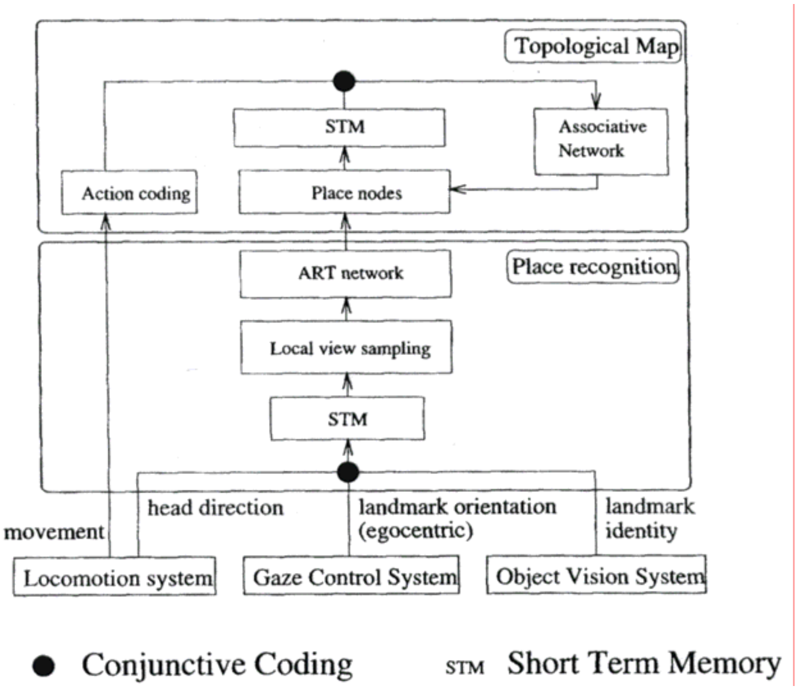

Figure 101 – Bachelder & Waxman's navigation model – Landmarks are identified by current view, relative landmark orientations and head direction within place recognition modules, and history of movement within toplogical map representation (Trullier et al. 1997)

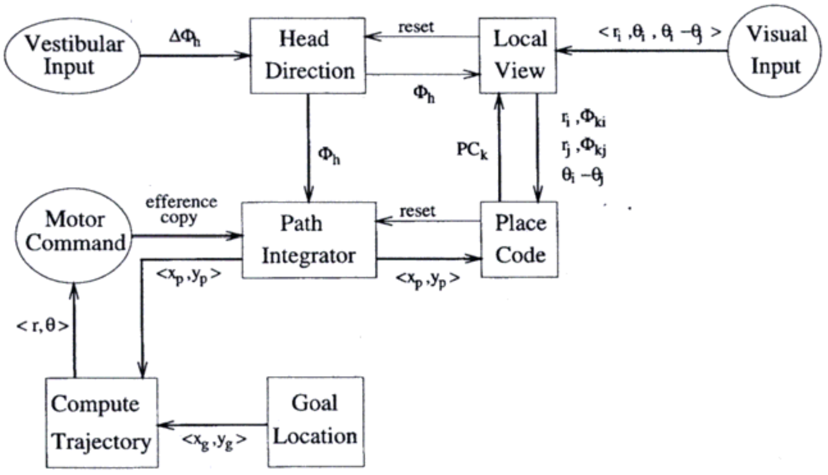

Figure 102 – Wan's navigation model combines multi-modal data – Place is represented as the conjunction of multi-modal inputs which couples path integration, vision, and head direction. (Trullier et al. 1997)

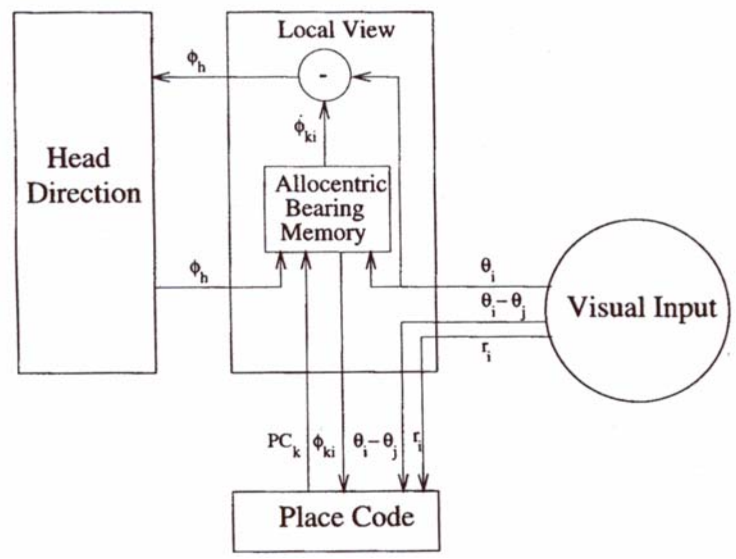

Figure 103 – Wan's model includes animal learning whereby the animat learns the relationships between egocentric and allocentric bearings of specific cues as viewed from a number of locations. When the animat has located itself and defined its goal, it computes the direction and distance to the goal for motor output (Meyer 1997)

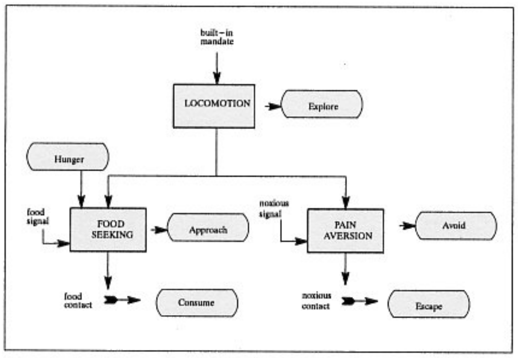

Figure 104 – Booker's animat control architecture – Each instinctive innate centre operates through a stimulus releasing mechanism and an associated action influenced by current motivational state and other instinctive centres higher up the control hierarchy. Food seeking and pain-aversion centres are located at the same hierarchical level (Meyer 1997)

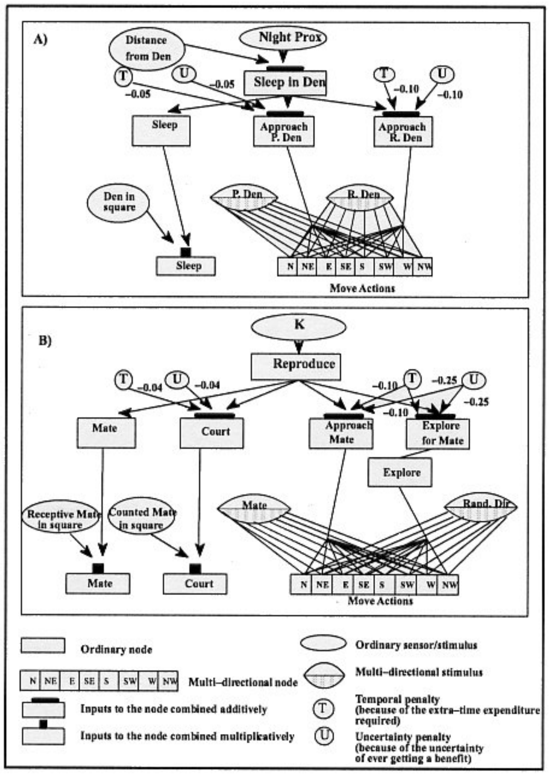

Figure 105 – Tyrell's free flow hierarchies which mediate between sleep-in-den and reproductive behavior (Meyer 1997)

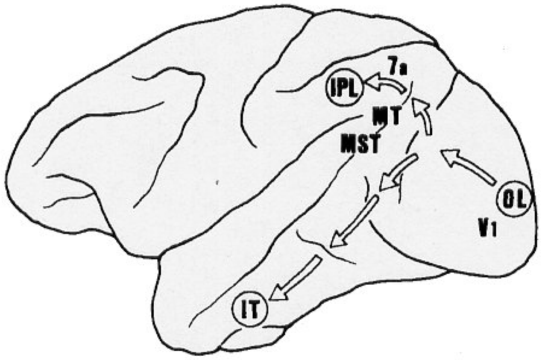

Figure 106 – Two cortical visual pathways of the human brain – dorsal pathway from occipital lobe to inferior parietal lobe mediates visual object location (navigation) while ventral pathway from occiptal lobe to inferotemporal cortex mediates object recognition (identification).

3.9. MEMORY – BEYOND BEHAVIOUR CONTROL

The CAM-brain (cellular automata machine) is a hardware-based electronic neural network of 38 million artificial neurons on 72 FPGA (field-programmable gate array). Each of 3276 modules of 1152 neurons is trained to perform a specific task using genetic algorithms to evolve connections within the module. Control genes are evolved to generate inter-module interconnections. This vast complex project is an attemot to model aspects of the human brain.

Brain tissue is of two main types: grey matter and white matter [Nauta & Feirtag 1979]. Grey matter consists of tight aggregations of neurons while white matter comprises mostly of nerve cell axons for long distance communication. . The brain appears to be constructed from a large number of distributed local but overlapping associative nets, some in parallel and some in series [Crick 1979]. The brain has a modular architecture of specialised components. The mammalian brain broadly consists of the prosencephalon (forebrain), the mesencephalon (midbrain) and the rhombencephalon (hinbrain). The prosencephalon consists of left and right cerebral hemispheres (telencephalon) connected by the diencephalon. The frontal 2/3 comprises the hypothalamus, the rest being the thalamus. The hypothalamus consists of the pituitary galan appendage and the septum. The telencephalon in the mammalian brain is the largest part of the forebrain - its surface is covered with the cerebral cortex which is heavily convoluted. The cortex is divided into several regions. The first region at the base of the hemisphere is the olfactory bulb and olfactory cortex. The edge of the cortex is the hippocampus. The rear section of the cortex analyses, codes and stores information. The latest evolutionary invention is the neocortex containing 70% of all neurons in the central nervous system in Man. The frontal lobes are responsible for the formation of intentions and activation of attention, ie. frontal lobes are not involved in sensory processing or motor functions. A simplified human brain architecture may be represented thus:

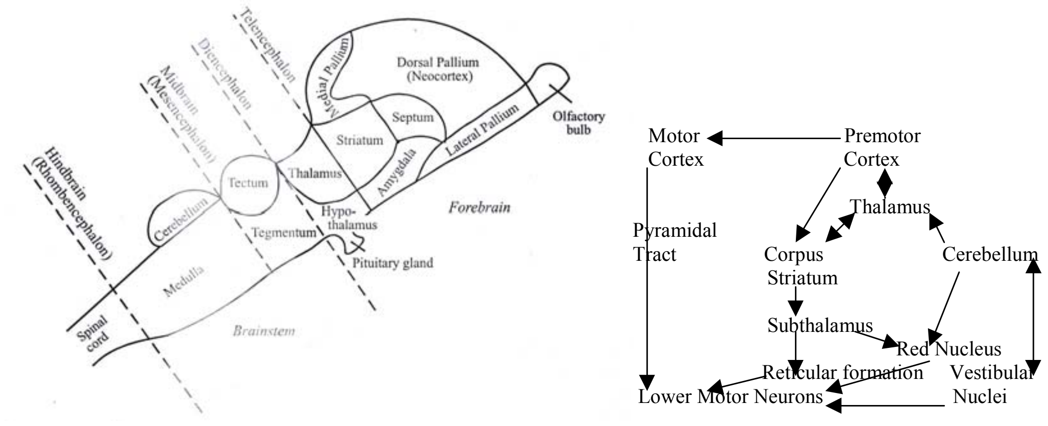

Figure 107 – Simplified Vertebrate Brain Architecture (a) anatomy (b) pathways (Prescott et al. 1999)

Mammals have two different types of memory which use two different brain structures for encoding information. The brainstem is associated with reflex behaviour and lower level learning. Memory traces in the human brain are localised to the cerebellum, hippocampus, amygdala and cerebral cortex. Working memory is associated with the prefontal cortex, particularly the principle sulcus, which accesses relevant stored information in the cerebral cortex – working memory provides moment-to-moment awareness and retrieves archived information (similar to a "blackboard"). Long term memory storage occurs in the cerebral cortex which is also the region responsible for higher cognitive processes in mammals. The temporal lobes on the inner surface of the underside including the hippocampus are responsible for the storage of new long term memories. The hippocampus comprises of pyramidal cells and lies in the temporal lobe of the cerebrum and is the origin of working memory. The hippocampus is associated with semantic/episodic information storage. Procedural knowledge is learned by repetition and is associated with the cerebellum – it is learned by Hebb's rule through association between presynaptic and postsynaptic activity. Declarative knowledge involves association of simultaneous stimuli about single events whereas procedural knowledge involves the association of related sequential stimuli. Declarative knowledge is learned by altering the connection strength between presynaptic activity and a third modulatopry neuron without post synaptic activity. The hippocampus is the temporary store with a characteristic timescale of weeks or months prior to transfer to the permanent store of the cerebral cortex [Kandel et al. 1992]. It has a regular structure and its principal region comprises a sheet of output neurons or pyramidal cells with interneurons of mostly stellate neurons. Pyramidal cells are analogous to Purkinje cells and interneurons are analogous to granule, Golgi, basket and stellate cells in the cerebellum. Hence, it is an associative memory like the cerebellum. The hippocampus learns its internal classification of input patterns for which it uses the granule cells with modifiable synapses. The hippocampus appears to be responsible for spatial memory of objects in the environment. It participates in the formation and retrieval of semantic memory and the cerebral cortex is associated with storage [Lynch & Baudry 1984]. In fact, the neocortex as a whole operates similarly to the hippocampus the neocortex comprises 70% of all the neurons of the human brain. The neocortex is dominated by pyramidal and stellate cells. Neocortical nets form internal representations of classes of objects with climbing fibres activating the classification process through long term potentiation. Memory is characterised by long term potentiation or persistent postsynaptic potentials following cortical and hippocampal activity. Objects are represented in memory by association of rhythmic and synchronised firing which bind activity in different parts of the cortex. The striate cortex (primary visual) cortex is essential for learning and memory of visual patterns. There is an area on both sides of the cerebral cortex – the inferior temporal sulcus - but dominant in the right hemisphere which is concerned with the recognition of faces. This region lies on the underside of both occipital lobes extending forward to the inner surface of the temporal lobes. The reason for such a highly specialised function being donated such a large region of the brain is that recognition of individuals lies at the heart of human sociability.

The memory system buffers the sensory-response system. Information acquisition initially comprises of pattern recognition with the stimulus coded in terms of features and their identification through classification. This involves matching with prototypes in LTM. Stored information is used to filter data to focus on patterns and to detect deviations from expectation. Humans remember and value information which confirms their previously held beliefs and expectations than they do to disconfirming information. Perception is directed according to their present and past experiences and this provides the basis for selective filtering. Human information processing and decision-making involves memory. Human memory comprises 3 major components: sensory memory, short term memory (STM) and long term memory (LTM). Sensory stimuli from the environment enter sensory memory by parallel excitation and undergo a process of filtering. Sensory memory can discriminate ~0.1s durations and is an unconscious process involving feature detection and pattern recognition. Short term memory (STM) is highly restricted in capacity being limited to 7±2 separate items or chunks. A chunk is a re-ordering and re-integration of items into an organised groups to increase the capacity of STM. STM provides immediate recall of actively rehearsed but limited information. It controls the focus of attention and operates serially on the filtered and preprocessed data in STM. Unless active rehearsal is consciously instigated, the limit of its duration being ~30s. One such rehearsal process is the articulatory rehearsal loop. STM provides active control of reasoning and memory processes. STM may be regarded as the activated portion of LTM and both are hierarchically organised into different levels of information reduction. STM acts as a selective filter to reduce the amount of information for processing and provide a scratchpad for working memory processes. The transfer of information from STM to LTM depends on attention limitations, interference of data and the formation of associations in LTM. Feature integration theory of attention suggests that sensory features are registered in parallel early in processing automatically and objects are identified separately at a later stage through focussed attention [Treisman & Gelade 1980]. Attention integrates separate features into unitary objects. Attention is necessary for the detection of targets defined by the conjunction of properties. Separable features are detected by parallel search. Conjunction requires serial search and provides identification. Information is lost from STM unless it is rehearsed by rote repeating. If sufficiently rehearsed for 5-10s the information trace may enter LTM. In the human brain the hippocampus consolidates STM information into LTM. Human long term memory (LTM) has a linked list structure and although of very large capacity must be finite (no known limit). This provides the basis for holistic reasoning such as reasoning by analogy. LTM suffers from slow decay rather than the rapid decay characteristic of STM. Retrieval from LTM is by parallel associative search triggered by STM contents. Information about the environment is then internalised to reflect the properties of the environment through three modes [Shepard 1984]:

(i) STM to reflect the current context;

(ii) LTM to reflect past contexts experienced by individual organisms;

(iii) Genetic code to reflect species evolution to match the environment.

A mixture of methodologies is probably required to provide a reasonable emulation of human behaviour involving multiple modes of representation. Whereas declarative knowledge can be combined with other declarative statements to allow inferencing, procedural knowledge appears to be inaccessible (in humans) to conscious manipulation. Declarative knowledge comprises general and specific facts about the world while procedural knowledge defines how to perform actions in the world and mental strategies. Data is usually declarative while programs are usually procedural. Procedural knowledge is only accessible as far as their outputs (eg. skilled motor behaviour). This suggests a dual representation in human cognition. ACT* (Adaptive Control of Thought) is a generalised implementation of the semantic network as a general model of human knowledge [Anderson 1989]. It comprised three memories: working memory (STM), procedural memory and declarative memory which together comprise LTM.

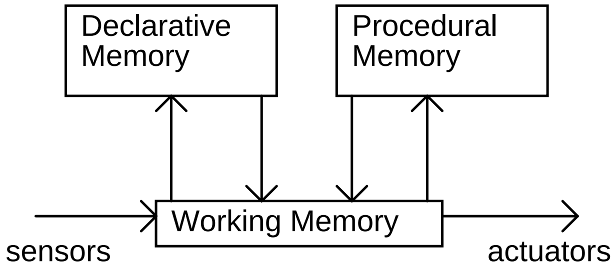

Figure 108 – LTM Diagram

Procedural memory is concerned with motor behaviour, motor sequences and motor skills (learned through association and reinforcement). Declarative memory is required for planning – it contains if-then expectations of the outcomes of possible actions (semantic memory learned through instrumental learning) and previous experiences (in episodic memory). All incoming information is interpreted in relation to knowledge already stored in memory and sensory information provides the cues to trigger the relevant information. Procedural memory comprises independent production rules which are invoked by matching the contents of working memory and activating actuation behaviour. Such procedural knowledge is not available to introspection and is unconscious. This memory uses GPS-type operations [Hunt 1989]. A separate declarative memory consists of factual knowledge organised as a semantic network of schemas. The schemas are cognitive structures comprising classes of action sequences which act as filters to organise sensory data. This distinction has its basis in that although declarative knowledge is accessible to introspective reporting procedural knowledge is not. All interaction occurs in working memory which provides the interface to the environment. This provides a means for both cognitively mediated processes and responsive behaviours generated by the state of the environment [Bandura 1971]. Declarative memory is accessed by spreading activation from working memory through the declarative semantic network of past knowledge and stored inferences. Working memory then comprises the portion of the knowledge structure that is currently active and activation depends on the perceptual input (to simulate attention focussing). Consciousness is identified with working memory which contains goals and perceptual inputs and provides an attentional focussing function. This is a contentious issue. Sperry (1969) was one of the first to elaborate consciousness (as a subjective mental phenomenon) as an emergent property of cerebral activity which in turn may exert a directive holistic (gestalt) form of control over cerebral excitation patterns. Conscious awareness is a dynamic emergent property of cerebral neural excitation different from the collected sum of specialised neurophysicochemical events. It is a higher level holistic property of a control hierarchy – the neural circuitry is generally genetic activated by the arousal system and are constantly changing in interaction. The integrated effect of lower control levels is higher level emergent and dynamic consciousness which in turn generates internal inputs in causally interacting with the lower levels of control, i.e. there is mutual interdependence between all neural levels of the cerebral cortex. Crook (1983) declared that consciousness is in fact unnecessary for intelligent behaviour and that consciousness is an empathetic process. It requires knowledge of other agents' experiences and in that way resembles the basis of speech acts. This view of consciousness is fundamentally a social one and that consciousness is identical with self-awareness and that this is learned through social interaction with the social environment and context. Consciousness as an empathetic process requires the agent to possess shared knowledge of other agents to reduce the emitted information content requirement of communication Declarative facts are activated whenever they are matched to the conditions of production rules. This pattern matching is processing intensive and serial representing conscious control. Knowledge compilation is the process of transferring declarative knowledge into procedural knowledge by embedding the declarative knowledge into production rules to simulate automated skills learning. As production rules are modular they can be added without affecting other rules in memory. This eliminates the need for pattern matching to declarative memory which is inherently computationally expensive. Learning provides the mechanism for increasing the agent's behavioural repertoire as a result of experience through the consolidation of STM into LTM. Multiple productions may be collapsed into single production rules in a process similar to chunking into macro-operators. In SOAR chunking occurred in working memory whereas in ACT* it occurs in the production memory. Inductive learning of new production rules occurs through generalisation/discrimination to classify from a number of examples presented to the system. This process is automatic and syntactic and applies to production rules only. The generalisation process involves the substitution of variables for constants which appear in both the condition and the action. Inductive learning by analogy is the process of learning used in declarative memory. This process is consciously controlled and applies to declarative networks only enabling learning from a single example. SOAR employed conflicts as a means for learning such that impasses defined new problem to be solved. In this way, problems are generated and stacked (subgoaling) and solved in turn. Each solution is a “chunk” of information encoded as new production rules.

Declarative memory may be organised into hierarchies of schemas representing concepts. Schank & Hunter (1985) advocate the use of schemas and MOPS to provide semantic knowledge associated with episodic memory. Schemas are collections of generic units of information with variable slots associated with generalised plans or behaviours. Scripts are schemas for frequently occurring sequences of events enabling temporal ordering of events to be represented [Schank 1972]. They specify general courses of action by incorporating a series of interrelated decision rules. They provide meaning to information and provide expectations and resolve ambiguities. Semantic structures may be represented as conceptual dependencies to represent meaning whereby slots are given for actions (Propel, move, ingest, expel, grasp) eg Mtrans (to transfer mental information), Ptrans (to alter location), Atrans (to transfer possession), Attend (direct sense organ to stimulus), Speak (to communicate), Mbuild (to create or combine thoughts) actors, objects, and directed components (from, to, on, etc). Each action has semantic content attached to permissible variables in each slot. Each sentence is stored as a case grammar:

```
Verb:  agent
       object
       instrument
       recipient
       direction
```

Schemas contain packaged and related information according to its relevance to a situation. Default knowledge is used to fill in gaps. Such schemas are nested such that some subschemas are shared between several schemas. Schemas are essentially frames which remain unchanged after performance of actions. MOPS are smaller finer grained units than schemas and scripts and form the building blocks of scripts. MOPs may appear in more than one schema. Each schema describes typical activities in terms of experienced primitives of semantics, procedural and episodic knowledge. Frames or schemas are organised as hierarchical directed labelled graph networks (semantic nets). Two conceptualisations may be related by consequence representing cause and effect enabling the representation of complex conceptualisations. Nodes signify elements and arcs signify relations. Demons monitor possible occurrences of likely consequences of a given situation and are activated by certain data or words. Wood's ATN can be used to grammatically parse sentences. Large amounts of knowledge may be structured as "chunks" or units of interrelated concepts applicable to particular events or situations and world knowledge. Such hierarchical structure allows the dynamic reconfiguration for behavioural diversity. All perceptual information is interpreted with respect to all ready existing knowledge in memory. Sensory information triggers the appropriate schemas.

Declarative memory is subdivided into episodic memory (experiential) and semantic memory (encyclopedic). Episodic memory is autobiographical and is self-referential while semantic memory is a mental lexicon stored without reference to personal events during which the facts were learned. The evidence for the separation of episodic and semantic memory comes from a brain condition called Korsakov's syndrome. It is caused by damage to the hippocampus which results from chronic alcoholism. The syndrome is manifest as a gross antegrade amnesia whereby the patient cannot recall events that occurred since the brain damage but all other mental functions are intact. They are capable of learning new motor and cognitive skills but cannot remember episodic events. The encoding of new episodic memories does not pass from working memory into long term memory but procedural and semantic long term memory encoding is unaffected. Another indication is hypnosis. Hypnosis is a social interaction whereby the subject is responsive to suggestions while altering the subject's conscious awareness. Perceptual processing continues normally and the subject can still learn new procedural and semantic knowledge though posthypnotic episodic memory retrieval is disruptive. Indeed as the key to consciousness is self-reference and identity whereby links must be forged between general mental representations and mental representations of the self as an agent this suggests that Korsakov sufferers possess static identities determined by episodic memory accumulated up to the point of brain damage. Stanfill & Waltz (1986) emphasised the importance of episodic memory which uses similarity based inductive learning based on pattern matching of data without the use of rules. Such approaches require the intensive use of memory in parallel. Accessing episodic memory for comparison with current data involves 4 basic operations:

(i) counting the number of times various features occur;

(ii) generating a similarity metric from the feature count;

(iii) computation of the dissimilarity of each item;

(iv) retrieving the best match.

The time required to produce such memory intensive reasoning is O(log n)^2. Similar past circumstances are used to make decisions and much of human reasoning is of this form.

The mind as a cognitive program would be characterised by:

- large size and complexity;
- multi-branched and holistic;
- formally incorrect;
- complexity of procedural interaction.

Charniak (1988) attempted to characterise the size of "mind" and the consequences that follow. The mind's dictionary is its lexicon of words estimated at 50,000 for a well educated person plus a similar size for concept names. The mind's encyclopaedia of basic factual common sense knowledge may be estimated at 100,000 similar to a large single-volume encyclopaedia. In distinguishing between fact and value, a goal structure of similar size is suggested. Episodic memory based on a 70y lifespan with a third of it sleeping and 10% retention of a single memory image per minutes generates a total of $2\text{x}10^6$ episodic memories. Overall this yields $2.3\text{x}10^6$ cognitive elements. The next stage concerns the perceptual and motor systems. Visual pattern storage may be estimated at 200,000 items as the dominant modality (a well educated Chinese person can recognise 50,000 Mandarin Chinese ideograms while an expert chess player can recognise 30,000 board situations). A similar number may be estimated for the other modalities combined. Motor skills are assumed to be patterned similarly to sensory patterns with 200,000 elements, yielding 600,000 sensorimotor categories in total. In total them, the human mind is perceived to store $3\text{x}10^6$ elements. The total brain volume is 1.4l of which 0.5l comprises the cerebral cortex. With an estimated 10^{10} neurons each with 10^3 synapses, there are 10^{13} synapses in total within which memory is believed to reside. Assuming that 90% of these synapses are dedicated to specialised functions, this leaves 10^{12} synapses available for the cognitive/perceptual/motor representations of 10^6 elements, i.e. 1 Mb per element (equivalent to 60 pages of alphanumeric text). The regularities of psychological structure being of cultural origin are not likely to correspond well with the regularities of the brain which is biologically determined. Comparison may be made with the original SDI program which was to be implemented in 10^7 lines of code. The SDI programme software was envisaged to be impossible to debug or verify. On average there is around one bugs per 100 statements in a program even after debugging and this average increases with program size – the SBM PC microcomputer operating system suffered from a number of non-trivial bugs. Rigorous proofs based on formal deductive logic are applicable to only small programs of 500 lines or less. Program verification is computationally complex and program consistency is relative rather than absolute. Numerical programs are particularly difficult involving floating point arithmetic generating round-off errors. Any large mind program that may or may not exist may not be verifiable as debuggability would require exhaustive checking of a combinatorially branching program. Defects will tend to propagate unless the program is hierarchically structured with cleanly interfacing modules. An evolutionary-built program with a messy structure would limit such modularity and verifiability. Human declarative reasoning may well be incorrect under all problem inputs but may be correct under a subset of identifiable problem inputs to yield acceptable solutions, i.e. to avoid intractability, reasoning processes are flawed (bugs) but probabilistically correct under likely conditions. It is conceivable that there may exist a program that yields human-like intelligent behaviour which is far simpler than the human mind, perhaps through exploiting symmetry principles and redundancy-resolution to reduce the information requirements for system specification.

3.10. BIOLOGICAL LEARNING

Simple animals such as insects behave through standard responses to stimuli. More complex vertebrates augment these responses with a world model which they use to select sensory information to plan courses of action to achieve its goals. The addition of learning provides the ability to adaptively use the world model to predict and compare predicted and actual outcomes of courses of action to update the knowledge structures for further use. The world model allows planning of future behaviours through simulation – simulation is reversible unlike action in the real world. Learning may thus be used to impart knowledge and skill to robot behaviours by accommodating the variability of the environment. Learning provides for greater complexities in behaviour beyond instinctive reflexivity. Learning is a much desired capability and may be considered to be an indispensable component of intelligent behaviour [Carbonell 1989; Wilkes 1953; Carbonell et al 1983]. A learning machine consists of a learning protocol specifying the way in which information is obtained and secondly, a deduction procedure which generates recognition algorithms for the learned concept in polynomial time [Valiant 1984]. The deduction procedure must in each case output a close approximation to the learned concepts such that it never outputs a positive when it should not but almost always outputs a positive when it should. Once learned the concepts should be represented in memory in some way by being integrated into the original knowledge base. To acquire new knowledge efficiently, the system must all ready possess a degree of prior knowledge. The creation of new symbolic categories through concept (symbolic) learning is a difficult process and differs from lower level association learning. Concept learning involves the ability to generalise to universals from particulars and reduce information redundancy. Conceptualisation relies on the agent's history of the system. Learning involves as irreversible internal changes in the system's structure that are adaptive such thatb they enable the system to do tasks drawn from the same population more efficiently at subsequent times according to some performance metric [Kolokouris 1986; Russell 1991]. There are two main forms of symbolic learning: inductive learning (by example) and analytic learning with the latter including explanation based learning and analogical learning. Traditional AI systems utilise deductive inferencing and do not learn inductively. The major limitations of these methods of learning is that they require the environment to be well behaved in that the planning agent is the only agent effecting changes in the world (closed world assumption). Since other agents do interact with the dynamic world this approach will generate uncertainty and errors. A further limitation is that learning through knowledge acquisition is dependent on the human programmer rather than by direct interaction with the world.

3.10.1. Learning Automata

Any system that alters with time may be modelled as an automaton in which time is modelled discretely and sequentially. Automata are abstract mathematical models of systems whose states are determined by their past histories, i.e. state transitions occur in response to inputs which in turn generate outputs. These state transitions occur on the basis of input stimuli which select the dynamics while the current state determines the effects of the input. The output depends on the current input, current internal state and the structure of state transitions. Finite-state automata have their histories stored in finite internal states as memories sufficient to predict the effects of the past on the future given the present input. Any system which processes information can be regarded as autopmata. A computing machine is an automaton which is a physical system whose dynamic evolution transforms its internal state from a set of input states to a set of output states. The output is a definite function of the input which is directly observable. The human brain is an example of a set of finite-state automata forming an interconnected array of 10^{11} two-state automata (neurons). Thought (learning, memory and processing) may be modelled as complex and numerous transitions from state to state [Aleksander 1973]. Although our capacity for thought is large, it is nonetheless finite at $2^{10^{11}}$ states assuming binary neurons. All finite state automata are defined by a finite number of internal states, and usually a finite number of inputs and outputs. The number of internal states, inputs and outputs are usually regarded as fixed and given the initial state and the input sequence, the automaton's behaviour is completely predictable (deterministic). All elements of the system are regarded as operating at a uniform rate (synchronous). Other forms of automata have been studied: non-deterministic automata where automata have a choice of possible actions at each state transition [Rabin & Scott 1959]; fuzzy automata in which imprecise fuzzy instructions have to be translated into more precise commands [Chang 1972; Datta 1985]; reversible automata with the ability to return to previous states [Bavel & Muller 1970]; and, asynchronous automata which have concurrent inputs generating partial ordering of sets of inputs [Shields 1985]. Cellular automata are the major tool for artificial life studies but are not considered here further. Kim (1989) proposed an extended, finite, deterministic automaton model for intelligent systems:

$$E=<S, f, \Sigma, \Omega, *, M, N, s_0, F, \sigma_0, w_0>$$

where:

S = *set of finite internal states*=$\{s_0, \ldots s_n\}$
f = *functional requirement representing goal states*
Σ = *finite set of inputs from environment*=$\{\sigma_0, \ldots \sigma_k\}$
Ω = *finite set of outputs to the environment*=$\{w_0, \ldots w_k\}$
* = *concatenation operators for compounded inputs/outputs*
M = *next-state transition function* $M{:}S\times\Sigma^*\rightarrow S$
N = *next-output function* $N{:}S\times\Sigma^*\rightarrow\Omega^*$
F = *final state*

M includes a high level condition-action decision rule-based knowledge system R, plans T and operators O. The output function N is concerned with automata self-reflective knowledge of procedures for low level control. The kernel is a blackboard encompassing F and data D from the input function M. The set of plans T determines how the knowledge base R is utilised by selecting operators O. The rule set embodies the input transition function and output transition function through the condition-action format. Kim (1989) stressed the importance of selectionist learning new rules such as classifier systems.

Learning automata have the characteristic ability of improving their behaviour over time in a manner similar to neural networks [McLaren 1966; Narendra & Thatchatar 1974]. They are modelled as stochastic systems with little or no a priori information. Stochastic automata have probabilistic state transition functions which are dependent on the input variables. Stochastic variable structure automata are able to update their action probabilities according to the input received from the environment to improve their performance. A random environment acts as a probabilistic teacher while learning essentially optimises an explicitly unknown functional. Learning automata initially have a choice of actions with equal probabilities of occurrence, one of which is selected at random, and the response of the environment alters the distribution of action probabilities through reinforcement algorithms thereby improving the automaton's performance recursively. The sequence of internal states forms a Markov chain – as in a Markov chain, the state of the system is determined by a probability rather than deterministically. The use of stochastic automata rather than deterministic automata reduces the number of states required. The stochastic automaton is a sextuple:

$$S = <\Sigma, S, \Omega, p, R, N>$$

where:

Σ = input from the environment
S = set of internal states $S_1 \ldots S_s$.
Ω = set of output actions, $\Omega_1 \ldots \Omega_r$ where $r<s$
p = state probability distribution determining the state choice at time t such that $\Sigma p_i=1$ for all t
R = reinforcement learning algorithm which generates $p(t+1)$ from $p(t)$ so $p(t+1)=R(p(t),S(t))$
N = next-output function (usually deterministic)

The automaton is embedded in a random stationary environment with feedback. The actions of the automaton Ω form the inputs to the environment and the responses of the environment Σ form the inputs to the automaton. If the environment responses are binary such that 0 is a non-penalty/reward response and 1 is a penalty response, we have a P model. Other models exist: e.g. if the automaton input is continuous in the range [0,1] we have an S model [Chandrasekaran & Shen 1968; Viswanathan & Narendra 1973]. This is operant conditioning type learning. The probability of an environmental response due to an input from the automaton is initially unknown, i.e. there is no information available about the environment initially, and $p_i(t)=1/N$. The penalty probability set $c_i= (c_1 \ldots . c_r)$ characterises the environment and is initially unknown; $(1-c_i)$ is the probability that a reward or non-penalty response is evoked, causing reinforcement of the chosen actions. The stochastic automaton learns by updating its action probabilities p in response to inputs from the environment to improve its performance by increasing the probabilities of the correct responses. The average penalty is defined as:

$$R(t) =< \Sigma(t) \mid p(t) >= \sum^{r} p_i(t).c_i$$

Learning occurs if the average penalty is lower than a random penalty (when actions are chosen with equal probability) and learning consists of reducing the average penalty, i.e.

$$< R(t) > \leq M_o$$

where:

$$M_o = \frac{c_1 + \ldots + c_r}{r} \text{ = average random penalty}$$

Optimisation consists of minimising the average penalty: $< R(t) >= \min\{c_i\}$

In fact, sub-optimal performance is more realistic (ε-optimality) such that $< R(t) > \leq \min\{c_i\} + \varepsilon$ [Najim & Oppenheim 1991]. Indeed, ε-optimality are strongly convergent in stationary random environments. Reinforcement schemes update action probabilities: $p(t+1) = R(p(t), \Omega(t), \Sigma(t))$

The greatest probability is assigned to the state with the smallest penalty probability. If an action at time t results in a reward input, the action probability $p_i(t)$ is increased at the expense of other action probabilities $p_j(t)$ where $i \neq j$. If a penalty input results, $p_i(t)$ is decreased with other components $p_j(t)$ increasing. If the action probabilities remain unaltered rather than decreased, this is inaction rather penalty. For an action Ω at time t [Narendra & Thatchacher 1974]:

$p_j(t+1)=p_j(t) - f_j(p(t))$ for $\Sigma(t)=0$

$p_i(t+1)=p_i(t) + \Sigma f_j(p(t))$

$p_j(t+1)=p_j(t) + g_j(p(t))$ for $\Sigma(t)=1$

$p_i(t+1)=p_i(t) - \Sigma g_j(p(t))$

There are a number learning schemes. The linear reward-penalty scheme (L_{RP}) is optimal is all stationary environments:

$f_j(p)=ap_j$

$g_j(p)=bp_j + (b/r) - 1$ where $0<a, b<1$

The linear reward-inaction (L_{RI}) scheme ignores penalty inputs from the environment so that action probabilities are unchanged in those cases (benevolent automaton):

$f_j(p)=ap_j$

$g_j(p)=0$

L_{RI} reduces the number of stages during which learning occurs but reduces the speed of convergence [Shapiro & Narendra 1969]. It is ε-optimal in all stationary environments. Genetic algorithms are similar to the L_{RI} learning procedures in that candidate solutions that survive into the next generation correspond to reinforcement and those that are discarded corresponded to zero reinforcement. L_{RI} procedure may also be used in artificial neural networks such that each connection is treated as a probabilistic switch – open indicates input 0 and closed indicates input 1. Learning alters the probability distributions of open/closed to maximise the expected reinforcement signal. L_{RI} computes the reinforcement signal to the switch configuration and the switch probabilities updated accordingly. The symmetric L_{RP} scheme is ε-optimal in restricted random environments such that $c_1 < 1/2$ or $c_2 < 1/2 < c_1$:

$g_j(p)=f_i(p)=ap_j(1-p_j)$

One non-linear learning scheme is represented by:

$f_j(p) = ap_j^2(1-p_j)$

$g_j(p) = bp_j^2(1-p_j)$ where $0<a<4$, $0<b<1$ for $j=1,2$

The learning rate of non-linear models are faster than those for linear models [Chandrasekaran & Shen 1968]. This scheme is ε-optimal in all stationary random environments. The A_{RP} learning algorithm (associative reward-penalty) is an associative reinforcement (competitive) learning systems which treats each neuron of a neural network as a stochastic automaton [Barto & Amandan 1985]. It is similar to the Robbin-Monroe gradient descent approximation procedure and is a generalisation of the perceptron algorithm and the Widrow-Hoff rule (adaline). It reduces to the non-associative L_{RP} algorithm under the appropriate conditions. A_{RP} provides pattern classification capabilities similar to the Boltzmann machine. A decision rule minimises the probability of mis-classification in terms of Bayesian a posteriori probabilities. Reinforcement employs gradient descent to alter the probability distribution of its responses:

$$p_i(t) = p_i(t-1) + \gamma(t)\left(\frac{\partial\phi_i}{\partial p_i}(1-\Omega_i) - \frac{\partial\psi_i}{p_i}u_i\right)/\,p_i$$

where:

ϕ_i=change in action probability for reward
ψ_i=change in action probability for penalty

Such an algorithm allows the learning automata to cope with dynamic and changing environments with robustness as it can discriminate between environmental states to maintain separate action probability vectors for each environmental state. The A_{RP} algorithm is ε-optimal. Combinations of linear and non-linear terms can improve the convergence of learning. While reinforcement schemes update the action probabilities, Bayesian techniques may be used to simultaneously estimate penalty probabilities and give fast convergence to optimal actions. Learning automata may be used for neural network synthesis, a discrete neural model being represented by:

W = <N,M>

where:

N = <S,A,B,I,O,T,C,P_1,P_2>
M = connectivity matrix among neurons in set N={(i,j,k)}
i,j = neuron address
k = synaptic link label
S = set of all neural states
A = set of output symbols
B = set of all input symbols
I = set of inputs
O = set of outputs
T = state transition function: $S\times I\rightarrow S;P_1$.
C = action function: $S\times I\rightarrow O;P_2$.
P_1 and P_2 = set of probabilities associated with each state transition (action)

These learning automata can function in stationary random environments only. A non-stationary environment will appear inconsistent because the automaton receiving the same input at two different times will generate the same actions which may generate different world states. This can happen if (a) the agent input-world state function is not one-to-one; (b) the agent has faulty sensors or effectors; (c) the world has a probabilistic transition function [Kaelbling 1991]. The only way to overcome this is to make the algorithms efficient in time such that the automaton's time constant is shorter than that of its environment by using simplifying assumptions about the environment to maintain environmental regularity. This usually involves using multiple autopmata with one automaton operating in its own stationary environment and switching between a number of stationary environments. This represents close to optimal behaviour. Hierarchically structured automata avoid the problem of excessive memory requirements for large numbers of possible responses. By dividing up the parameter space, faster convergence can be achieved by limiting the number of actions for each automaton. The action of each automaton at a certain level triggers automata at lower levels. The first level of the hierarchy consists of a single automaton with N internal states. The second level consists of N automata of N actions, and so on to the k^{th} level, totalling N^{k-1} automata. Action u_{ij} activates automaton A_{ij} at the next level, etc.

3.10.2. Neural Networks

Non-associative learning is associated with primitive response variation to single event stimuli such as habituation, dishabituation and sensitisation. Associative learning is based on causal inference in deriving relationships between two or more phenomena – the basis of concept formation. There are two types of associative learning – Pavlovian (classical) conditioning and Thorndikian (operant) conditioning. Reinforcement learning is based on punishment/reward signals to learn complex behaviours by favouring actions which generate the maximum numerical reward value produced the environment. Neural networks lie at the core of brain function in providing control of body movements through learning. The brain is essentially a sensor-effector control system. Most data concerning human brain function comes from cerebral thromobis (stroke) or other brain damage aphasias [Geschwind 1979]. The primary purpose of the brain is to analyse sensory data and control movement, i.e. an input-output processing device. The human brain comprises an interconnected system of 10^{11} neurons of white and grey matter which act as complex analogue processor units. Sensory neurons transmit information from the periphery to the central nervous system (CNS) while motor neurons transmit information from the central nervous system (CNS) to the effectors. There are two types of neuron – interneurons and output neurons. Interneurons process the information in the central nervous system between the sensory and motor systems of the body. The nervous system is divided into the central and peripheral systems. A typical motor neuron synapses with a very large number of axons from ~10^3-10^5 interneurons typically. A somatic motor fibre terminates on skeletal muscle while autonomic motor fibres terminate on cardiac muscle, smooth muscle and glands. The somatic part of the CNS is concerned with control of the voluntary muscles while the autonomic part is concerned with the control of the viscera, vessels and internal organs.

ANN's (artificial neural networks) are artificial analogues of neural architecture of the human brain. The first neural networks were the McCullogh & Pitts (1943) which modelled networks of neurons connected by simple weightings to emulate synapses. Each neuron was a AND, OR and NOT gate to perform logical computations on binary inputs which were thresholded to give all-or-nothing action potentials. Hebb (1949) the introduced the Hebb learning rule to provide causal representation of perceptual signals as patterns of neural activity by modulating synaptic weights The more often two neurons are associated with each other's activity, the greater the synaptic connection strength between them. Indeed, Hebbian learning has a neurobiological basis in long-term potentiation. Churchland & Churchland (1990) proposed that the parallel distributed architecture of the neural network is the appropriate cognitive architecture as the critical requirement for intelligence. Such do not manipulate symbols according to specified rules but may learn this cognitive skill as one of a range of capabilities. By virtue of its parallel and distributed operation, higher level emergent properties arise from the neural network's complexity. In this way, neural networks are not equivalent to universal Turing machines (though they may include the universal Turing machine as a subset of its operation). Connectionism is an holistic and nonsymbolic approach in that information representation is mapped directly into its architecture. Unlike the symbolic paradigm, the connectionist approach has an epistemological basis being founded on bottom-up processing through self-organisation. The chief difficulty with ANN's is that the effect of scaling to large sizes is to increase learning times exponentially. However this may be overcome by partitioning networks into modular hierarchical systems. Each module may comprise of a large number of neurons. And the various modules operate independently and in different ways. Parallel processing allows a large number of activated modules to influence each other in parallel. Hence modularity is just as important in connectionism as it is in symbolic models to overcome the complexity problem.

There are ~10^{11} neurons in the brain with very high fan-in and fan-out of ~10^3 synaptic connections per neuron [Hubel 1979; Nauta & Freitag 1979; Stevens 1985]. Such a degree of connectivity is much smaller than that required for a crossbar switching network which would otherwise make the fan-in and fan-out impossibly large at ~10^{22}. The brain comprises a hierarchically organised system of self-organising neuronal network modules [Szentagothai 1978]. Evidence for hierarchical organisation comes from the increasing receptive field along the brain's sensory processing path which increases from 1^o of arc at area 17 to 100^o of arc at area 19. The brain's primary function is that of a control system to produce highly complex behaviour. Neural networks are physical devices that embody transformation functions for sensorimotor integration of behaviour. It accomplishes this by virtue of its structure of highly interconnected assembly of relatively simple processing units (neurons) in a highly redundant fashion. A biological neuron is around 10^6 times slower than a silicon switch. The biological neuron performs non-linear analogue integrator computations rather than digital recursive algorithms while neural transmission is digital through pulse modulation over the axons at ~300 pulses/second. This is some 10^6 times slower than electronic signals. All electrical impulses are identical so that information is encoded in the spatio-temporal relations between neurons. The brain derives directional mappings, which define the relationship between a set of causes and a set of effects. Similarly, the artificial neural network transforms sets of input signals into sets of output signals. Neural networks comprise of a system of large numbers n of interconnected neurons each of which is a non-linear processing unit of limited capabilities. The synaptic links between neurons are unidirectional and each has an associated connection weight w_{ij} forming a non-linear network of intercommunicating elements. The values of the connection weights are obtained by learning from a set of training patterns. Neurons may be modelled as distributed RC networks with the cell cytoplasm providing a series resistance and the lipid cell membrane providing the shunt capacitance. An RC circuit sums the inputs according to the weighted synapses. Voltage and time dependent intrinsic currents model other non-linear and noisy effects:

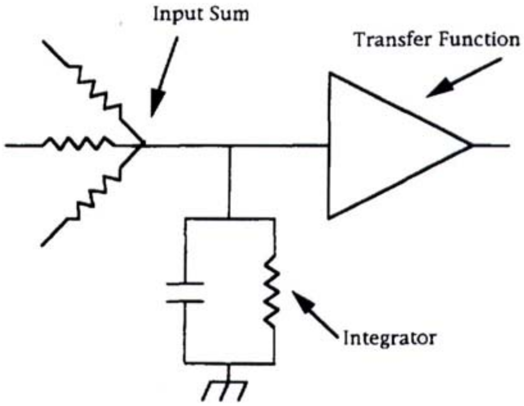

Figure 109 – Voltage and Time Dependent Intrinsic Currents Model of a Hopfield Neuron (from Gallagher et al 1996)

However, this is probably an oversimplification as neurons have complex dynamics as multi-input devices with many non-linearities from in-out cross-couplings [Lewis 1983]. The extended leaky integrate-and-fire neuron is a more sophisticated neuronal model (Figure 110)

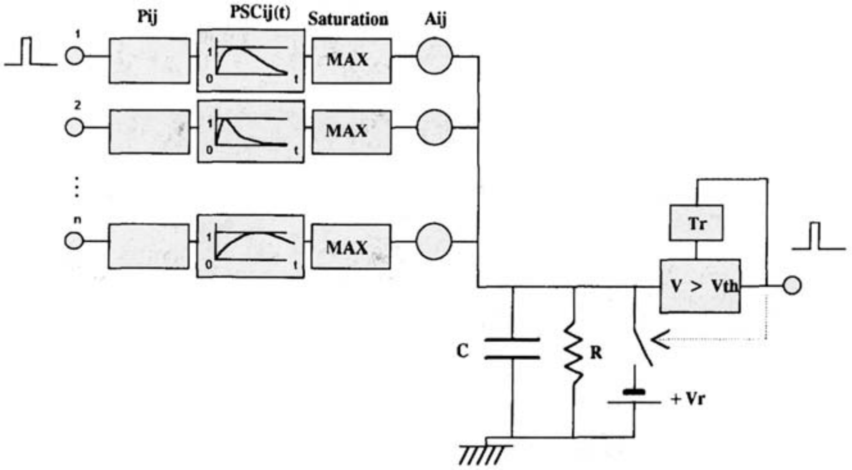

Figure 110 – Extended leaky integrate-and-fire neuron (Bugmann 1997)

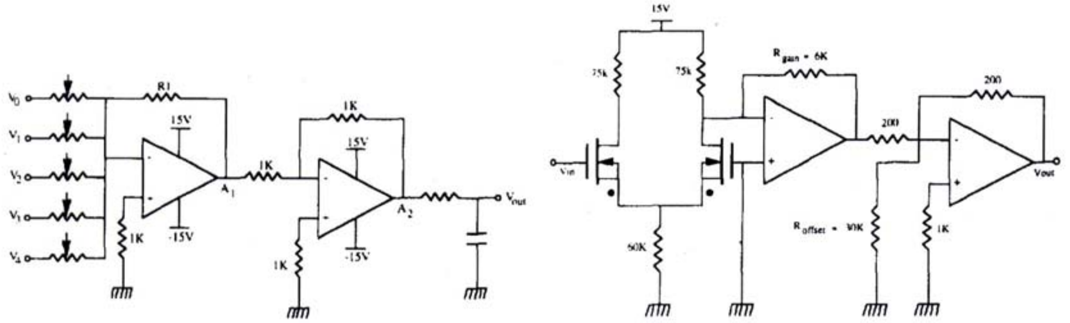

Figure 111 – a) Adder/integrator stage of each neuron includes an op-amp adder with inverter and passive RC network to integrate weighted sum; b) neuron map stage comprising matched n-mosfet pair and two amps to fine tune sigmoid function (Gallagher et al 1996)

The human brain is primarily a control system which produces goal-seeking biologically advantageous behaviour to find food, avoid danger, compete for territory for survival and attract sexual partners and care for offspring in the pursuit of gene propagation. Neural networks have been used for robotic control which are modelled on biologically plausible neurophysiological models, eg. emulation of the cerebellum and motor cortex. Neural networks of the brain are the means by which animals exert motor effects on the world – in the case of a human being, the most sophisticated and important of limbs are the arms. The simplest type of neural network control is the monosynaptic reflex such as the stretch reflex. It comprises of a sensory neuron connected directly to an effector neuron. The inside and outside of the neuronal cell has different concentrations of charged ions which creates a potential difference across the plasma membrane [Catherall 1984]. Neural transmembrane resistivity is typically 5 kΩcm filled with cytoplasm of resistivity of 50 Ωcm – the nerve fibre thus acts as a leaky conductor. The cytoplasm contains 400 nmole/kg water of K ions and 50 nmole/kg water of Na ions established by an Na pump which generates a resting potential of –70 mV (cf. normal cytoplasm concentration is 440 nmole/kg water of Na ions and 20 nmole/kg water of K ions). An electric signal causes membrane pores to open to external Na ions generating a depolarising action potential along the nerve fibre. The activity potential of a neuron acts as a crude memory and an inhibitory connection back to itself provides negative feedback to effect an exponential decay in activity potential $p = p_0 e^{-kt}$ where k=decay constant which determines the recovery period [Feldman & Ballard 1982]. The Na pump restores the resting potential within a finite refractory period of >3 ms – this limits mammalian biological reaction speeds to ~0.1s, i.e. a maximum information transfer rate of 30 bps. As signal voltages drop exponentially with distance, nerve fibres with lengths >1μ are encased in myelin sheath segments separated by nodes of Ranvier every 1 mm. The Hodgkin-Huxley model of nerve conduction suggests that the propagation velocity of an action potential is proportional to the square root of axonal diameter for unmyelinated nerves and proportional to axonal diameter in myelinated fibres. Output cells have graded analogue input with a digital all-or-nothing output. Interneurons however, have both graded analogue inputs and outputs – the graded potential output spreads out a short distance attenuating as it does so. All neural impulses have the same amplitude – neural information is frequency coded.

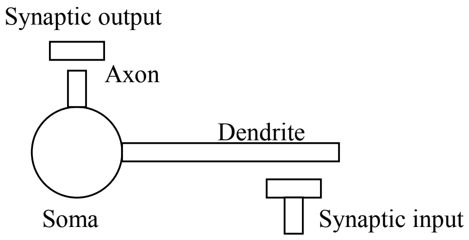

Figure 112 – Synaptic Schematic

The insertion of mediating interneurons between effector and sensory neurons allows the implementation of associative learning and more complex behaviours. Even simple animals possess small neural circuits for habituation, eg. the leech possesses only 40 neurons – four neurons are sensory neurons which interconnect to 30 interneurons which in turn output to a 10-neuron motor layer. Learning and generalisation in neural networks takes place by neuroplasticity of the connection weights (synaptic weights) which may be described by learning rules such as Hebb's rule (which may operate in the basal ganglia [Jabri et al. 2000]. Learning alters the quantity of chemical transmitter release [Kandel & Schwartz 1982]. This is determined by the concentration of free Ca++ in the presynaptic terminals – Ca++ enables the synaptic vessicles to bind and release transmitters in the presynaptic terminal. Environmental stimuli cause long-term neural changes and alters the levels of RNA transcripts encoding neurotransmitters which in turn alters the synaptic function. Impulses actively alter gene expression and increase mRNA synthesis and neurotransmitter protein synthesis over the long term. Hence, neural plasticity is provided by altering transmitter phenotypic expression in the nerve cell soma.

Neural networks are based on learning and the storage of that learned information in groups of simple computational units connected topologically. Beer et al (1990) used a neural model for simulated cockroach. Artificial neural networks (ANN) are based on electronic neurons arranged into regular layers connected by weighted links. Individual neurons react specifically to a given set of stimuli – this set represents the receptive field of the neuron. The information is effectively stored by the specificity of the neuron. The receptive field acts as a matched filter for those stimuli. A neural network thus learns to model input-output patterns and essentially produces an output:

$$y(t) = f\left(\sum_{j=1}^{n} w_j(t)x_j(t)\right)$$

where:

w_j= *association connection strength*
x_j=*input to neuron*

The function f is related to response mapping rules such as the firing frequency of the neuron to its membrane potential. If the weighted sum of those inputs to each neuron is above a given threshold level of excitation, the neuron fires and outputs a signal to the next neurons linked to it. ANN's process patterns of activation rather than symbols. Each input is propagated to the output layer through the network.

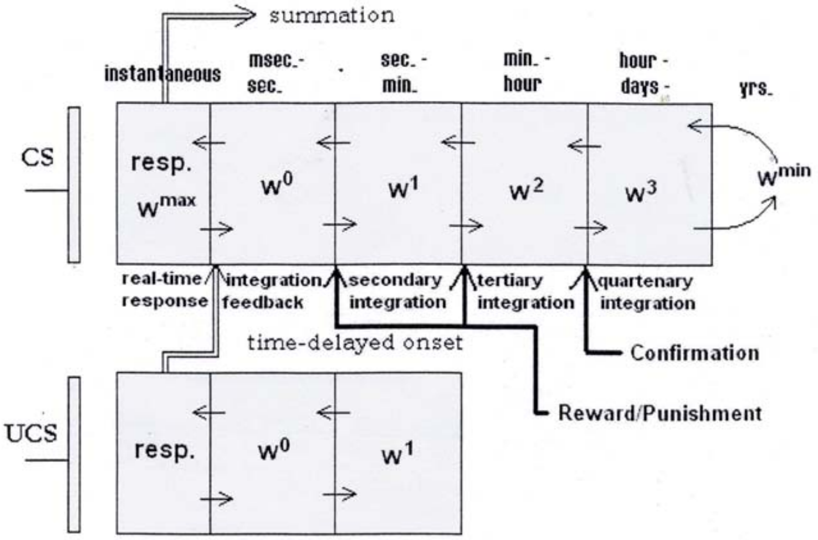

Figure 113 – Adaptrode model of adaptive biological synapse (Mobus 2000)

The commonest ANN architecture is the multilayer perceptron comprising of an input layer, a small number of hidden layers (usually one or two), and an output layer. Such ANN's can model continuous non-linear functions. Training the network involves adjusting the connection weights to minimise an error measure (usually sum of squared errors). The usual learning algorithm adopted is the backpropagation algorithm which is based on an error function of weights and thresholds. The application of neural networks in combination with fuzzy logic offers the possibility of coping with uncertainties by learning. The artificial neural network performs a weighted summation of its inputs and its output depends on that summation exceeding a threshold – it may be regarded as a form of Kalman filter [Ruck et al 1992]. Multitudes of these "neurons" connected together form the network but are commonly configured into layers. The units represent highly simplified representations of neurons in the brain, and their connectivity represents highly simplified representations of the topological connectivity of neurons in the brain. The most important simplifications are that the computational neuron is represented as a thresholded summer circuit, and the computational connectivity is generally limited to layered topologies only. There are many unsolved problems in neural network design – the choice of learning rules, the values of learning parameters (though this may be alleviated somewhat by representing the neural net as a Kalman filter), the separate phasing of learning and operation, the issue of over-learning, the choice of node connectivity, the choice of number of layers (though three layers can theoretically approximate to any non-linear mapping function), and the choice of number of nodes for the hidden layer(s). To illustrate the general form of the neural net, we present it as a Kalman filter. The multilayer perceptron (MLP) adopts a fully interconnected node structure between layers. Ruck et al (1992) showed that the backpropagation algorithm for assigning weights in the multilayer perceptron in batch mode is a degenerate form of the extended Kalman filter for training the weights in instantaneous mode. The Kalman filter attempts to estimate the state of the system, i.e. the MLP weights $w(t_i)$ which the output of the network represents as the measurement or observations $z(t_i)$ of the filter.

$$\hat{w}(t_i + 1) = \hat{w}(t_i) + K(t_i)[y^d(t_i) - y(t_i)]$$

where:

$$K(t_i)=P(t_i)H^T(t_i)[H(t_i)P(t_i)H^T(t_i)+R(t_i)]^{-1}$$
$$P(t_{i+1})=P(t_i)-K(t_i)H(t_i)P(t_i)$$
$$H_{ij}={}^{\partial y_i}/_{\partial w_j}$$
$$y(t_i)=h(w(t_i))+v(t_i)$$

Observations are modelled as non-linear function of states corrupted by zero mean white noise where

$$< v(t_i)v^T(t_j) >= R(t_i) \qquad \text{for } t_i = t_j.$$

$$< v(t_i)v^T(t_j) >= 0 \qquad \text{for } t_i \neq t_j$$

Backpropagation may be represented in the form: $\delta w_{ij} = \eta \sum_{k=1}^{n} (y_k^d - y_k) \frac{\partial y_k}{\partial w_{ij}}$

Extended Kalman filter reduces to:

$$\delta w_i = ap \sum_{j=1}^{n} (y_j^d - y_j) \frac{\partial y_j}{\partial w_i}$$

where:

$$a = [HPH^T + R]^{-1}$$
$$p = P(t_i)$$

Filtered output is a weighted combination of past input signal samples. It is possible to estimate future values of time-correlated signals from present and past input samples without having the autocorrelation function. Filter weights converge to best least squares estimate similar to Kalman filter. Hence the backpropagation algorithm is a degenerate form of the extended Kalman filter with no ability to modify the training rate as $\eta = ap$ which may be useful in updating weights. However, the Kalman filter requires three orders of magnitude more floating point operations for a given accuracy. The backpropagation algorithm requires an external supervisor to specify the desired output states. However, the input itself can perform supervision using a multilayer encoder network where the desired output vector is precisely the input vector. A variant of the backpropagation algorithm is the A_{R-P} procedure which although slower does not require a separate pass to backpropagate the error information. Backpropagation is limited as a learning time scales poorly with $O(n^3)$ on a serial machine (where n=number of weights). A parallel hardware network with one processor per connection could reduce this time to $O(n^2)$. There is no evidence that biological synapses in the brain can be operated in reverse to implement backpropagation of derivatives but Stevens (1989) notes that reciprocated dendrite-dendrite synaptic contacts occur which may be used for feedback.

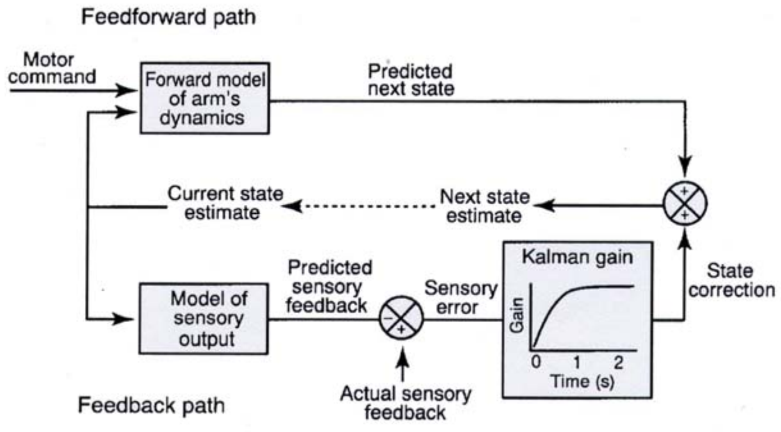

Figure 114 – Kalman filter model of sensorimotor integration (Wolpert 1997)

3.10.3. Chaotic Neural Networks

It has long been known that conventional neural network models can exhibit chaotic behaviour by virtue of their non-linearity. The neural activity trajectories for functional modules (nuclei) are characterised by chaotic rather than point or periodic attractors on a Poincoire map [Harth 1983]. Chaotic attractors in neural dynamics can store dynamic behaviours with several initial conditions allowing the system to randomly between solutions. Macroscopic neural activity exhibits irregular fluctuations without periodicity due to dissipative non-linear adaptive activity. Hence, the macroscopic neural behaviour is unpredictable. Chaotic dynamics arises a result of non-linear feedback with time delays with neurons possessing more than one firing level to produce a range of possible behaviours, i.e. higher level cognitive processes such a generalisation and abstraction are not predictable from microscopic neural dynamics. This differs from the view of cognition as derivable from sets of finite state machines. For a set of neurons, there is an emergence from chaos of a society of identical neurons with different dynamics and patterns of behaviour by virtue of the random nature of noise [Zak 1991]. The upshot is that macroscopic features evidenced by EEG traces cannot be predicted by measurement as the behaviour is determined by microscopic randomness of neural mechanisms. Wang (1991) constructed a simple 2-2 neural network with weights (**a**,**ka**,**b**,**kb**) and sigmoidal function $\sigma(x) = \frac{1}{1+e^{-x/T}}$ in which the outputs were fed back to the inputs. Wang found that there exist period doublings to the chaotic regime and the existence of strange attractors by using a homeomorphism from the network to a known dynamical system having these properties. This formally established that neural networks can exhibit chaotic behaviour. Further examples were given subsequently by Dracopoulos et al (1993). One example mentioned earlier was the Ikeda map defined by:

$$g(z) = \gamma + Rz.\exp(i(k - \frac{\alpha}{1+|z|^2}))$$

where:

z is a complex variable of the form x+iy, $I^2 = -1$
R, k and α *are suitable constants*

The dynamical system is then defined by $z_{n+1}=g(z_n)$. For a given set of parameter values, this mapping possesses a strange attractor - the Ikeda map. However, biological neural networks act as chaotic associative memories which provide robustness to noise. Neural chaos is an intrinsic property of normal brain function. Many linear systems exhibit a fixed point attractor point to which the dynamic trajectory converges over time. In contrast, non-linear systems have no such fixed points but exhibit trajectories that never repeat. They usually comprise an infinite set of unstable periodic orbits which define the strange attractor. Such chaotic dynamics can occur even in very simple systems. They are characteristic of non-linear systems, whereby a system's behaviour is not proportional to the variables that describe the system to the extent that the future dynamics of the system are not predictable in the long term. Chaos is characterised by great sensitivity to initial conditions such that small perturbations will generate large differences in subsequent behaviour. The basic tools for analysing such systems are diverse, for example *Lyapunov exponents* (which characterize the local rate of trajectory divergence) and *Poincarè map* are two such tools. The Lyapunov exponents provide a coordinate-independent measure of the asymptotic local stability of properties of a trajectory. The Poincarè map is a technique which models a continuous time system as a discrete system. A specified sub-manifold of the state space, called the Poincarè section, is selected. As the system evolves through time the trajectory repeatedly intersects the Poincarè section. Plotting the points of intersection creates a lower dimensional portrait of the system behaviour. ECG waveforms of normal human hearts show evidence of chaotic cardiac cycles in that cardiac activity is not perfectly periodic but exhibits the characteristics of chaotic dynamics [Babloyantz & Destexhe 1988; Garfinkel et al 1992]. Perturbations can be induced near the desired fixed point attractor to force the system onto stable trajectories.

Artificial neural networks are artificial analogues of neural architecture of the animal brain. They embody transformation functions for sensorimotor integration of behaviour. They comprise of highly interconnected assemblies of relatively simple processing units (neurons) in a non-linear network of intercommunicating elements. Learning procedures may be utilised in neural networks to extract internal representations by modifying the synaptic weights between neurons through gradient descent methods by computing the error derivatives with respect to the synaptic weights. They function as associative memories in storing associations between sensory and response patterns. Such representations of stimulus-response associations are content-addressable such that when a trained network is activated by an input pattern that is similar to the stored pattern it will respond similarly. Neural networks may be of the feed-forward type with unidirectional information flow (eg. the perceptron) or of the feedback type with bilateral connections (eg. the Hopfield net). The dynamics of the conventional (symmetric) Hopfield neural network may be characterised as using an energy-like Lyapunov cost function:

$$E = -\tfrac{1}{2}\sum_{i}^{n}\sum_{j}^{n} w_{ij}x_i x_j + \sum_{i}^{n}\theta_i x_i$$

where:

w_{ij} *is the weight from unit j to unit I*
θ_i *is the threshold of unit I*

For a Hopfield network under asynchronous updating it is easy to see that E monotonically decreases to a local energy minimum. This shows that the dynamics is particularly simple and that the Hopfield network will always stabilize to a fixed point. (The case for synchronous updating can be similarly treated with a different Lyapunov function and behaves in much the same way except that period 2 orbits are possible.) Memories are then associated with fixed points by some weight assignment algorithm, e.g. the outer-product rule. It is well known that the memory capacity of a Hopfield network for uncorrelated patterns is around 0.15n, where n is the number of neurons.

Deterministic feed-forward networks may not be an accurate analogy to biological computation regardless of the training algorithm employed. For the chaotic neural models it is often the case that no such simple Lyapunov function exists. Moreover memories are no longer associated with fixed points but rather with unstable periodic behaviours, in principle there are infinitely many potential memories. For these more complex systems is not clear how learning should progress. While different stimuli can produce different stabilized periodic behaviours the study of how progressive strengthening of these associations via some form of Hebbian learning can be plausibly implemented is one of the aims of this project. A variety of learning algorithms for more conventional models exist. *Hebbian learning* seeks to modify connection weights so as to correlate firing rates of connected neurons, often forming an association between the input and output neural activations: $\delta w_{ij} = u_i u_j$. For purely feed-forward networks, gradient descent of some error function can be employed via *error backpropagation*. The error backpropagation algorithm is a generalisation of the least squares error procedure for single layer networks. Although there is no evidence that biological synapses in the brain use backpropagation of derivatives, reciprocated dendrite-dendrite synaptic connections do occur possibly for feedback. Hence, traditional neural learning approaches are not appropriate. This includes learning algorithms such as the meta-backpropagation algorithm or even genetic algorithms. Furthermore, it is possible that memory storage in unstable periodic attractors may offer considerable enhancement of memory capacity over memory storage in point attractors (as in the Hopfield network model).As a result of his studies of neural activity in the olfactory bulb of the rabbit [Freeman 1991] hypothesised that chaos may play a crucial role in perception and recognition. EEG's on humans have shown evidence of chaotic behaviour. EEG's show patterns of brain activity that depend on functional state. Although EEG traces do not directly represent the underlying neural dynamics that are of interest here, they are instructive. Generally EEG signals are highly complex. Spikes of activity appear before an action is made indicating mental events. Phase portraits made from EEG's generated by computer models reflected the overall activity of the olfactory system of the rabbit at rest (control) and in response to a familiar scent such as banana. The resemblance of these portraits to irregularly shaped yet structured dynamics exhibited on a Poincare map revealed that brain activity under both conditions is chaotic but that the response to the known stimulus is more ordered and more nearly periodic during perception than at rest. Such changes of neural activity occur very rapidly - large assemblies of neurons switch to different stable dynamical behaviour in response to small changes of sensory input. The infinite number of periodic orbits in the strange attractor represent a vast range of possible behaviours. Chaotic attractors in neural dynamics enable the storage of dynamic behaviour with multiple input conditions allowing the system to randomly converge to different solutions. There is then a tendency of neural networks of neurons to exhibit dynamics that shift from complex chaotic activity to more stable periodic activity in response to low levels of input. The chaotic neural activity exhibits irregular fluctuations without periodicity due to non-linear dissipation of the dynamics. The chaotic motion makes possible an infinite number of unstable periodic trajectories with unstable trajectories switching from one periodic motion to another giving the appearance of randomness. Hence the neural behaviour is unpredictable. Chaotic dynamics arises as a result of non-linear feedback from re-entrant pathways which introduces time delays. Neurons possessing more than one firing level produce a range of possible behaviours. The probabilistic characteristic of chaotic oscillations allow the storage of temporal behaviour in a compact form. The upshot is that macroscopic features evidenced by EEG traces cannot be predicted by measurement because neural behaviour is determined by randomness in neural mechanisms. From an evolutionary point of view this phenomenon is significant. The extreme sensitivity to initial neural input conditions displayed by chaotic neural networks makes them unstable and unpredictable but it also makes them susceptible to control by very small inputs. This enables rapid switching among many different behaviours by small perturbations. A large number of different structures are possible as a consequence of small random inputs to which neurons are sensitive, possibly even without input from the outside environment. This implies that very low energy inputs are required to alter neural behaviour, much less than that required to alter the dynamics of deterministic neural networks. The chaotic network offers a more energy-efficient mode of behaviour generation.

Control may involve the elimination of multiple basins of attraction, stabilisation of fixed points or the stabilisation of unstable periodic orbits, and it is this last that is of interest. Generally the OGY method can be applied to the control of conventional feed-forward networks whose behaviour under iterated feedback has been trained to be chaotic. The strange attractor comprises an infinite collection of unstable periodic orbits which never repeat. The OGY (Ott-Gregobi-Yorke) control method forces the chaotic system to stay near selected unstable periodic orbits. First a Poincare section through the state space trajectory is constructed. The location of the unstable fixed points must be determined followed by the directions along which the trajectory converges (stable eigenvectors) and those along which the trajectory diverges (unstable eigenvectors) with respect to the unstable fixed point. The rates of divergence and convergence are obtained and finally the control parameter change required to effect the motion onto the stable direction towards the fixed point. When the system comes near the desired unstable fixed point in a periodic orbit in the section, a small perturbation is applied to the system control parameter under investigation to keep the motion near that fixed point. If u is a scalar control parameter over the time interval t_i, the aim is to vary u around the nominal value u_0 by a small amount δu to stabilise the output y(t), i.e.

$$\delta u = u - u_0$$

$$\delta y_{i+1}(t_i) = y_{i+1}(t_i) - y(t_0)$$

The output is related to the control input by a non-linear function F: $y_{i+1}=F(y_i,u)$.

The behaviour of F near the control point u is described by the Jacobian matrix: $J = \nabla F$. Hence, $\delta y_{i+1}(t_i) = J\delta y_i(t_{i-1}) + \delta u_i$ The idea of the OGY method is to choose control inputs δu_i to eliminate components of output δy_{i+1} that are further away from the control point region $y(t_0)$, i.e. unstable trajectories that diverge from the fixed point. If the Jacobian has linearly independent eigenvectors with real eigenvalues then it is possible to define a function f such that: $fJx = \lambda fx$ where λ_i=real eigenvalues of J. A linear approximation of the dynamics in the neighbourhood of the desired fixed point is valid allowing the use of local parameters rather than global parameters. The OGY theorem states that the constraint $f \,.\, \delta y_{i+1} = 0$ generates the first order control law:

$$\delta u_i = -\lambda f \, \frac{\delta y_i(t_{i-1})}{fu}$$

where:

$$f_j = 1{,}0 \text{ if } i = j, i \neq j \text{ respectively.}$$

This control law ensures the stabilisation of the periodic dynamics around the attractor adjacent to the corresponding unstable point. The parameter u is adjusted externally to modify the system dynamics from the nominal u_0. The position of the periodic orbit is a function of u. The point is that only very small control signals are required by virtue of the extreme sensitivity of the chaotic system to control inputs. The OGY method forces the chaotic motion to stay near selected unstable periodic orbits so stabilising one of these periodic motions by virtue of the chaotic system's extreme sensitivity to perturbations in initial conditions. Control by variation of a global slope parameter is difficult to achieve but it becomes easier when control variations are applied locally to a single layer rather than the whole network. Generally, it is notable that control becomes very much easier when the controlling parameter is a small signal applied to one of the inputs. This may be closer to being a biological analogy than the control of behaviour through global control. Other approaches are possible using different neural network models and different control techniques. Babloyantz's group [Sepulchre 1993, Lourenco 1994] for instance have controlled a network of oscillators coupled to their four nearest neighbours using both the OGY method and a delayed feedback control technique first suggested by Pyragas [Pyragas 1992]. Sole & de la Pride (1995) uses the GM technique [Guemez & Matias 1993, Matias & Guemez 1994] to control small fully connected networks of three or four neurons whose attractors are similar to a chaotic network first discussed by Wang (1991). These control mechanisms of relaxation onto a fixed point are external to the neural network. This is in contrast to the hypothetical biological process in which some form of Hebbian learning is used to cause frequently encountered sensory input to be associated with a particular (unstable) periodic dynamical regime. Consequently when the sensory input is re-encountered, the neural system relaxes naturally onto a particular (unstable) periodic behaviour which characterises the input. This switching of the dynamical behaviour is implicit in the neural structure rather than being externally imposed. Time delayed feedback can be used to stabilize high dimensional chaotic systems, in particular neural systems. Since time delayed feedback is easily implemented in neural circuitry this method becomes a serious contender for a biologically plausible implementation of an associative memory based on neural chaos.

3.10.4. Cerebellar Models

Neural modelling has been particularly applied to motor and perceptual tasks rather than logical reasoning and knowledge organisation as they provide a better description of low level processing such as pattern recognition and of motor skills which are all highly parallel processes. Hence, they have become dominant in machine perception and motor control. This is particularly true of biological motor control which has particular application in robotics. The control of muscular movement is one of the most important functions of the brain, and large sections of the mammalian brain is devoted to this important function. The primary motor cortex (area 4) stretches across the roof of the brain in front of the central sulcus of the cortex and initiates voluntary muscle movements. The motor cortex is responsible for voluntary motor control over the muscles of opposite sides of the body. It comprises large pyramidal cells with very long axons projecting towards the spinal cord. It receives inputs from other parts of the cortex, the cerebellum and the basal ganglia. The premotor cortex (area 6) lies in front of the motor cortex and coordinates and sequentially links units into complex motor activity. It has subcortical connections with the coordination nuclei of the thalamus through which it links to the cerebellum. Motor neurons lie within the spinal cord, hindbrain and midbrain (none exist in the forebrain). These comprise the lower motor system for local motor control of different parts of the body. Within the hindbrain, motor neuron projections originate in the inner two-thirds of the hindbrain reticular formation. Within the midbrain, motor projections originate in the superior collicus and red nucleus – the superior collicus receives input from the optic nerve and cerebral cortex while the red nucleus receives input from the cerebellum and the motor cortex. All motor fibres descend into the spinal cord conveying commands originating from a wide region of the brain, especially the cerebral cortex. The superior collicus obtains inputs from the optic nerve and large regions of the cerebral cortex. The red nucleus obtains inputs from the cerebellum and the motor cortex.

There are two main non-pyramidal systems controlling movement: the cerebellum and the extrapyramidal system. Pre-programmed automatic behaviour is characterised by information flow from the cerebellum to the thalamus on to the motor cortex. The thalamus relays to the motor cortex signals from the basal ganglia which initiates movement. All voluntary movement begins in the motor cortex which comprises large pyramidal cells. Such cells have very long axons that connect to the root of the spinal cord resulting in muscle constriction. In addition, it receives feedback from the muscles and joints. Voluntary movement requires sensory feedback analysis from the post central sensory cortex. Spatial analysis for goal orientation is performed in the tertiary parts of the parieto-occipital regions of the cortex. The premotor cortex is responsible for sequential scheduling of the separate elements of motor activity. The prefrontal lobes provide the intentions and goals to guide motor behaviour.

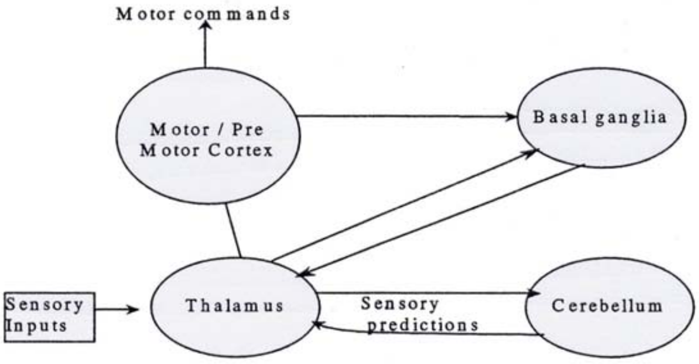

Figure 115 – Sensorimotor architecture where sensory inputs are relayed through thalamus to basal ganglia and cerebellum and premotor cortex, notor cortex and thalamus act as relays (Jabin et al 2000)

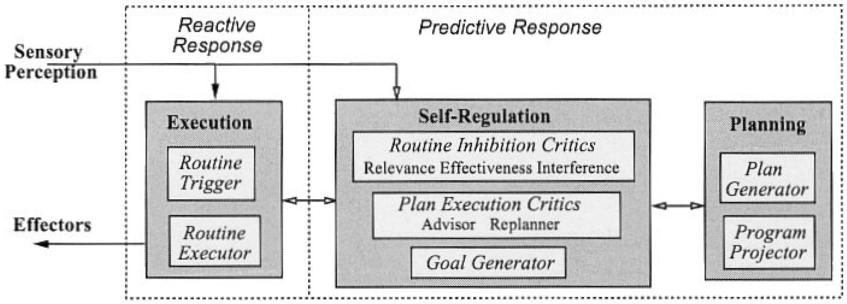

Figure 116 – Artificial frontal lobe (Levinson 1994)

The motor system comprises the upper motor pathway, lower motor neuron, myoneural junction and muscle. The upper motor neuron pathway begins in the precentral gyrus. The pyramidal tract descends from the precentral gyrus, some of the fibres passing into the midbrain, pons and upper medulla. In the lower medulla, the pyramidal tract passes from the brain stem to the spinal cord. The lower motor neuron is final pathway to muscle. The arrival of the motor neuron at the neuro-muscular junction causes the release of acetylcholine which crosses the nerve/muscle cleft and becomes attached to the acetylcholine receptor in the motor endplate. This causes depolarisation and muscular contraction. Cholinesterase destroys the acetlycholine and the muscle repolarises. Coordination deends on the smooth contraction of one group of muscles (agonist) and equally smooth relaxation of opposing muscles (antagonist). The pyramidal system comprises descending fibres of the motor cortex which cross to the opposite side and terminate in a distributed fashion to the voluntary nuclei of the cranial nerves and the spinal cord. The number of fibres of the corticospinal tract decrease down the spinal cord as they terminate. Peripheral nerves fibres thence are distributed from the spinal cord to the striated muscles, each innervating ~100-300 separate muscle fibres. Skeletal muscles are innervated by alpha motor neurons. Muscles have proprioceptive sensors (muscle spindles, Golgi tendon organs, joint receptors) for monitoring and controlling motor activity. Reflexes involve afferent fibres directly synapsing onto efferent fibres across the spinal cord while bypassing the brain. In the stretch reflex, muscle stretch causes the discharge of impulses in its neuro-muscular spindles. Afferent fibres pass to the spinal cord and synapse directly onto efferent fibres which conduct impulses from the spinal cord to the motor endplates of the muscle fibres which increase muscular tension. The spinal cord possesses afferent fibres for proprioceptive and tactile information to be passed from the muscles, tendons, ligaments and joints to the brain. They follow three divergent routes:

1. direct to the lower motor neurons for reflex arcing;
2. fibres to the spinal cord to relay to the cerebellar cortex (spinocerebellar tract);
3. fibres to the spinal cord to the medulla and thalamus and then to the somatosensory cortex.

The cerebellum is the major component of the rhombencephalon. It is attached to the pons and medulla by the cerebellar punducle. The cerebellum comprises ~10^{11} granule cells. The cerebellum is primarily responsible for muscular coordination and regulation of motor skills – indeed, unlike the case of the cerebral cortex, each cerebellar hemisphere controls its own side of the body. The cerebellar cortex is the corrugated outer layer of grey matter enclosing a core of white matter. The cerebellar cortex consists of neurons arranged in highly ordered repeated patterns. It comprises several types of neurons:

- long axoned input excitatory climbing and mossy fibres conduct impulses to inhibitory outputs;
- Purkinje cells of the cerebellar cortex with extensive dendritic inputs of 10^5 synapses forming a continuous parallel array layer with dendrites extending towards the cortex and axons extending towards the underlying granule cells;
- long axoned excitatory granule cells;
- short axoned inhibitory Golgi basket and stellate cells act as intermediate neurons between the input and output neurons.

The cerebellum predicts body movements and comprises 50% of the human brain's neuronal complement. The cerebellum essentially simulates muscular movement through feed-forward models of motor commands before feedback from the muscles is available. It provides the means for rapid movement through pure prediction – the biological analogue of the Smith Predictor. The internal world model emulates the effect of the command based on past experience. The model is updated when feedback becomes available. The forward predictor model simulates the effect of the commanded movement based on the body's current state. This model is supplemented by an inverse control model which represents the mapping of the motor commands to the muscular movements. There are a vast number of linked models stored in the cerebellum. Sensory feedback provides constant calibration of these models. The Purkinje cells prevent the transmission of impulses from the input fibres that would otherwise reach motor neurons and generate movement. The cerebellum can thus nullifies self-generated stimulation to differentiate from environmental signals. The cerebellum is also the organ in electric fish, echolocating bats and whales for processing signal echoes. Some secondary sensory neurons in the spinal cord and hindbrain channel directly into the cerebellum. Inputs to the cerebellum include the inferior olive of the brainstem which relays sensory information from the limbs. These inputs connect to the Purkinje cells of the cerebellum via the climbing fibres. During actual movement, sensory pathways via the climbing fibres are gated out. The timing of the gating detects the mismatch between predicted and actual movements. Gating cancels out sensory feedback from self-commanded movement. The cerebellum is divided into two cerebellar hemispheres fused at the mid-line. The central part of the cerebellum (vermis) is concerned with equilibrium and maintains balance when the centre of gravity of the body changes. The flocculonodular lobe comprises paired appendages (flocculi) at the posterior and represents the cerebellar portion of the vestibular system – it is evolutionarily the oldest part of the cerebellum. The anterior lobe receives most of the proprioceptive impulses from the spinocerebellar pathway. Both lobes are prominent in primitive vertebrates. The middle lobe is the main part of the cerebellum and receives input from the cerebral cortex – it is expanded as a neocerebellum in mammals. The fastigiobulbar tract travels from the cerebellum to the vestibular nuclei to close the vestibular circuit. Afferent fibres also come into the cerebellum from the vestibular nerve. The cortex of the cerebellum controls voluntary motor activity by commands issued from the motor cortex through the corticopontocerebellar pathway, and through feedback through the spinocerebellar tract from the muscles. It does not initiate motor actions but receives inputs from the motor cortex and proprioceptive inputs [Llinas 1979]. Each proprioceptive nerve in the skeletal muscles map onto climbing fibres organised into strips over the cerebellar cortex. Corrections to inaccuracies of muscular activity are dispatched as nerve impulses from the detate nucleus. Ataxias result from cerebellar damage which result in postural or gait disequilibrium, uncoordinated movements and overshoot movements.

Kung & Hwang (1989) considered that artificial neural networks were suitable for robotic path control as it represented a pattern recognition problem. The path control problem can be defined as the mapping between measured sensor signals into computed actuator commands. It may be implemented as a two stage process - one neural network produces the desired torque commands from the desired joint trajectory inputs, while the other neural network produces reference torque commands from the actual trajectory inputs. The first stage involves generalised learning. In the first stage, a neural net is trained to output the desired joint motor commands from the desired trajectory. These are used as inputs to drive the manipulator to produce the actual joint trajectory which allows fine-tuning of the system of the form p(τ) after the generalised learning phase. This second stage involves specialised learning in attempting to minimise the training set required for generalised learning. The actual trajectory provides the input to the second network. The second network derives the reference torque commands in attempting to minimise the error between the reference torques and the desired torques in a least squares sense. Nguyen & Widrow (1990) used a slightly different technique applicable to robotic manipulators. One multilayer artificial neural network is trained to identify the system dynamics (the emulator) while another multilayer artificial neural network learns to control the actual system dynamics (the controller), both using the backpropagation learning algorithm. Once the emulator has been trained to match the manipulator dynamics, it is used to train the controller network. The controller network learns to drive the emulator from an initial state to a desired state over a number of trials by minimising the error function. The overall objective was to train the controller network to produce the correct control signals to drive the manipulator to the desired trajectory given the current trajectory state. The emulator acts as a feed-forward component to the control system while the controller acts as an error feedback component effectively implementing an adaptive computed torque control method. This prompted the use of the two-layer perceptron neural network as equivalent to a linear control strategy (i.e. the PD feedback control law) and the three-layer perceptron neural network as a non-linear adaptive feed-forward strategy [Yabuta & Yamada 1990].

Of particular interest are artificial neural network models of the cerebellum which store pre-learned of associative motor activation patterns, eg. CMAC (cerebellar model articulation computer) [Albus 1975a,b; Kawato et al 1988, 2000; van der Smagt 2000]. Feed-forward systems predict future behaviour based on learned models of the current state and the behaviour of the system (emulation). Feedback provides the basis for determining actions required based on the difference between the current state and the goal state (control.) In the case of control of the human body, the feed-forward model emulates the behaviour of the musculo-skeletal system. The feed-forward model provides predicted feedback faster than real feedback from peripherals. Cerebellar neural network models provide the basis for emulation to overcome the limitations in feedback delay from proprioception. The emulator takes the current state of a system as its input together with the commands to be implemented and predictively outputs the resultant state of the system. This provides the basis for future feedback – the emulator "forward models" the system's behaviour to provide rapid error-correction (similar to model-referenced adaptive controllers or Kalman filter-based controllers). This also provides the basis for internal representations in mental imagery in higher level off-line planning without the need for continuous matching with the real world and its implementation within a dangerous world but which guide future goal-oriented behaviour [Clark & Grush 1997].

The cerebellum is attached to the midsection of the upper spinal cord under the visual cortex at the rear of the brain [Albus 1979a,b; Marr 1969; Llinas 1979]. It is the cerebellum that is involved with the control and learning of coordinated and skilled movement of limbs, hands and eyes and is comprised of 10^{11} granule cells. Indeed, it has been suggested tthat it plays a wider role in learning and memory [Thompson 1986]. It is highly convoluted and the cerebellar cortex comprises of neurons arranged in highly ordered repeated patterns. The main input to the cerebellar cortex are the mossy fibres which convey information from the cerebral cortex (sensory-motor information), proprioceptive sensors in the muscles, joints and skin, the reticular formation (alert information) and the vestibular-ocular system (balance information). Each proprioceptive neuron in the skeletal muscles map onto climbing fibres organised into strips over the convoluted cerebellar cortex. Proprioception refers to sensory inputs during muscular movement which provide negative feedback information for closed loop control. The gamma-motor neuron senses muscular elongation while alpha-motor neuron senses muscle tension. The association cortex sends desired joint movement commands to the motor cortex which produces the required motor torque. The actual joint movements are sent back from the proprioceptor sensors to the motor cortex for error feedback control. The cerebellum also receives information about the joint command torques and the results of actual movements from the proprioceptors and on the basis of this constructs an internal neural model of the neuroskeletal system to provide internal feedback loop (motor learning).

Excitatory mossy fibres input to inhibitory output Purkinje cells via Golgi cells, granule cells, basket and stellate cells acting as intermediate cells. Mossy fibres come in two types: those carrying commands from higher control levels, and those carrying feedback information. The mossy fibres make excitatory contact. Granule cells are the commonest cells of the human brain with $3x10^{10}$ such cells in the cerebellum alone (about the same number of cells in the human-specific cerebral cortex). Each granule cell is contacted by around 10 mossy fibres and each granule cell give off a single output axon which travels to the cerebellum surface at which point it splits into two parts in opposite directions along the convolution of the cerebellum. Hence, these axons form a parallel sheet of parallel fibres. Each granule cell receives input from several different fibres and no two granule cells receive input from the same combination of mossy fibres. Hence, a precise variable can be transmitted reliably over multiple imprecise channels. One of the cell types contacted by the parallel fibres are Golgi cells which form a dendritic tree which is excited by the parallel fibres. Each Golgi cell puts out an axon which branches extensively making inhibitory contact for negative feedback with 10^5 granule cells in its neighbourhood to pevent granule cell overload from mossy fibre inputs by raising their thresholds. Only a small ~1% of granule cells can fire at any one time as the Golgi cells suppress the outputs of all but a few maximally stimulated cells by inhibitory contact. Every input pattern is transformed by the granule layer into a smaller subset of active parallel fibres. These active parallel fibres make excitatory contact with Purkinke cells and basket and stellate cells. Each Purkinje cell has 10^5 synapses with parallel fibres. Each parallel mossy fibre synapses with 300 synaptic contacts with Purkinje cells and each Purkinje cell can learn 300 different mossy fibre input patterns. These cells have weighted synapses and each Purkinje cell performs a summation of its inputs and produces the output to the cerebellar cortex. Basket and stellate cells are invertors which synapse exclusively with Purkinje cells and give them their negative which. Granule cells generate parallel fibres which act through weighted connections on the Purkinje output cell varying its firing rate. Climbing fibres also enter the cerebellar cortex one for each Purkinje cell for adjusting the strength of the weighted connections through excitation thereby controlling the patterms between granule and Purkinje cells. Climbing fibres connect cells of the inferior olive to the Purkinje cells in a one-to-one relationship and each olivary cell responds to a cerebral instruction for an elemental motion. When olive cells fire they send an impulse to the corresponding Purkinje cells while the mossy fibre input provides contextual information and the axons of the Purkinje cells are the only output from the cerebellar cortex and they induce motor movement. Hence, every action is defined as a sequence of firing patterns in the olive. The Purkinje cells can learn all the situations in which the olive cell movement is required. In this way, the cerebellum can learn the context in which the action is appropriate and involves it automatically. Collectively, the Purkinje cells can store multiple contexts in a distributed fashion. Granule cells and Golgi cells together act a pattern discriminators. The command circuit is generally of the form receptor/cortex-olive-Purkinje cell-effector in a negatoive feedback loop. This frees the cortex from routine functions. Muscle tissue contains special spindle cells that act as receptors which detect contraction or stretch in a muscle. The muscle itself is triggered by efferent nerve fibres from the spinal cord. The spinal cord carries efferent-motor impulses down to the muscles while sensory impulses are carried up the cord. Efferent fibres to the muscles are of small ~5μ diameter variety in the ratio of 1:3. The larger fibres are fast conductors of muscle impulses while the small slow-conducting fibres connect to spindle fibres and modulate their receptivity. Hence, small fibres alter the amount of kinaesthetic feedback to the central nervous system. Purkinje cells prevent impulses from input fibres that would otherwise reach motor nerouns and generate behaviour. The red nucleus of the cerebellum which is particularly well developed in primates also appears to exhibit plasticity of synapses for motor learning. This is the site of the construction of the inverse dynamics model of the musculo-skeletal system. Once learned, this model can compute motor commands directly from the desired joint trajectory. This provides the means for learning through error correction.

The first neural network model for robotic control was CMAC (cerebellar model articulation controller) [Albus 1975, 1979a,b]. It used input-output measurements stored in a look-up memory to provide feed-forward control signals. CMAC is based on the architecture of the cerebellum which plays an important role in motor learning by storing associative muscle activation patterns for the control of the limbs and body. One of the first neural networks type controllers was introduced into robotics by Albus (1975) and was inspired as a cerebellar model of the brain processes of voluntary movement [Albus 1979a,b]. CMAC (cerebellar model arithmetic computer) was based on this model. It was similar to Rosenblatt's perceptron and was defined by a series of mappings from input-output measurements (eg. end effector position to joint torques as a function of time) - the mapping generated through associative learning was used as feed-forward information to calculate the control signals. It was modelled thus:

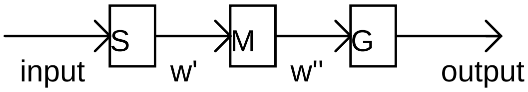

where:

S=input fibre
M=mossy fibre
G=granule cell
P=output Purkinje fibre=f(S)

A unique address was given for each granule cell by virtue of its connectivity to mossy fibres. Each granule cell was associated with a weighting factor w=w'.w''. Purkinje cells output the sum of weights activated as the output $P = f(S \rightarrow M \rightarrow G)$. CMAC offered good performance in spite of non-linearities. The weights were derived by moving a robotic arm through a set of paths. If a large number of paths were stored, it was possible to generate interpolated paths with little computation but the scheme does require a large storage capacity for the static variables of joint angle, joint rate and joint acceleration. Furthermore, the scheme is highly configuration dependent. ANN's are particularly useful as observers for parameter identification of non-linear plants [Narendra & Parthasarathy 1990, Widrow & Lehr 1990]. They are useful for control problems that are so complex that analytical solutions are not readily available. System identification may be regarded as a form of pattern recognition. Pattern classification involves making a set of feature measurements x_i which must be categorised into a number of possible pattern classes w_j. They are generally limited to around 10 inputs beyond which training becomes inefficient and suffer from interference. Most muscle activity involves the coordination of agonistic and antagonistic muscle stiffness. The cerebellar cortex undergoes plastic changes in synaptic connections during motor learning involving all skeletal muscles of the body. The cerebellum receives a wide variety of sensory inputs including proprioceptive, touch, vestibular, visual and auditory stimuli through its mossy fibres.

Due to the requirement for an inverse dynamics model, Kawato et al (1988) suggested that limb kinematic transformations are precisely the computations that are carried out neurally in the human brain in voluntary movement. Areas 2,5 and 7 of the sensory-association cortex perform coordinate transformation from the desired joint angles θ^d and the generation of joint motor commands τ with the actual movements θ fed back to the motor cortex by proprioception. Sensorimotor integration shifts positional frames of reference in spatial coordinates to derive the muscular trajectories. Area 2 produces motor torque commands while areas 5 and 7 provide the Cartesian trajectory and their transformation into joint coordinates. The parietal-association cortex receives visual and sensory information about the world. Neural circuits are capable of computing the non-linear transformation and the inverse dynamics model of the musculoskeletal system through motor learning of inverse dynamics reference models in the cerebellum, i.e. stimulus-response associations. This motor learning provides a non-linear mapping between the torque input τ and the trajectory output θ^d by cross-correlation. The learning of such an inverse dynamics model is computationally difficult so Kawato (2000) has proposed a cerebellar feedback error learning model which is consistent with ocular-following responses – the tracking motion of the eyes important in gaze stabilisation - in the monkey cerebellum. Spatial analysis for motion tracking is performed within the parieto-occipital areas of the cortex. The cortico-cortical (cerebral cortex to cerebellar cortex) pathway corresponds to the feed-forward arc. Parallel fibre inputs to the cerebellum code for the desired limb trajectory and proprioceptive feedback from the muscles. Climbing fibre inputs encode motor feedback commands which activate Purkinje cells which encode sensory error signals for the motor commands. The Purkinje cells utilise the long-term synaptic plasticity as a temporal averaging mechanism triggered by the stimulus from the climbing fibres. The most important point is that a single internal model cannot cope with all varieties of manipulated objects. The cerebellum employs multiple modules to cope with such variety, each of which comprises three components: a feed-forward model, an inverse model, and a predictor. The brain contains many of these module pairs of forward (predictor) and inverse (controller) models. The module recruited is based on the prior probability computed by the predictor which determines the accuracy of its predictability based on the feed-forward model. This then invokes the appropriate inverse model.

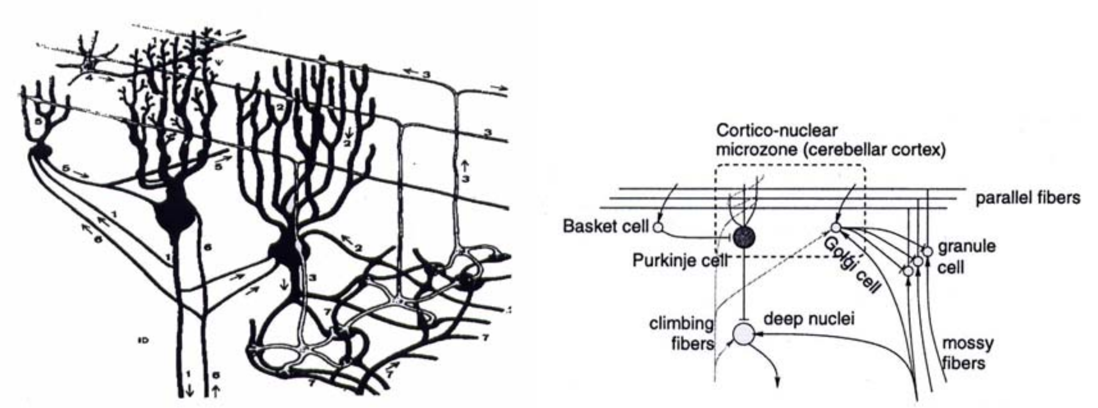

Figure 117 – a) Cerebellar wiring – 1. Purkinje cell, 2. Golgi cell, 3. Granule cell/parallel fibre, 4. stellate cell, 5. basket cell, 6. climbing fibre, 7. mossy fibre (Daya & Chauvet 1999); b) Components of cerebellum (van der Smagdt 2000).

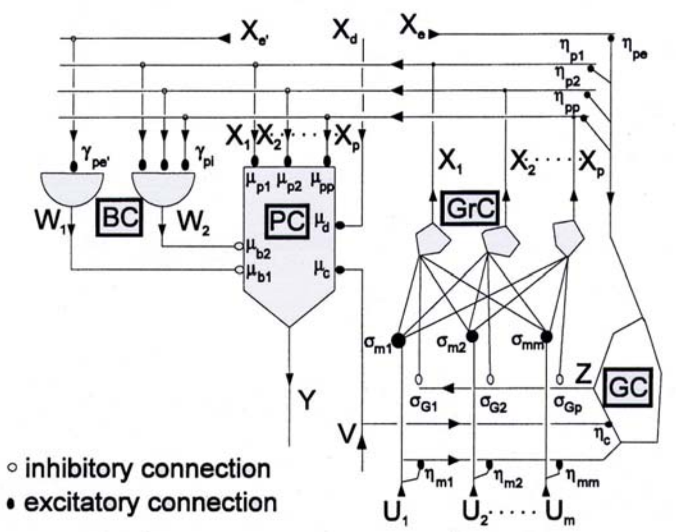

Figure 118 – Pukinje circuit representing basket cells (BC), Purkinje cells (PC), granule cells (GrC) and Golgi cells (GC); inputs are mossy fibres and outputs are parallel fibres (Daya & Chauvet 1999)

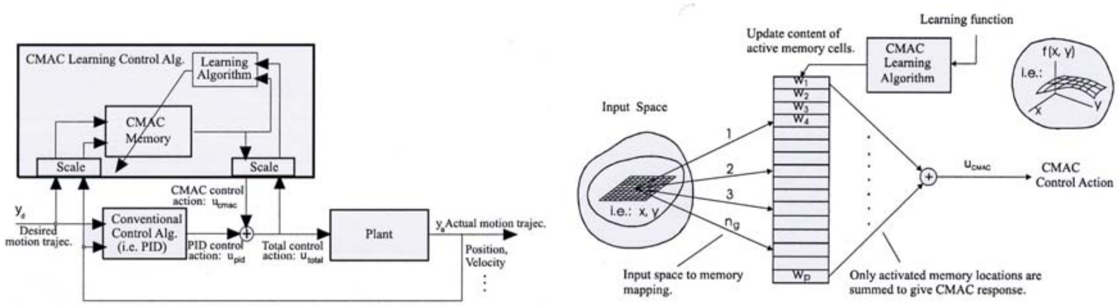

Figure 119 – a) CMAC based control algorithm; b) CMAC neural network (Abdelhameed et al 2002)

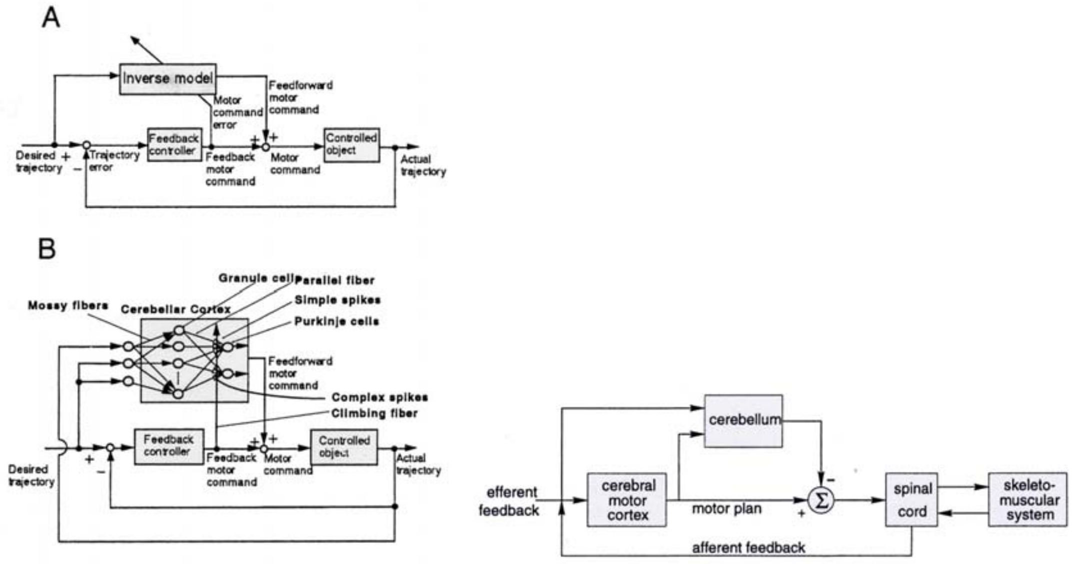

Figure 120 – a, b) CBFELM model of the cerebellum (Kawato 2000); c) Cerebellum as a feedforward filter (van der Smagt 2000).

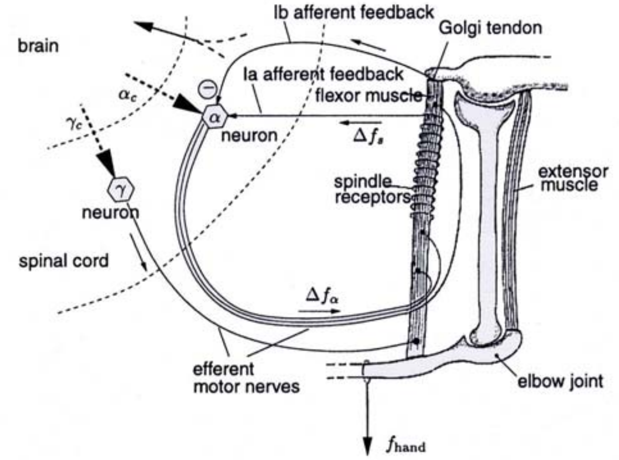

Figure 121 – Simplified model of upper limb control system (van der Smagt 2000)

3.10.5. Associative Learning

Human reasoning is demonstrably associative in that it is content addressable and connectionist architectures are implicitly associative. This implies parallel operation afforded by interactive parallel activation by simultaneous communication of excitatory and inhibitory signals. Learning of motor control may be conducted using a predictive Hebbian reinforcement learning algorithm within the basal ganglia [Jabri et al. 2000]. These patterns are generally learned through post-natal development in humans and encoded into neural structures.

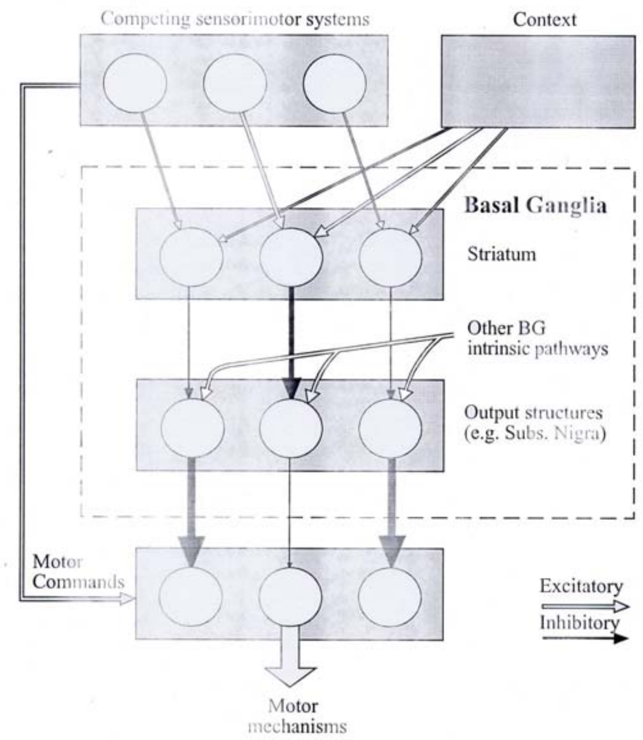

Figure 122 – Basal ganglia as action selection mechanism in which winning sensorimotor mechanism inhibits outputs from basal ganglia which disinhibits motor processes (Prescott et al. 1999).

Central pattern generators for the control of appendages may be implemented through finite state machines or nervous nets. The simplest reflex is a direct connection involving no more than one or two synapses between sensory neurons which recognise stimuli and motor neurons which actuate responses. There are two types of biological learning: associative and non-associative. Non-associative learning involves experience gained from a single type of event, eg habituation. Habituation involves the magnitude of response to a stimulus decreasing with repeated exposure to it. Repeated exposure evokes no response. Habituation is a presynaptic process that results from synaptic depression due to decreased transmitter release. This habituation is exhibited in marine molluscs. Dishabituation involves the sudden restoration of habituated response following a strong stimulus. Sensitisation involves the enhancement of responses to a stimulus following the presentation of a previous strong stimulus. This strengthens its defenses and avoidance reflexes to formerly neutral stimuli now deemed noxious. Sensitisation results from an increased transmitter release and so is the opposite process to habituation.

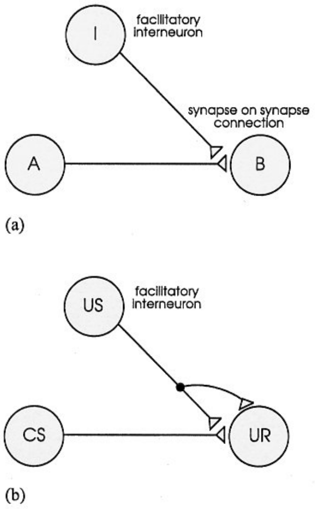

Figure 123 – a) Sensitisation may be modelled as a facilitatory interneuron I with synapse-on-synapse connection; b) Classical conditioning can be modelled similarly (Damper et al 2000).

None of these involves the pairing of stimuli essential to associative learning with reinforcement. Reinforcement learning involves learning behaviours through trial-and-error interactions with a dynamic world in order to extract its regularities [Kaelbling et al 1996]. The agent receives input i that is a limited representation of the current state of the environment, i.e. the environment is partially observable. The agent then chooses an action a as output which in turn changes the state s of the environment. This state transition is fed back to the agent as a reinforcement signal r. The reinforcement signal in biological systems is generally associated with pleasure stimulus which encourages the behaviour, or pain which discourages it. The agent's behaviour is determined by choosing actions that tend to increase the overall reinforcement signal. The agent must learn to find a behavioural policy π in mapping states to actions that maximise the reinforcement measure over time $<\sum_{t=0}^{n} r_t >$. Unlike supervised learning, the agent is informed of the immediate rewards on selecting the action, but given no information on long term advantage. These agents must therefore learn over time the actions that are optimal through experience. Hence, reinforcement strategies are faced with the problem of uncertainty in selecting actions to give high rewards. This policy will change over time as the environment is non-stationary. The n-step optimal policy is to select the best action given that it has n steps to act and increase its reinforcement: $\lim_{n\to\infty} < \frac{1}{n}\sum_{t=0}^{n} r_t >$.

There is a fundamental trade-off between exploitation of past high paying actions and exploration of new actions. Various options are possible, but the Boltzmann machine computes the expected reward $<r_a>$ for selecting action a according to a probability:

$$p(a) = \frac{e^{<r_a>/T}}{\sum_{a'\in A} e^{<r_{a'}>/T}}$$

where:

T=temperature

Reinforcement learning has to explore the problem of finding the optimal action selection policy when the uncertainty exists concerning resultant future reinforcement signals.

Classical (Pavlovan) conditioning is one of the most important principles of animal learning. Classical conditioning occurs when repeated pairings of a neutral stimulus with another stimulus-response causes the neutral stimulus only to evoke the response. Classically conditioned responses (CR, e.g. salivating) are developed according to the temporal contingencies of perceived stimuli (conditioned stimulus, eg. a bell and unconditioned stimulus, eg. food), while instrumental (operant) conditioning are determined by the laws of effect (reinforcement-based response learning are based on rewards/penalties). Classical conditioning deals with unconditioned stimulus (UCS) that generates an unconditioned response (UCR), and a conditioned stimulus (CS) that is associated with a UCS. This CS-UCS relation creates a CR that involves the specific generation of the UCR by the CS. If some given CS precedes another UCS that creates a response, this CS will trigger the UCR. CRs are central to the optimisation of animal behaviour. Such learning techniques can be applied to mobile robots to learn obstacle avoidance to eliminate undesirable "bumps". As learning proceeds, the number of positive reinforcements increases in relation to negative reinforcements: exploring changes towards exploiting [Wilson 1996].

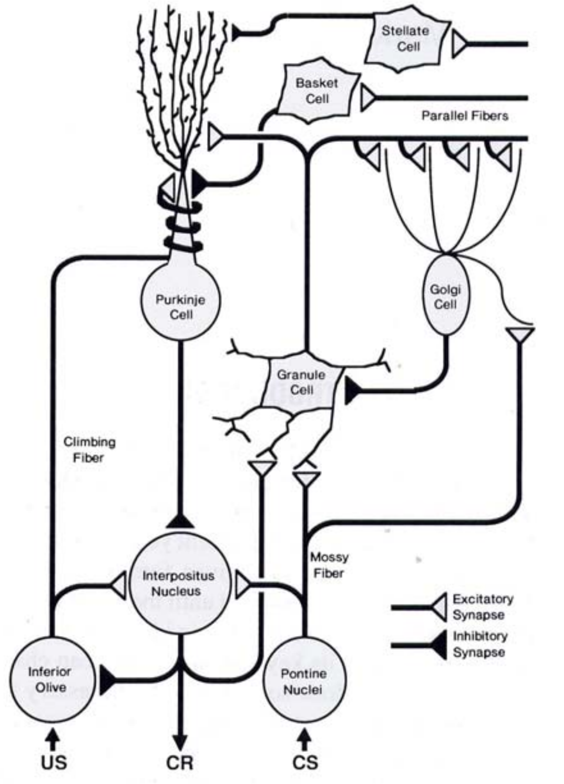

Figure 124 – Neural circuit for conditioned eye-blink reflex – CS information travels from pontine nuclei through mossy fibre projection in cerebellum; US reinforcement pathway projects from

inferior olive to cerebellum via climbing fibres; efferent pathway driving CS projects from interpositus nucleus to motor neurons and also to inferior olive to cerebellum via mossy fibres (Gluck et al 1995)

Verschure et al (1992) implemented robot learning through a neural model of classical conditioning as the dominant form of knowledge acquisition to enable animals to adapt to their environment. It was implemented through connectionist strength adjustment. with the learning rule defining how the weights are updated, and have the general form:

$$w_i(t+1) = w_i(t) + cr_i(t)$$

where:

c=learning constant
$r_i(t)$*=reinforcement signal*

Each time an action is selected by an agent, it changes the state of the environment which in turn emits an observable reinforcement signal r. The reinforcement measures the agent's performance. Hebb's learning rule has the form $r_i(t) = x_i(t)y(t)$ but cannot produce the temporal characteristics of classical conditioning. Verschure et al (1992) developed neural models of classical conditioning using the Widrow-Hoff delta rule:

$$r_i(t) = \big(z(t) - y(t)\big)x_i(t)$$

where:

$y(t) = \sum_{j=1}^{n} w_j(t)x_j(t)$ = *trace of response*
z(t)=teacher input signal (UCR or CR)
$x_i(t)$*=temporal trace of CS*

y(t) is a feedback term representing expectation as the weighted average of past values of the output (similar to a Kalman filter output). Rosenblatt's perceptron is similar but represents binary scheme. If $z(t) = \lambda$=asymptotic value of the connection weight then this scheme corresponds to the Rescorla-Wagner model of classical conditioning [Rescorla & Wagner 1967]. The Widrow-Hoff rule offers a model of causality which enables it to emulate classical conditioning (Figure 125). The Widrow-Hoff rule has been discovered with 98% learning efficiency by genetic algorithm search with the strings encoding weight changes [Jones 1991]. It resembles stochastic approximation by the Robbins-Monro algorithm. Genetic algorithms have been used to design ANNs and indicate how evolution can give rise to learning rules [Jones 1991]. Supervised learning was applied to a fully connected single layer feed-forward net with signal output units with thresholds. A gene string encodes the weight space dynamics thereby invoking a potential learning procedure. The idea is for the GA to evolve the learning rule that maps from the input patterns to the output patterns. The fitness of the learning rule was determined by its performance in a number of tasks. Each learning rule is assigned a fitness between 0 and 100%. At the start, all individuals in the population have a fitness spread around 50% representing random preferences. After a few generations, differential reproduction finds small differences and the fitness of the best individuals rises until a maximum fitness of 80-98% with a mean of 92% is achieved. The Widrow-Hoff delta rule with a fitness of 98% and variations of it has been discovered by such means. Diverse environments as characterised by a number of tasks can generate the evolution of general learning procedures to adaptation to cope with the variable environment.

The adaptive network learned to associate CS's with specific US's from predefined US-UR reflexes and these substituted CS's can trigger the original response. The connections between the US input patterns and the UR output patterns were prewired and not modifiable as they represented basic reflexes. The CS input pattern was correlated to the US pattern by modifiable weights representing synaptic plasticity to provide the learning mechanism to allow the system to associate the CS patterns with the US patterns. The CS replaced the US in involving the CR similar to the UR in that the two responses occur via the same response pathways being activated by two different stimulus pathways. The adaptive elements of the neural model increase the response rate to increased stimulation during learning. They demonstrated successfully robotic avoidance behaviour. When the robot collided with obstacles the collision sensors triggered avoidance behaviour. Over time this collision sensor became associated with a range sensor input. All collisions triggered the sense-act reflex as the basis for the associative mechanism to accumulate experience over time as control gradually shifted from the collision sensor to the range finder. Essentially the range sensor input alone eventually triggered avoidance behaviour. During this process the robot behaviour got smoother as the robot began "anticipating" obstacles - this was an emergent property those arose when the response was triggered by lower activity due to weight increases. The weights connecting the CS and US fields provided the internal activation pattern representation of the system - environment interactions as higher order field at a more abstract level than the reflex pattern. Sutton & Barto (1981) introduced a more sophisticated model of classical conditioning which may be implemented similarly:

$$\bar{x}_i(t+1) = \alpha\bar{x}_i(t) + x_i(t)$$

where:

$$\begin{aligned}&0 \leq \alpha\\ &\bar{y}(t+1) = \beta\bar{y}(t) + (1-\beta)y(t) \quad \textit{where} \quad \beta < 1\\ &w_i(t+1) = w_i(t) + c\big(y(t) - \bar{y}(t)\big)\bar{x}_i(t)\end{aligned}$$

Each stimulus x_i has associated with it a separate stimulus trace $\bar{x}_i$ and similarly for the output signal y(t). The constants α and β relate to the size of the temporal integration window. These trace signals are prolonged to simulate the temporal characteristics of classical conditioning. They model the latency effects of chemical concentration traces in neurons that persist after the CS ends and the US starts. These trace signals are prolonged to simulate the temporal characteristics of classical conditioning. A weight increases when a CS occurs before the UCS and creases it if the predicted UCS does not occur. There are thus three stages to stimulus conditioning [Vico et al. 2001]:

1. acquisition phase – the UCS onset follows the CS offset yielding UCR and an increase of the association between the CS and the response;
2. prediction of UCS – the CS generates a CR that is followed by the UCR generated by the UCS, the response is correctly anticipated and the weight does not undergo any change;
3. extinction phase – the CS generates a CR but the UCS does not follow after the CS offset which causes a decrease in the synaptic weight

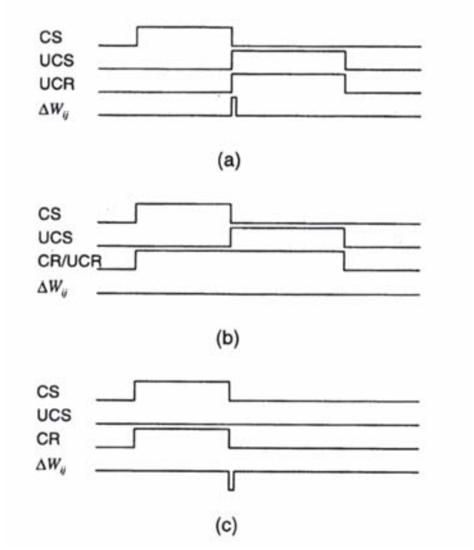

Figure 125 – Three stages of stimulus conditioning: a) acquisition; b) UCS prediction; c) extinction (Vico et al 2001)

Neurotransmitter binding to receptor sites of the neuron regulates the postsynaptic concentrations of cyclic AMP and cyclic GMP to mediate slow postsynaptic potentials of the neuron. They cause a transient increase in synaptic efficacy by altering the firing threshold of the postsynaptic neuron by altering the conductance of Ca^{2+} ions in the neuron electrolyte by depolarisation. This means that simultaneous presentation of the CS and the US is not required and that the interstimulus interval between the CS and US can be simulated. The connection weights are modified by reinforcement only if the current value of y differs from the trace value $\bar{y}$ representing the average of the past values. This rule detects correlations between traces of input stimuli and changes in output. It is equivalent to the Widrow-Hoff/Rescorla-Wagner rule except that it substitutes output for the teacher signal and thereby retains the output substitution of Hebb's rule.

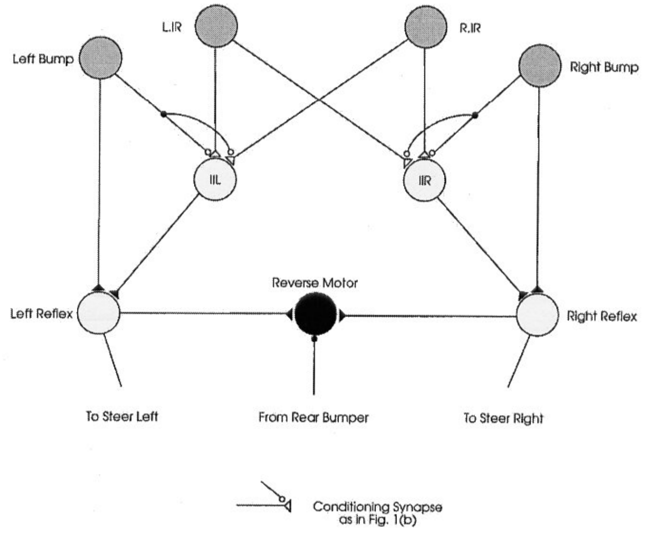

Figure 126 – Conditioning may be added to a mobile robot control network through conditioning connections to synapses of ipsi-lateral sensory neuron (Damper et al 2000)

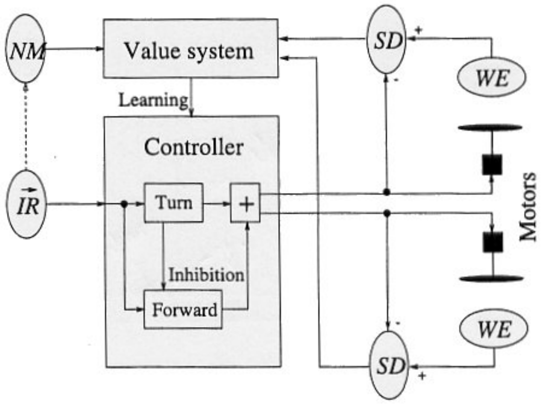

Figure 127 – Self-organisation-through-perception architecture which comprises a value system which uses the difference between motor activation and wheel sensors to recognise collisions (Salomon 1998)

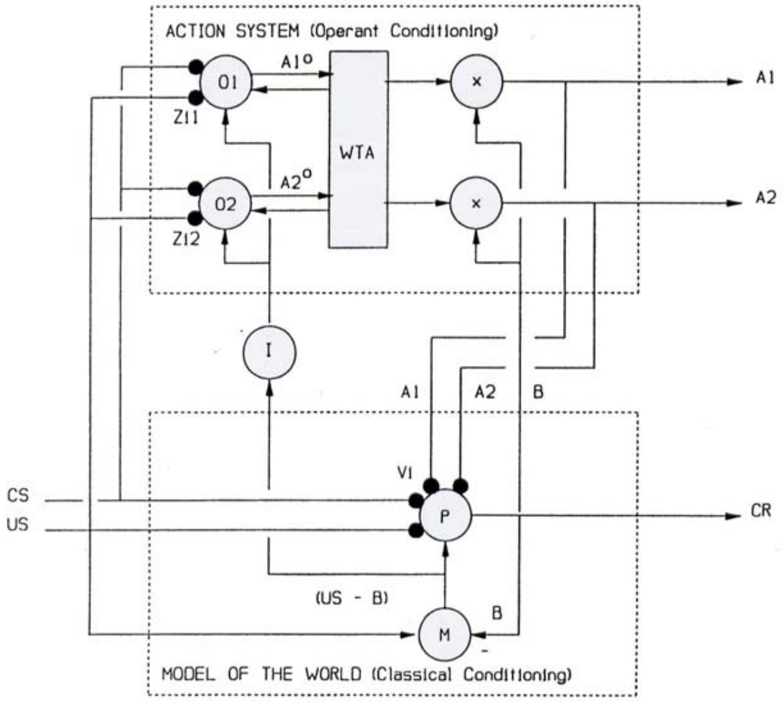

Figure 128 – Neural network for avoidance learning – classical conditioning builds an internal world model which generates predictions of US; operant conditioning establishes action system which determines behavioural strategy by selecting avoidance responses; mismatch between actual and predicted US modulates classical and operant association (Schmajuk et al.1999).

3.10.6. Concept Learning in AI

Barlow (1970, 1983) defined intelligence as "the capacity to detect order among different sets of information, essentially reducing the redundancy between the sets of information". Again, this implies inductive inference. Dreyfus & Dreyfus (1987) considered learning to be critical for common-sense knowledge. Learning involves organisational change to the agent's structure giving it adaptive plasticity to modify future behaviour as a result of the past history of experience based on some performance metric, i.e. learning provides the means for the agent to act more intelligently in new situations based on its past experiences [Russel 1991]. Learning introduces a means for modifying and extracting the range of behaviours available to an agent on the basis on new information about the environment as it attains its goals. Learning consists of two basic forms: knowledge acquisition as the learning of symbolic information coupled with the ability to apply that information in an effective manner, and skill refinement as the gradual improvement of motor and cognitive skills by repeated practice. Knowledge acquisition is a conscious process and is characteristic of declarative knowledge but skill acquisition is unconscious and characteristic of procedural knowledge. Indeed cognition is based on two abilities: the ability to store information (memory) and the ability to modify information (learning).

One of the first learning programs was Samuel's draughts playing program [Samuel 1963]. Learning was implemented by simulating the mechanism of evolution by natural selection. A scoring system evaluated the strength of a position. Two versions of the program played each other in competition with the scoring systems that were slightly different. The winning version was retained and was used as the basis for constructing new variations. The HACKER plan generator invoked a limited form of learning. A knowledge base of generic plan solutions was stored and used for new circumstances. New successful plans were generalised by variablisation of constants and chunked into macrooperators (as in STRIPS) and added to the knowledge base. Simulated execution then tested the plans. If algorithmic planning is viewed as an automatic programming system where subroutines comprise the subgoal hierarchy then three possible outcomes can occur: success in generating a plan, failure when no plan is found after exhaustive search, or non-termination as the plan expands ad infinitum never converging to a solution (the equivalent of Turing's halting problem). If a plan is produced it correctly solves the problem but if failure or the non-termination state is generated no solution exists. Since any algorithm with its inputs can be encoded as a Turing machine, the planning problem is undecidable and no upper bound exists on the time required to solve a problem due to the existence of the possibility of infinite loops [Chapman 1987].

The amount of time required to program an agent with the equivalence of human intelligence would be very large: the human brain has an information capacity (assuming that the neuron is a finite state machine) of 10^{15} bits since each of the 10^{11} neurons of the brain can store 10^4 bits. Assuming 10% capacity is used, to program a machine with 10^{14} bits at one line of code per hour (rate of production of software from conception to installation) with an average line of code some 500 bits long would require 100 million man-years. This may be compared with the information input rate of the human vision sensors of 250 Mbps which may input 10^{14} bits in 20y. True intelligent capability implies the need for a vast repository of encyclopaedic knowledge about the world which humans acquire through learning. This exercise illustrates the importance of learning [Bock 1985]. Charniak (1988) questioned whether such a "mind" program would be verifiable or debuggable due to its size and complexity but that a functionally identical program may be possible but with a much simpler structure – indeed, this is an implicit assumption in AI. One of the first learning programs was Samuel's draughts playing program in 1960. Learning was implemented by simulating the mechanism of evolution by natural selection. A scoring system evaluated the strength of a position. Two versions of the program played each other in competition with the scoring systems that were slightly different. The winning version was retained and was used as the basis for constructing new variations. The HACKER plan generator invoked a limited form of learning. A knowledge base of generic plan solutions was stored and used for new circumstances. New successful plans were generalised by variablisation of constants and chunked into macrooperators (as in STRIPS) and added to the knowledge base. Simulated execution then tested the plans.

There are two main forms of symbolic learning: inductive learning and analytic learning with the latter including explanation based learning and analogical learning. Inductive learning involves generating a general concept representation from a set of positive and negative examples. The positive examples induce generalised descriptions of the concept for categorisation while counter-examples prevent over-generalisation. In this way it learns discriminant properties of concepts and thereby generalise. The source of examples may be a teacher or the environment. The environment will produce examples randomly and is the only learning strategy available to autonomous systems. This is a form of unsupervised learning requiring large degrees of inference: passive observation of the environment and active experimentation on the environment as discovery should both be employed to suit the learning environment. In real systems, noise has a tendency to distort well-defined data necessitating the use of statistical techniques to enable classification. A set of tests seperate all examples of the concept from all known concepts. Autonomous learning employs incremental learning whose concepts are derived from data previously stored but which may later be modified by further data as necessary. This is a fundamental problem in hypothesis generation and confirmation. Popper views inductive hypotheses as unprovable (in the positive sense) but disprovable (in the negative sense). Bayes theorenm is the usual approach to induction: if evidence E exists from trials, then the probability of hypothesis H may be defined. A fundamental philosophic problem concerns that if a property is observed to be true in a finite number of observations, it does not necessarily follow that this property will hold in the future (White 1988). This is not only a problem of logic, but also of the changing external world.

The ability to generalise through extracting important common features of a class from a set of observations is a fundamental component in learning. Bundy et al (1985) provided a comparative overview of symbolic learning techniques. Concept learning is considered a subset of rule learning. Concepts are generated by forming symbolic descriptions of the desired concept from examples and non-examples and use this for prediction. Rule learning involves modification of existing rules and mostly involve conjunctive rules since any disjunctive rule can be split into conjunctive rules. Much modification involves weakening a rule by weakening its instantiation or by giving it stronger conditions. The main control loop for rule learning involves:

- identifying the fault within the rule using a critic, and
- modifying the rule to remove the fault using a modifier.

Faults are identified by running the existing rules on a problem and identifying where they behave correctly (positive training instance) and where they behave incorrectly (negative training instance) in the resulting rule trace. Positive instances generalise the rules while negative instances correct them. Negative instances may be errors of commission whereby a rule is fired because it was insufficiently constrained or errors of omission where a rule does not fire because it was incorrectly constrained (or did not exist). The modifier requires information on the type of response (positive or negative instance), the rule responsible and the context in which it fires (i.e., the variable bindings). The credit assignment problem of deciding which rules are responsible for aspects of the program's behaviour is determined by an ideal rule trace, i.e. the correct sequence of rules. This ideal trace is compared with the actual trace to locate the divergence point where the fault rule occurs. At that point the rule that fired traces an actual trace R_a which differs from the rule that traces the ideal trace R_i exhibits error of commission with respect to R_i which exhibits an error of omission. The error of commission is then corrected. The program reconstructs the ideal trace by pruning the actual trace search tree. There are 3 methods of modifying the faulty rule:

1. imposing order to the rules to keep the actual trace in line with the ideal trace such that $H \rightarrow C$ is fired over $H' \rightarrow C'$;
2. adding extra conditions to the rule hypothesis to make it more specific such that $H \rightarrow C$ is replaced with $H \& H' \rightarrow C$ a special case of which is the instantiation of a rule where $H(X,Y) \rightarrow C(X,Y)$ is replaced by $H(X,X) \rightarrow C(X,X)$;
3. updating a rule hypothesis to account for new instances such that $H \rightarrow C$ is replaced by $H' \rightarrow C$ where H' is derived from concept learning methods.

Method 2 involves finding the difference between the rule which under certain variable binding conditions gives the correct response. It is then necessary to find a discriminant H' that is true under the correct response conditions but false under the incorrect conditions to form a new rule $H \& H' \rightarrow C$.

Mitchell (1983) characterised inductive generalisation as a bi-directional breadth-first search. A matching predicate matches generalisations to specific training instances of positive and negatiove examples. It attempts to map feature vectors of ganeralisation g onto feature vectors of instances i such that their respective restrictions are consistent (ie. identical). The search space is characxtertised by the "more specific than" relation such that G1 is more specific than G2 only if G1 matches a subset of instances that G2 matches: {i M(G1,i)} {i M(G2,i)}. This partially orders the generalisations in the hypothesis space. The specific to general breadth first search generates a set of alternative plauisible hypotheses S={s} such that s is consistent with observed instances i and that no generalisation is more specific than s. This set of increasingly general hypotheses is derived from positive training samples. The general tro specific breadth first search generates a set of hypotheses G={g} such that g is consistent with observed instances i and that no generalisation is more general than g. This set of increasingly specific hypotheses eliminates from the set s and prunes the search tree - it is derived from negative training instances. This bidirectional search produces a version space of generalisatiuons x such that x is more specific than the members of G but more general than the members of S.

Analytical techniques allow learning from a few examples and use deduction and background domain knowledge rather than induction. It involves extending existing knowledge to situations that have high similarity. Learning by analogy consists of tranforming existing knowledge for one domain for use in related domains. This method is most appropiate for increasing the efficiency of rule-based systems where reasoning and inference performs well in well structured environments. Analogy is based on the assumption that if two situations of different domains are similar in some respects, they will be similar in other respects, ie. reasoning from a model. Winston (1980) suggestes that several properties are essential for learning by analogy. Firstly, situations must be represented by relations between sets of parts. Between two situations corresponding matches must be established between imporetant constraint relations, eg. causal relations CAUSE and CAUSED_BY. Important relationships between objects or classes of objects are relations such as "a-kind-of" and "has-property". In a frame-based representation, agent-act-object may be made explicit with variable slots. Matching is restricted by pairing only frames of the same class. Similarity measure is a measure of overlap and is defined between features or frame classes A and B.

Explanation based learning (EBL) is an analytical and knowledge-intensive form of learning. It addresses the problem of generalising from single training examples (as opposed to inductive generalisation which requires many training examples). It can improve problem-solving performance from single examples in the context of background knowledge [Minton et al 1989, Mitchell 1983]. The ability to generalise implies the ability to explain how a training example is a member of the concept being learned. It analyses why a single example is an instance of a concept by deriving the relevant features of the example as sufficient conditions of the concept. Generalisation is a search through the state space of possible concept definitions for the target concept. The purpose of the knowledge is to constrain that search. The concept description enables instances of the concept to be recognised efficiently. Domain knowledge encapsulating specific knowledge about objects is used to guide the learning process. EBL has its roots in programs such as STRIPS and HACKER and subsequently SOAR which used chunking through macro-operator formation. EBL provides a declarative justification of the generalisation obtained from its domain knowledge similar to truth maintenance. Explanations essentially constitute proofs requiring a clear specification of the input and output. The input has several components: the target goal concept definition to be learned as necessary and sufficient conditions for an example to be an instance of the concept; a training example of the concept; a domain theory of rules and facts for constructing explanations of how the training example is an instance of the concept; an operationality criterion depicting the utility of the example. The attributes required in the explanation define the concept. An explanation is constructed on how the example satisfies the goal concept definition. Generalisation determines the sufficient conditions under which the explanation holds. Explanation is a proof that the instance is a valid example of the concept. The example guides the search and improves the problem-solving performance by analysing the example features in terms of how they satisfy the concept definition. Next, backpropagating the goal concept occurs step by step through the explanation structure. This determines the necessary and sufficient conditions for a rule to be used to infer a formula. This explanation is generalised into an operational description of sufficient conditions by backtracking from the concept through the proof structure substituting constants with variables. The chain of inferencing provides the explanation structure on how the example is an instance of the goal concept definition. Hence there is a similarity between theorem-proving (chunking) and EBL. A specific example of the generalisation of an explanation is given by: "generalisation-to-n" for repetitive events. PRODIGY is an EBL-based system which learns domain-independent, metalevel control rules from a problem-solving trace; analogy replays solutions to similar past problems; multilevel plan extraction and experimentation (simulation) are invoked when plan execution monitoring detects departure from expectations.

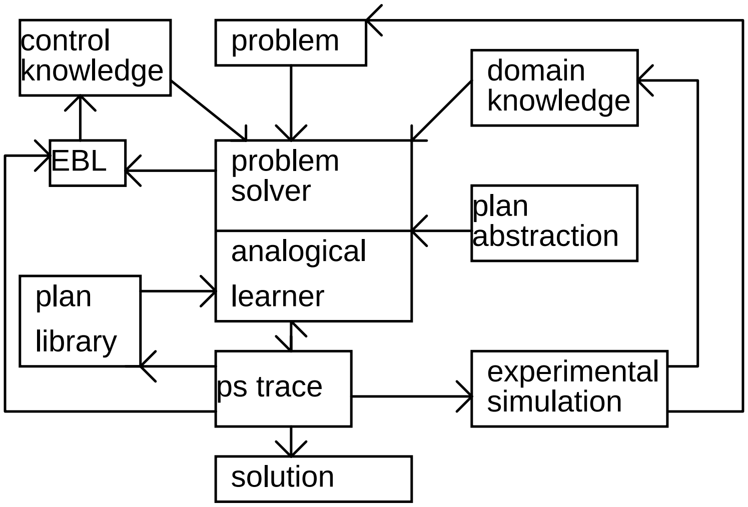

Figure 129 – PRODIGY, an EBL-based architecture

The task domains include blocks world expert, a 3D robotics construction domain and task-specific experts. The problem solver used predicate logic and performed inferencing to manipulate knowledge and applied operators to generate actions on the world. Metalevel control knowledge was separate from domain specific knowledge as it constrained and guided problem-solving searches. The construction of an explanation effectively created a search control rule by specialising concepts in proving that an example was an instance of the target concept. This reduced the cost of matching. The utility of the search rule is determined by the improvement in the search time by the rule. Segre (1991) showed an explanation based learning could be applied to learning mechanical robot assembly plans. Acquiring new concepts involves using a priori declarative domain knowledge to analyse previous solutions to problems to discover the reasons for their success. The explanation is then generalised for future problems. It builds macro-operators compiled from chunks of problem solving knowledge which modify the problem space explored, i.e. they act as sophisticated search control heuristics to reduce search for solutions to problems based on sample solutions. Domain knowledge is stored as schema data structures stored in a schema library and it is used to construct solutions to problems and produce new schemas. The current world state comprises a collection of such schemas. A schema planner searches through the problem space to produce a plan which is simulated. This process of planning and execution is interleaved. The learning element comprises two modules: the understander and the generaliser. The understander uses natural language story-understanding methods (i.e. scripts) to construct a causal model represented as a semantic graph of relations between objects. A verifier constructs the explanation by forward chaining theorem proving to find the schema representing the goal. The generaliser produces a new composite schema by generalising the goal schema to cover a greater number of contingencies. This is done by applying transformation operators which preserve validity to modify the schema explanations such as "replace ground values by variables". Learning universal functions such as the XOR function which is widely regarded as difficult and is rarely used in natural systems.

Zhou (1989) presented the classifier system with memory (CSM) as a model of human cognition. It was similar to the classical classifier system as a general purpose rule-based genetic learning system. However, classifier systems offer degraded performance under rapid environmental changes as classifier systems lose inactive but potentially useful rules under such circumstances, i.e. a problem of forgetfulness. CSM provides memory to allow long-term cumulative learning and robustness to rapid environmental changes. CSM differentiates between the temporary STM where rules are manipulated and the permanent LTM.

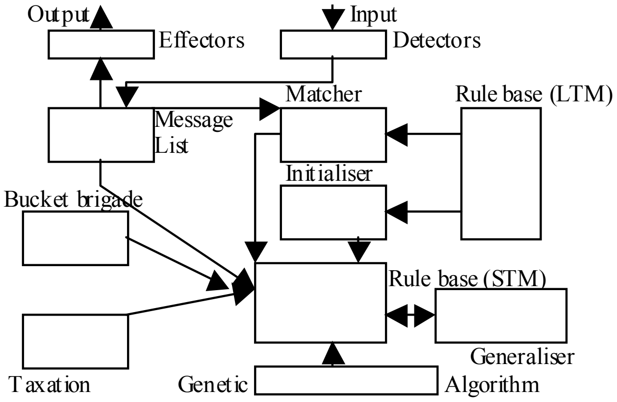

Figure 130 – Classifier system with memory (Zhou 1989)

CSM attempts to recognise a problem and searches its LTM for a ready solution. If no solution is available due to incomplete knowledge, CSM uses analogical problem solving to utilise any similar solutions in LTM. A matcher unit uses analogical reasoning to select solutions of similar past problems from LTM through partial condition matching to measure prototypicality and relevance. If there is no relevant solution, the system then used the standard classifier system with the genetic algorithm component to solve the problem. An initialiser provides a small initial seed population for STM from LTM to heuristically guide learning in promising directions for the genetic algorithms where they are used according to match (if no match, then the initial population is totally random). Once solved, successful rules are generalised and chunked and transferred from STM to LTM for future use. Successful rules are generalised by the generaliser dropping conditions and variablising constants. The use of a separate LTM of inactive rules inhibits the problem of forgetfulness allowing the preservation of past accumulated knowledge. STM is the working memory containing only active knowledge relevant to the current problem and is subject to modification. Hence, recency of use is no longer the dominant criterion of usefulness. CSM was successfully tested with a simulated robot for navigation tasks in the presence of obstacles. CSM was capable of extracting components of specific rules and constructing task-independent knowledge which could then be used for other navigation tasks of greater complexity.

3.11. GENETIC & NEURAL APPROACHES TO LEARNING

Algorithms based on biological analogies have become commonly adopted for solving difficult problems: genetic algorithms [Holland 1975; Goldberg 1989] (and variants such as genetic programming [Koza 1993] and evolution strategies [Rechenberg 1973] – collectively known as evolutionary algorithms) provide the basis for the evolutionary search of solutions to NP-hard problem spaces; neural networks (highly simplified idealisations of biological neural architectures) provide the basis for the generation of non-linear functions for robust pattern recognition and for control systems based on non-linear dynamics. Evolutionary algorithms are a range of machine learning mechanisms inspired by biological evolution. The genetic algorithm typically codifies solutions to the problem at hand as numeric strings, which can be thought of as artificial genotypes (bit strings were once the dominant representation, integer or real number encodings are now widely used). A population of these strings are subjected to a selection process whereby the best solutions are more likely to be chosen. A new population (the next generation) is created by applying various operators to the strings selected from the current population. These operators include: mutation (involving random changes to some of the symbols in the string), cross-over (involving combining material from 2 or more 'parent' strings), inversion (involving reversing the order of a randomly selected sub-string), and many other special operators specific to individual problems and solution encodings. The mapping between the strings and a solution to the problem can be straight forward (a simple list of values for variables) or complex (the definition of a 'developmental' process whereby the string is converted into a solution). Selection operates according to fitness function, or fitness metric, which measures the relative merits of competing solutions defined in the population. Over a series of evolving generations, increasingly fit solutions to the problem are generated. Genetic programming relies on a very similar scheme but adopts a high level language representation scheme (such as Lisp) which can be represented as tree structures. Genetic programming representations do not always utilise mutation as it can be difficult to define suitable operators that do not result in illegal programs most of the time. The processes involved in an evolutionary algorithm are illustrated in Figure 131. Here the problem is to evolve the control system for a robot, hence the fitness measure is related to the behaviour generated in a robot by a control system specified by an artificial genotype string.

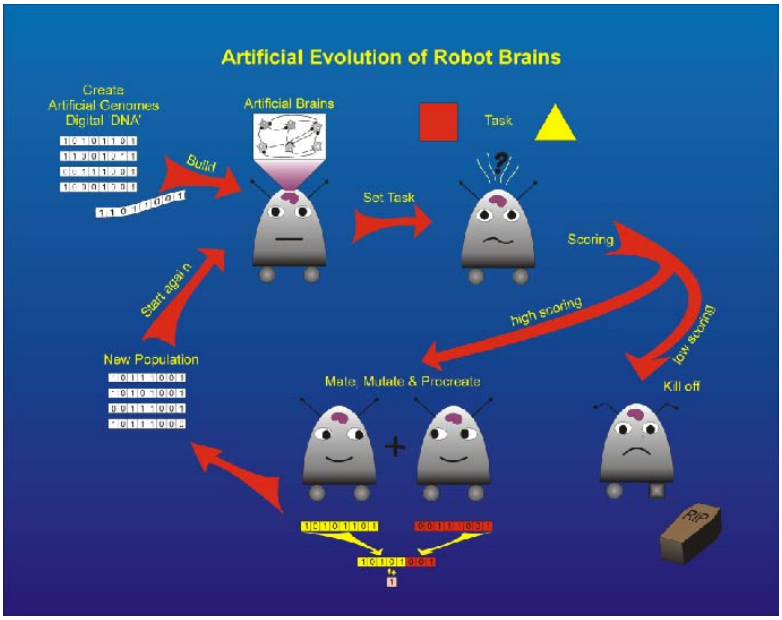

Figure 131 – An illustration of an evolutionary search algorithm applied to the problem of developing a robot control system.

Classifier systems are a machine learning architectures that employ a genetic algorithm where bit strings codify situation-response rules similar to an expert system [Booker et al. 1989] and have been used to implement behaviour control through reinforcement learning – they are algorithmically equivalent to Q-learning [Dorigo 1993]. They often employ a bucket-brigade algorithm to assign credit to successful rules which contribute to the solution of the problem. There has been some recent work on developing more powerful stripped down version of the original classifier system algorithms [Wilson 1994, 1995]. In some circumstances learning is a highly effective means for guiding evolutionary search by altering the shape of the search space [Hinton & Nowlan 1987]. The genetic landscape represents the phase space of all possible genetic combinations. Learning over the single individual's lifetime is much faster than evolution over many generations within species and involves less energy expenditure than the creation of a whole organism. Learning requires long individual lifetimes while evolution favours rapidly passing generations. Both require feedback of some form – reinforcement learning gives right/wrong feedback signals while evolution gives survival/death feedback signals [Ackley & Littman 1991]. A neural net which requires precise connection schemes may be represented as a space of high fitness in the genetic search space. Rather than specifying the phenotype directly, the genes code for an adaptive process to acquire the desired function. If the phenotype specifies some of the connection requirements while other decisions are required by learning, then the evolutionary search space is restructured by providing a zone of increased fitness around the spike of the required network connection scheme. Evolutionary reinforcement learning involves a genetic code specifying the initial values for the weights of an action network that maps between sensorimotor innate behaviours and the weights of an evaluation network that maps between the sensors and a scalar-value of good/bad reinforcement output. The action network weights are adjusted through complementary backpropagation reinforcement learning but the weights are fixed for the learning cycle being determined genetically. The inherited evaluation function converts long-scale feedback into short-scale feedback. When the reinforcement signal is positive, the output vector becomes more probable for the same input while negative reinforcement reduces the probability of the output vector for that input. Learning provides local hill climbing as an inner loop to guide the global genetic search of evolution. It was found that learning is dominant initially but after a period of reproductive generations, learning becomes incorporated genetically. In this context, learning and evolution together are much more successful than either alone. More recently there has been a growing interest in using evolutionary algorithms to develop more complex adaptive systems by evolving the way in which the system is plastic. The system's plasticity, the way in which key properties change in response to the environment, underlies the mechanisms used for learning and other forms of adaptation. Various stands of this research, especially in connection with the development of robot control system, are covered by later sections of this document.

The **evolutionary robotics** approach is an automatic design procedure whereby a robot's controller, and possibly its overall body plan, is progressively adapted to the specific environment and the specific problems it is confronted with, through an artificial selection process that eliminates ill-behaving individuals in a population while favouring the reproduction of better-adapted competitors. Such a process calls upon some evolutionary procedure as outlined above and illustrated in Figure 131. It involves a population of genotypes (i.e. of information that evolves through successive generations) and a phenotype (i.e. the robot's control architecture, its body plan, and its behaviour) that is encoded in any one genotype. A dedicated fitness function is used to assess the behaviour of each individual in the population and to direct the selection proper. Dedicated operators such as mutation and crossing-over give rise to new genotypes in the population and permit robots of ever-increasing fitness to be generated, until the process converges to some local or global optimum. In the majority of applications, the evolutionary procedure is performed in two stages: fitter phenotypes are first sought through simulation and are then downloaded in turn on a real robot to check their fitness with respect to real world constraints. However, in some other applications, the evolutionary procedure takes place through evaluations performed directly on the robot and fitnesses are directly assessed through real world interactions. In both cases, software controllers can be evolved. They may be implemented as control programs (in a high level language or in machine code), as a variety of production-rule systems, or as neural networks. Finally, within the so-called evolvable hardware approach [Sanchez and Tomassini 1996; Higuchi et al. 1997; Thompson 1997], genotypes code for the configuration of hardware controllers and body plans, and fitnesses are also assessed through real world interactions.

Dorigo (1993) implemented behaviour based control methods using genetic algorithm based classifier systems. The classifier system provided a mechanism for reinforcement learning based on the performance of robotic behaviour and generate feedback concerning the success or failure of actions. At the heart of the classifier systems architecture is a set of condition-action rules constructed from “chromosomes” based on symbol strings of the form: if <sensor data> and <current goal> then <action> formed from a three-valued alphabet A=(0,1,*). The classifier rule is composed of a number of chromosomes which represent the condition-action information as a string of three chromosomes. The first two chromosomes comprised the condition part while the third is the action part. The classifier may send messages to the effectors directly or to other classifiers. A set of such rules comprises the knowledge base and new rules may be generated from past experience. A message list collects all the messages to and from the environment via an input-output interface of detectors and effectors. Environmental messages in the message list are matched against the condition parts of the classifier rules and those that match are activated in parallel. A feedback mechanism provides reinforcement learning via reward/punishment. An apportionment of credit algorithm distributes the payoff to the rules that performed the actions. The GA provides a means for rule discovery through modification of the population of solutions. A fitness function relates rule solutions to performance. The value of this fitness function is determined by the rules' usefulness as calculated by the apportionment of credit algorithm. The worst classifiers are replaced by m new classifiers. In this way, an internal world model is built up corresponding to factors in the external environment. The apportionment of credit algorithm classifies rules according to their usefulness (determined by the fitness function). A variable real valued strength is associated to each classifier. When a classifier invokes an action on the environment a payoff is generated the value of which depends on the effectiveness of the action. For multiple and sequential rules, the bucket brigade algorithm passes reward back through the decision chains so that the payoff flows backwards. The system may be implemented as several classifier subsystems acting in parallel where each classifier subsystem learns a simple behaviour (such as avoidance of obstacles, grasping of objects, etc) through interaction with the environment. Sensors monitor the environment while rules determine which actions are to be taken. The system as a whole learns to coordinate its activities at a higher hierarchical level. As hierarchical structures learn faster than single level distributed architectures, an hierarchical structure of behaviours is produced. The hierarchical structure has the lowest level classifier system learning basic behaviours from sensory data while those at a higher level learning to coordinate those behaviours with no direct access to the environment. Such coordination learning leads to more complex behavioural sequences. Each classifier system receives excitatory and inhibitory signals from the classifier systems connected to it. Classifiers at the same level compete. This enables the system to distribute behavioural knowledge across the architecture. The experimental unit was successful at following lights and avoiding heat sources. The coordination level learned a switching procedures between the two low level behaviours. Essentially it may be regarded as a coarsely structured neural network. Indeed, it may be that the higher level coordinating behaviours could have subsequently evolved symbolic reasoning capable of manipulating the classifier rules as a superstructure built onto the instinctive behavioural base. This approach was extended and applied to a small robot AutonoMouse which had two eyes. Each classifier system of 100-500 rules was implemented on a distributed set of transputers as its brain, each with a specific behaviour arranged hierarchically (Alecsys) [Dorigo 1995]. It successfully exhibited simple light following, evasion of sounds, homing, and coordination of switching between these simple behaviours.

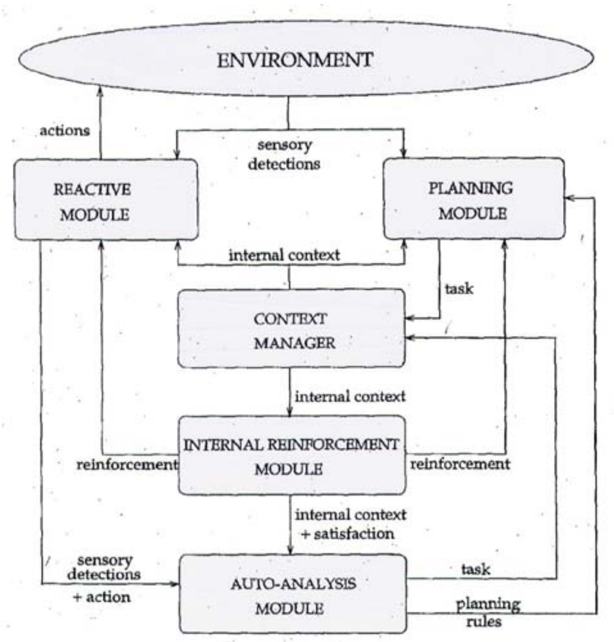

Figure 132 – MonaLysa Architecture (Donnart & Meyer 1996)

3.11.1. Evolutionary Neural Networks

Evolutionary neural networks apply biological analogies more strongly by hybridising the two approaches – the structure of the neural network is encapsulated into a genetic code and populations of these neural network structures evolve according to their success in providing appropriate control systems for robots. This is a powerful technique as it exploits the evolutionary process to design the topology of the neural network which traditionally must be partially designed by hand – a black art at best. GA's can be used to search for optimal weights of an ANN to correctly classify training examples. A linear string of weights of each link can be encoded as real numbers. The ANN population is generated with random weightings. Each ANN has an associated fitness indicating its classification accuracy determined by the fraction of set of training examples it correctly classifies. Each generation is tested on the training set (or perhaps a random subset of the training set) to determine the fitness of each individual. Representation is characterised by four properties:

(i) neuronal activation level

(ii) synaptic weights relating neurons in the network

(iii) dynamic trajectories of activation patterns across neurons

(iv) genetic representation of neural architecture.

The ANN structure is responsible for generating functionally fit behaviour triggered by environmental input and it is determined by the adaptation of the synaptic weights. The development of the species adaptation algorithm (SAGA) provided the basis for using this technology in open-ended contexts by allowing variable length genomes [Husbands et al. 1997]. Evolutionary robotics uses this approach with great success though its scaleability to complex behaviours is currently not clear. Edelman's neural Darwinian theory of neuronal groups attempts to incorporate selectionism into neural dynamics directly. Edelman's theory of neural networks asserts that an extremely rapid selection process occurs within the brain during behavioural learning similar to the way in which the immune system adapts the body to the environment through internal selection of a large population of diverse antibodies. During human development, neonates generate a large diversity of random movements which are narrowed on the basis of punishment/reward learning, a child can sit by 6 months, crawl by 12 months and walk by 18 months. The brain with its vast interconnectivity can generate a vast number of potential firing patterns representing a vast range of behaviours but only those patterns which are useful are selected by a "value system" (emotional/motivational) structures in the brain stem which project and ascend diffusely over the entire cerebral cortex, eg. the reticular formation. Certainly, the brain stem reticular formation has great proliferation of dendrites and projections into cerebral cortex characterised by highly anisotropic connectivity. This value system evolved to recognise events in the world that are of adaptive value. Firing patterns which generate useful behaviours in the world are strengthened by signals from the discriminatory value system. Even more abstract levels of behaviour and thought are based fundamentally on the value system. It is not currently clear how this may be incorporated into the design of space systems but robots have been reported incorporating aspects of this approach – eg. Darwin III. Darwin IV has been developed at the Neurosciences Institute based on these noyion as a "neurally organised mobile adaptive device (NOMAD)". It generates spontaneous motion but when its value system detects light it selects and strengthens those synaptic activities that generate the desired movement. From spontaneous behaviour, light tracking behaviour emerges.

De Garis (1991) used GA's to evolve neural network modules (GenNets) to construct hierarchical control systems. Each artificially evolved neural net is regarded as a functional black box which performs a specific behaviour, i.e. a specific control system. Fitness is a measure of the quality of the behaviour. Each neural net module input is connected to each sensor and each output unit is connected to each motor. Synaptic weights and thresholds can be coded and evolved as genetic algorithm chromosome values. The GA has better convergence properties than gradient-based techniques such as backpropagation in the face of noise, reducing training times by two orders of magnitude. The GA searches the solution space globally while gradient techniques perform local searches in the vicinity of the current solution. It finds the best combination of connection weights rather than just each weight independently. A population of genetic algorithms has a proportion (typically 20%) which survive through cross-over into the next generation undergoing 10% random mutation.

A typical fitness criterion for a wheeled rover is given by:

$$F = V + (1 - \sqrt{DV}) + (1 - I)$$

where:

V=average rotation speed of opposing wheels
DV=difference between signed speed values of opposing wheels
I=activation value of the sensor with the highest input

This maximises speed, movement in a straight line and avoidance of obstacles. The fitness function should compute information that is available to robot through its sensors only. The N neurons of the neural net are labelled 1 to N (N~10-15 being typical) and N^2 interconnections if the net is fully connected. A number of binary places P are allocated to specify the weight values and one bit to specify its sign. For N^2 connections, the binary string requires $N^2(P+1)$ bits to specify all the connection weights. It is possible therefore to specify the GenNet from its chromosome representation. GenNet modules may be constructed into hierarchical structures to exhibit arbitrarily complex behaviour. Each GenNet with its specific behaviour may be evolved separately and then organised appropriately. Control GenNets are used to direct and command functional GenNets in that the outputs of the control nets serve as inputs to the functional nets. Control GenNets are evolved once the functional GenNets have been frozen. Such hierarchical control may be multiply layered. This methodology has been applied to an artificial lizard, LIZZY which reacts to three types of creatures each of which emits a characteristic frequency: mate, predator and prey. Its behaviour varied according to the signal it received and its strength (indicating its distance) via its a suite of detector, logic and effector GenNets. GenNet weights were stored in ROM and the weights were used to connect up and control physical detectors and effectors.

Handley (1993) introduced Genetic Planner which used GA-based artificial selection to breed computer programs that generate plans to purposively control a mobile robot. The genetic planner used crossover recombination and artificial selection through a fitness measure to breed computer programs. A world model comprised a set of procedural operators which operated on the set of predicates that defined the world model. The world model predicates were of the form: at (object,x,y), with-respect-to(object1,object2). The procedural operators included: “turn left/right 90”, “move forward a fixed distance”, etc. The fitness function was defined as a distance between objects measure. The agent used the procedural operators to generate robot action plans utilising a goal-oriented fitness measure and a world simulation. The genetic planner used the operators to execute candidate computer program plans and the fitness function evaluated the state of the world after each simulation. As the genetic planner did not reason about the world, it randomly generated plans and ran them until it came up with effective plans.

An extension of these approaches which offers great interest is incorporating the process of development of neural networks from genetic coding. This relies on encoding of the mapping between the genetic coding and the expression of that coding in the structure of the neural network. This emulates the process of embryonic development in providing information amplification from the genetic code to the connectivity matrix for a neural network. Initial work appears to suggest that the more complex the coding mechanism, the more difficult the evolutionary convergence. Husbands et al (1995) stressed the need for adequate genetic coding of phenotypic behaviour. In the simplest case, the genetic algorithm codes a direct description of the circuit network wiring. In biological systems, the genome codes a more complex interpretative process through development into a phenotype. A number of possibilities for sensorimotor control are applicable, particularly using genomes to code for neural network topologies. The genotype may encode a graph generation grammar with each node of the tree being specific developmental genograms, such as DIVIDE, GROW, etc where the rewriting process is considered a developmental process. Rewrite operations include: PAR which divides the cell into two cells, CUP which cuts links, WAIT which delays rewrite operations. NCAGE (Natework Control Architecture Genetic Encoding) is a programming language that permits the specification of sensorimotor controllers based on neural networks. It utilises a three-letter alphabet A={1,0,STOP}. Genomes are NCAGE programs which are compiled to create the controller. The genome is a sequence of subroutine calls. Cells place neurons into different positions. Mutation, crossover, and inversion operators act on the genome. Alternatively, parameters to systems of ordinary differential equations, may be encoded into a genotype. Neurons are positional in a 2D plane and exert attraction forces on dendrites from other neurons. The dendrites grow in the force field generated by the distribution of neurons. If a dendrite moves within a given distance of a neuron, it establishes a connection to it. The motion of the dendrite is modelled as the vector sum of the attractive forces on the dendrite (dependent on an inverse square law of distance of the end of the dendrite to the centre of the neuron), a viscous resistive force and bias force at the end of the dendrite inversely proportional to the cube of the current length of the dendrite.

$$\frac{d^2 r_{ij}}{dt^2} = \sum_{k=0}^{n} F_{ijk} - K\frac{dr_{ij}}{dt} + \frac{G_{ij}}{l_{ij}^3}$$

where :

$F_{ijk} = \frac{S_{ik}}{r_{ijk}^2}$, *$K$, G_{ij}, S_{ik} are genetically determined constants.*

A genotype encodes the products of the equation of motion – a bit string representing neuronal parameters.

The earliest work on developmental approaches to evolutionary neural networks was due to Kitano [1990]. He reduced the size of the genome using a graph generation grammar to encode neural network topologies which he demonstrated to have a better scaling behaviour than direct encoding schemes. Belew [1993] also used a grammar to simulate developmental processes. His scheme is context sensitive -- but it is restricted to pre-specified neural network topologies. Perhaps the most successful scheme to date is due to Graua [1994]. He encoded the developmental process as a grammar tree controlling a cellular developmental process. In this approach the process starts with a single cell that divides, producing daughter cells that can in turn further divide. How much division occurs and how the cells connect together at each stage is controlled by the genotype. Cells inherit their connections from their 'parents' and no context sensitive development is possible. This method allows a restricted form of specification and reuse of 'modules' and has been successfully applied to fairly complex evolutionary robotics tasks, such as the development of modular and non-modular controllers for hexapod and octopod walking robots. Nolfi (1994) used a developmental model for neural networks based on the growth of inter-neuron connections. These controllers were composed of a collection of artificial neurons distributed over a 2-dimensional space with growing and branching axons (Figure 133, top). Inherited genetic material specified instructions that controlled the axonal growth and the branching process of neurons. During the growth process, when a growing axonal branch of a particular neuron reached another neuron a connection between the two neurons is established. At the bottom of Figure 133 the network resulting from the growth process displayed in the top of the figure can be seen. This is the resultant network after the elimination of isolated and non-functional neurons. Axons grew and brunched only if the activation variability of the corresponding neurons was larger than a genetically specified threshold. This simple mechanism is based on the idea that sensory information coming from the environment has a critical role in the development of the connectivity, and other properties, of the biological nervous system and, more specifically, that the maturation process is sensitive to the activity of single neurons . Therefore the developmental process was influenced both by genetic and environmental factors (i.e. the history of sensory states experienced by the network influenced the process of neural growth). A potential problem of this approach is that the number of genes in the genome grows with the number of neurons to be encoded which leads to relatively poor scaling behaviour.

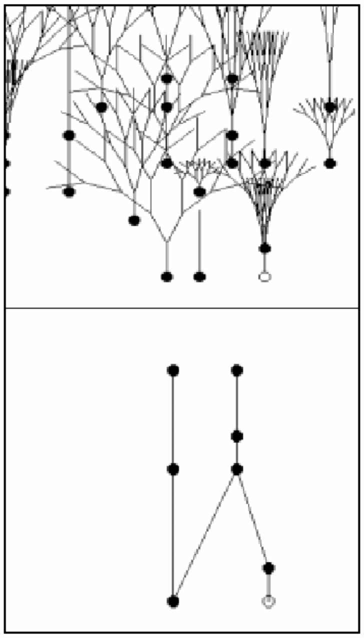

Figure 133 – Development of an evolved neural network using the Nolfi scheme. Top: the growing network. Bottom: the resultant network after pruning of isolated non-functional axons and neurons.

Michel and Biondi (1995) introduced a model for development which uses morphogenetic mechanisms to evolve neural control structures for autonomous agents. Dellaert and Beer (1994) proposed a model based on Boolean networks, representing gene regulation networks, to evolve autonomous agents with developmental processes. They were able to demonstrate the evolution of simple behaviours with this scheme. More recently Eggenberg (1997) proposed a biological inspired model to develop the structure of artificial neural networks. The model is based on an artificial genetic regulatory system which controls the development of the neural network. The model allows for different cell types which are the result of different intercellular communication processes. Having different cells was shown to lead to differentiated development of connection patterns.

3.11.2. Evolutionary Robotics

As outlined earlier, evolutionary robotics is concerned with the use of evolutionary algorithms to develop key aspects of autonomous robots including their control system and/or body and sensor designs. To this end the key elements in the approach are:

- An artificial genetic encoding specifying a robot control systems/body plan/sensor properties etc along with a mapping to the target system
- A method for measuring the fitness of the robot behaviours generated from these genotypes
- Methods based on the application of selection and 'genetic' operators to produce the next generation from the current

Potential advantages of this methodology include:

- It is possible to explore potentially unconstrained designs that have large numbers of free variables.
- It is possible to use the methodology to fine tune parameters of an already successful design.
- It is possible, through the careful design of fitness criteria and selection techniques, to take into account multiple, and potentially conflicting, design criteria and constraints (e.g. efficiency, cost, weight, power consumption, etc).
- It is possible to develop highly unconventional and minimal designs.
- It is possible to explicitly take into account robustness and reliability as major driving force behind the fitness measure, factors that are particularly important for space applications.

An early example of work in this field is described in [Harvey et al 1993]. In this research, recurrent dynamic neural networks are developed using incremental evolutionary learning algorithms to provide the basis for visual navigation behaviours. Such tasks were chosen because extroreceptive visual and proximal data are essential for adequate navigation beyond primitive behaviours. The initial behaviour required was for a robot, using only its visual sensors, to manoeuvre itself into the centre of an empty room from a random starting position and orientation. A range of control architectures and sensor layouts are represented as a population of genotypic chromosomes. These chromosomes interbreed with mutations according to a task-oriented fitness function which imposes a selection pressure on the population. Each member of the population is evaluated through the fitness function over a number of successive trials. In order to do this in simulation, noise must be incorporated into the model to ensure that hardware sensors and actuators are accommodated. In order to ensure generalized behaviour, a range of unstructured, dynamic environments should be used in the simulation. It was argued that genetic programs which allow recursion are not a suitable candidate for evolved control systems as the program may never halt unless an arbitrary time-out signal is imposed. Genetic programs are also “coarse” which limit the genetic exploration process across a coarse-grained fitness landscape with limited capability of accommodating noise. Such computational approaches do not exploit the fact that physical devices such as autonomous robots are in fact dynamic systems. Neural networks offer an appropriate medium for a cognitive architecture in that they can model any non-linear input-output function which may be searched using a genetic algorithm. Traditional feed-forward neural networks have no internal state and can only model reactive behaviours. It was proposed that suitable neural networks suitable for generating more complex behaviours should have internal feedback and be characterised by physical constraints such as physical signal propagation delays through analogue circuits with real-valued signals through the connections – this defines a recurrent neural network which can model any class of dynamic system such as networks of augmented finite state machines. In this work, each neuronal unit is modelled as a noisy linear threshold device. Each sensor is associated with a single input node and two output units are associated with each motor. Signals to the motors may be represented as a real value in the range $-1.0<s<1.0$ which is divided into 5 equal segments depending on which segment the signal falls into – stationary wheel, full speed forward, full speed backwards, half-speed forward, half-speed backwards. Each network also possesses an arbitrary and unspecified number of hidden units. The genetic encoding specifies the properties of the units and the connections and connection types between them. Each genotype is interpreted sequentially – first inputs are coded, then internal and output units are coded, each with a marker for addressing purposes. Each node is encoded by a representation of its properties such as threshold, then a number of groups to represent connections from that unit (excitatory or inhibitory), and a target unit addresses either relatively or absolutely. Cross-over is achieved through the SAGA (Species Adaptation Genetic Algorithm) principle which allows variable length cross-over between homologous genotype sections by minimising length changes [Harvey 1997]. It is only through the use of increasing genotype lengths that the genotypic space can itself evolve beyond finite structures, i.e. a form of species evolution which allows local searches whilst retaining existing adaptations. Mutation rates must be optimised to minimise deviation from fitness maxima while maximising exploration across high ridges, i.e. around 1 mutation per genotype per generation. Selection should be constant, i.e. expected number of offspring of an individual depends on its fitness ranking rather than ratio of its fitness to average fitness. The control systems for successful evolved robots were analysed in some detail [Husbands et al 1997] and shown to be in several respects optimal. This early work at Sussex University showed the feasibility of the approach by developing working controllers for a number of simple robots engaged in simple navigation tasks. Important, and influential, aspect of this early work were the concurrent evolution of control networks and sensor layout (positions and properties of the simple vision sensors used), rather than imposing a controller on fixed sensor morphologies and vice versa, and the use of a dynamic length encoding to evolved unconstrained control architectures – the size and connectivity patterns were in the complete control of the evolutionary algorithm.. At the same time as methods were being developed at Sussex, parallel approaches were being developed at EPFL, Lausanne, and CRC, Rome as well as related techniques at Case Western University, USA. Theoretical frameworks were also developed at MIT and other institutions.

3.11.2.1. Simple mobile robots

The Khepera robot [Mondada et al. 1993] was used widely in the early 1990s in various strands of evolutionary robotics research. The Khepera, originally designed and constructed at EPFL, Switzerland, is a circular miniature mobile robot with diameter of 55mm, a height of 30mm, and a weight of 70g that is supported by two wheels and two small Teflon balls. In its basic configuration, it is equipped with eight infra-red proximity sensors – six on the front, two on the back – that may also act as visible-light detectors. The wheels are controlled by two DC motors with incremental encoders that move in both directions.

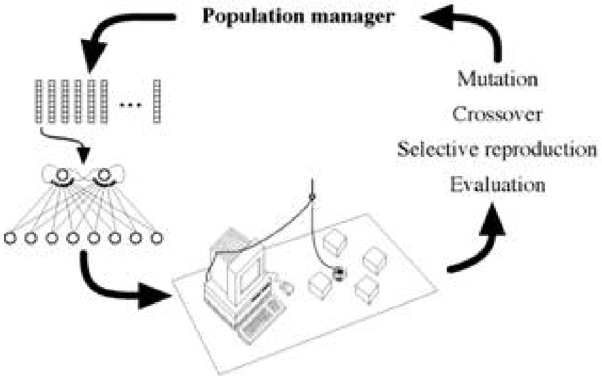

Figure 134 – Evolutionary robotics experimental setup using a Khepera robot as performed at EPFL.

The first successful evolutionary robots experiments at EPFL employed a Khepera robot in the setup illustrated in Figure 134 and Figure 135. A population of bit strings encoded the connection weights and node thresholds for a fixed architecture feed-forward neural network. Each member of the population was decoded into a particular instantiation of a neural network controlled which was then downloaded onto the Khepera robot. This controlled the robot for a period of time as it moved around the environment shown in Figure 135.

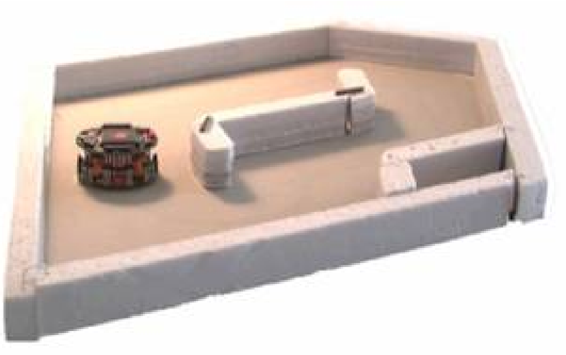

Figure 135 – The obstacle avoidance course used with the evolutionary robotics experimental setup illustrated in Figure 134.

The following fitness function was used:

$$F = V + (1 - \sqrt{DV}) + (1 - I)$$

where:

V=average rotation speed of opposing wheels
DV=difference between signed speed values of opposing wheels
I=activation value of the IR sensor with the highest input

This maximises speed, movement in a straight line and avoidance of the walls of the environment. After about 24-36 hours of real-world evolution using this setup, controllers were evolved that successfully generated efficient motion around the course avoiding collisions with the walls.

Using a specially constructed software simulator, Jakobi et al. (1995) evolved both obstacle-avoidance and light-seeking behaviours in Khepera. The simulation was based on a continuous two-dimensional model of the real world physics and allowed the calculation of the dynamics of the robot's sensory inputs in response to its motor signals. Recurrent networks of threshold units that were evolved in simulation evoked qualitatively similar behaviour on the real robot, especially when the levels of noise present in the simulation had similar amplitudes to those observed in reality. This work showed that simulators could be used in evolutionary robotics, although many had doubted this prior to this work [Brooks 1992]. Simulators are now widely used in real-world evolutionary robotics applications. Simulation has the potential advantage of speed as well as controllability and ease of access to data for analysis.

Unless special generalized noisy simulations are used, as in the piece of work described directly above, there will usually be a mismatch between simulated and real sensory data. Simulations often not describe the physical laws of interaction in the real world and physical sensors and actuators are subject to uncertainty. However, an alternative approach that has proved successful [Nolfi et al. 1994b] is to evolve first in the imperfect simulation and then continue evolution in the real world. After initial evolution in the simulated environment, only a few additional generations are usually required to achieve successful behaviour in the real world – a kind of adaptive fine-tuning. It is required that the level of noise in the simulation matches that of the real world. Most of the evolutionary process is performed on the simulator with a further few generations of evolution on the robot itself. The use of simulation significantly speeds up the learning process. Typical tasks included obstacle avoidance and goal-seeking. Grasping behaviour also has been achieved using a 5-input/4-output neural network. The inputs were from two frontal sensors, left and right sensors, and a gripper sensor. The outputs included two wheel velocities, pick up procedure and release procedure. The fitness function included parameters concerned with robot-target distance, target position relative to the robot, presence of an object in the gripper and its release, etc. Learning gripping behaviour was more complex than most other rover-type behaviours as it involved a multitude of sensors and tasks including moving the robot close enough to the target to be grasped. The fitness function became complex and subgoals were embedded in them to produce the required behaviours. The species adaptation algorithm (SAGA) may overcome this problem by utilising variable length genomes [Husbands et al. 1997]. Initially, short simple genotypes may be used to generate elementary behaviours, and the length increased until sufficiently complex behaviours are evolved. A specialised cross-over procedure is used with these variable length genomes. Further work has involved the evolution of visually guided behaviours on GA-based recurrent neural network control systems incorporating excitatory and inhibitory mapping between vision and actuators [Husbands et al. 1997b]. A number of circular receptive fields of given radius and angular position rϕ are represented on the genome which provides the basis for interpreting visual inputs through the construction of vector flow fields representing robot motion. Incremental evolution whereby increasingly challenging tasks are solved sequentially rather than always starting from random populations offers better performance.

To evolve the capacity of moving in the environment while avoiding obstacles, Miglino et al. [Miglino et al. 1995a; Lund and Miglino 1996] used a two-layer feed-forward neural network with no hidden units and a fitness function with three components, which were respectively maximized by speed, by straight direction, and by obstacle avoidance. With the help of a genetic algorithm, the synaptic connections and thresholds of the neural controllers were first evolved through simulation. Then, the corresponding networks were downloaded onto a Khepera and proved to be efficient. With an alternative approach, Floreano and Mondada (1994, 1996) had the whole evolutionary process take place entirely on the robot without human intervention, and two-layer Elman neural networks [Elman 1990] were used as controllers. In the work of Nolfi (1996b), the parameters of a feed-forward neural network with no hidden units were evolved to control a Khepera robot that had to explore his environment, to avoid walls and to remain close to a cylindrical target when it found it. Experimental results showed that the evolved individuals were successful in the real world and that, by intensively using an active perception strategy, they could overcome the problem posed by the fact that the walls and the target were hard to distinguish in most cases. Several research efforts have aimed at evolving neural controllers for the Khepera robot through developmental approaches that call upon various biomimetic processes – like cell division, cell differentiation, or cell adhesion - to gradually build a neural control architecture. Controllers for obstacle-avoidance, wall-following, light-seeking or light-avoiding behaviours have thus been evolved by Eggenberger (1996) and Michel [Michel 1996; Michel and Collard 1996].

Interactions between reinforcement learning and evolution have been exploited in the work by Grefenstette and Schultz (1994), which calls upon the use of the SAMUEL classifier system [Grefenstette and Cobb 1991] for evolving collision-free navigation in a Nomad 200 mobile robot. The task consisted of learning to reach a fixed goal location in a predetermined time, starting from a fixed initial position within an environment that contained obstacles whose positions were changed at each trial. Yamauchi & Beer (1994) used a Nomad 200 mobile robot to evolve highly dynamical recurrent neural controllers capable of identifying one of two landmarks. The robot's trajectory was controlled by a fixed behaviour-based control system that allowed the robot to find a wall and follow it. Network parameters - like time constants, thresholds, or connection weights - were genetically encoded as vectors of real numbers, of which each element was indivisible under crossover. The fitness function of each individual in a population of 100 networks was evaluated in simulation and assessed the average capacity of the network to correctly identify the landmarks over six test trials. After 15 generations, an individual capable of correctly recognizing the landmarks in simulation was generated. When transferred on the real robot, it correctly classified the landmarks in 17 out of 20 test trials. Beer has since gone on to show how such dynamical networks can generate fairly sophisticated behaviour such as sequence learning without having to explicitly bolt on 'learning' procedures and build in domain knowledge.

3.11.2.2. Visually Guided Behaviours

Several pioneering experiments on evolving visually guided behaviours have been performed [Cliff et al. 1993; Harvey et al. 1994; Jakobi 1997a,b; Husbands et al. 1997, 1998] in which discrete-time dynamical recurrent neural networks and visual sampling morphologies are concurrently evolved to allow a gantry robot (as well as other more standard mobile robots) to perform various visually guided tasks. The Sussex gantry robot is shown in Figure 136.

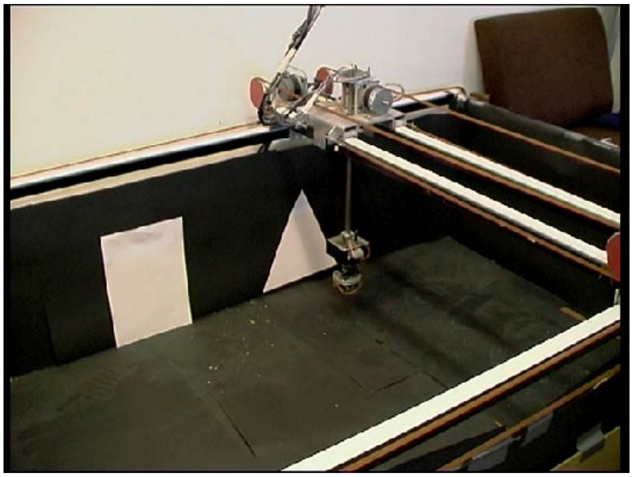

Figure 136 – The Sussex gantry robot.

Figure 137 – The camera head for the gantry robot. The black box at the top of the structure houses a ccd camera. This points down on a rotatable mirror angles at 45 degrees. The plastic disc at the bottom is connected to a multidirectional sensor and hence acts as a contact sensor.

A ccd camera points down towards a mirror angles at 45 degrees as shown in Figure 137. The mirror can rotate around an axis perpendicular to the camera's image plane. The camera is suspended from the gantry allowing motion in the X, Y and Z dimensions. This effectively provides an equivalent to a wheeled robot with a forward facing camera when only the X and Y dimensions of translation are used. The additional dimension allows flying behaviours to be studied.

The apparatus was initially used in manner similar to the real-world EPFL evolutionary robots setup illustrated in Figure 133. A population of strings encoding robot controllers and visual sensing morphologies are stored on a computer to be downloaded one at a time onto the robot. The exact position and orientation of the camera head can be accurately tracked and used in the fitness evaluations. A number of visually guided navigation behaviours have been successfully achieved including navigating around obstacles and discriminating between different objects. In the experiment illustrated in Figure 136 starting from a random position and orientation the robot has to move to the triangle rather than the rectangle. This has to be achieved irrespective of the relative positions of the shapes and under very noisy lighting conditions. Recurrent neural network controllers were evolved in conjunction with visual sampling morphologies as illustrated in Figure 138. Only genetically specified patches from the camera image are used (by being fed into the control network by connecting to input neurons according to the genetic specification). The rest of the image is thrown away. This resulted in extremely minimal systems only using 2 or 3 pixels of visual information, yet still able to very robustly perform the task under highly variable lighting conditions.

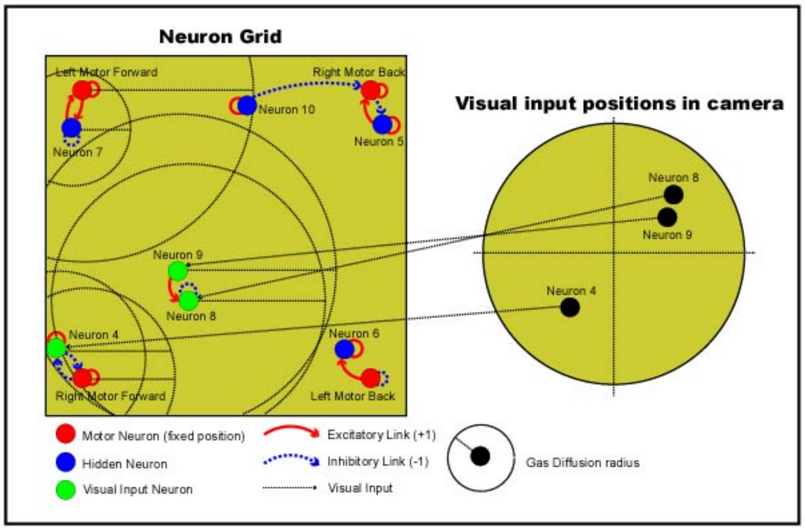

Figure 138 – Concurrent evolution of control network and visual sampling morphology for the gantry robot. See text for further details.

More complex visual behaviours took several days to evolve on this setup which made it very difficult to perform detailed comparisons and statistical analysis of the pros and cons of different approaches and neural architectures. To counter this problem Jakobi developed a new approach to simulations. His minimal simulations [Jakobi 1998b] are ultra-fast by using crude models of carefully chosen sets of basic interactions between the robot and its environment. These employ multi-layered structured noise on all modelled aspects. This means there are no stable simulation 'features' that an evolved robot could come to rely on (and therefore fail when transferred to the real world). Highly robust generalized behaviours are forced by this technique as they are the only ones that can score reliably. One way of looking at it is that the task is made more difficult in the simulator by being over-specified in terms of robustness to noise. This style of simulation has now be validated on a range of different robots. In each case the robots evolved in simulation crossed the reality gap – they transferred perfectly to the real world. The speed with which such simulations run mean that detailed comparative studies have now been undertaken which has allowed a deeper understanding of some of the issues underlying successful evolution.

More recent work has focused on highly evolvable (i.e. take very few generations to achieve fit behaviour) minimal neural networks based strongly on current research in neuroscience. In particular the concept of a diffusing gaseous modulatory neurotransmitter has been abstracted into a class of artificial neural networks known as GasNets [Husbands et al. 2001]. Such networks have been shown to be highly reliable and robust to noise and damage [Smith et al. 2003].

3.11.2.3. GasNets: From Neuroscience to Engineering

This section describes one style of artificial neural network from a class of networks whose operation is strongly inspired by those parts of contemporary neuroscience that emphasize the complex electrochemical nature of real nervous systems. So-called GasNets incorporate virtual diffusing gaseous neuromodulators and are used as artificial nervous systems for mobile autonomous robots. They are being investigated as potentially useful engineering tools and as a way of gaining helpful insights into biological systems.

The basic GasNet networks used in many recent experiments are discrete time step dynamical systems built from units connected together by links that can be excitatory (with a weight of +1) or inhibitory (with a weight of -1). The output, O_i^t, of node *i* at time step *t* is a function of the sum of its inputs, as described by Equation 3.10.1. In addition to this underlying network in which positive and negative `signals' flow between units, an abstract process loosely analogous to the diffusion of gaseous modulators is at play. Some units can emit virtual 'gases' which diffuse and are capable of modulating the behaviour of other units by changing their transfer functions in ways described in detail later. This form of modulation allows a kind of plasticity in the network in which the intrinsic properties of units are changing as the network operates. The networks function in a 2D plane; their geometric layout is a crucial element in the way in which the 'gases' diffuse and affect the properties of network nodes.. This aspect of the networks is described in more detail later.

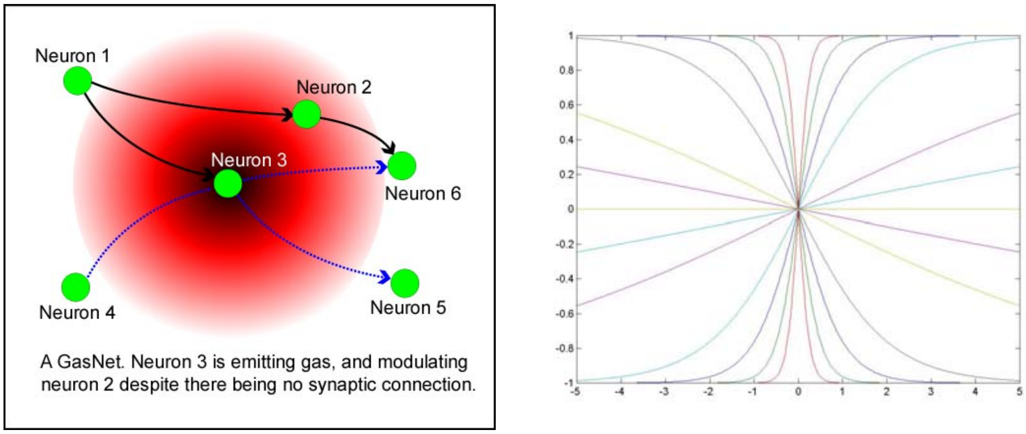

Figure 139 – GasNet operation depends on the geometric layout of the nodes in a 2D plane. The right side of the diagram shows how the shape of the tanh transfer function depends on the gain parameter kit.

$$O_i^t = \tanh[k_i^t(\sum_{j\in C_i} w_{ji}O_j^{t-1} + I_i^t) + b_i]$$

Where C_i is the set of nodes connected to node *i*, I_i^t is the external (sensory) input to node *i* and b_i is a genetically set bias. Each node has a genetically set default transfer function gain k_i^0. The right hand side of Figure 139 (right) shows the shape of the function *tanh(kx)* over the range [-5,5] for a discrete set of values of *k* between -4 and 4. It is this gain parameter that is modulated by the diffusing virtual gases in the networks. This means that while the gases are active the shapes of the node transfer functions are being altered from time step to time step. The mechanism for this is explained in the next section.

3.11.2.3.1. Diffusion and modulation

The virtual diffusion process is simple in order to be computationally fast so that GasNets can be used to control robots in real time. For mathematical convenience there are two gases, one whose modulatory effect is to increase the transfer function gain parameter and one whose modulatory effect is to decrease it. It is genetically determined whether or not any given node will emit one of two 'gases' (gas 1 and gas 2), and under what circumstances emission will occur (either when the 'electrical' activation of the node exceeds a threshold, or the concentration of one of the gases (genetically determined) in the vicinity of the node exceeds a threshold). The electrical threshold used in the experiments described later was 0.5, the gas concentration threshold 0.1. Allowing these two highly biologically inspired possibilities is important – it provides a mechanism for rich interaction between two processes, the 'electrical' and the 'chemical'. A very abstract model of gas diffusion is used. For an emitting node, the concentration of gas at distance d from the node is given in the equation below. Here, r_i is the genetically determined radius of influence of the ith node, so that concentration falls to zero for $d>r_i$. This is loosely analogous to the length constant of the natural diffusion of NO, related to its rate of decay through chemical interaction. $T_i(t)$ is a linear function that models the build up and decay of concentration after the node has started/stopped emitting. The slope of this function is individually genetically determined for each emitting node, C_0 is a global constant [Husbands et al. 1998].

$$C_i(d,t) = C_0 e^{-2d/r_i} \times T_i(t)$$

At each time step the gain parameter, k_i^t , for the node transfer function at each node, is changed (or *modulated*) by the presence of gases at the site of the node. Gas 1 increases the value of k_i^t in a concentration dependent way, while gas 2 decreases its value. Concentration contributions from nodes within range of any given site are simply added together. The modulatory effects of the two gases are then summed to calculate the value of k_i^t at each time step. Each node has its own default rest value for the gain parameter, the virtual gasses continually increase or decrease this value. Referring to the bottom part of Figure 139, this modulation can potentially have drastic effects on a nodes's transfer function, dramatically increasing or decreasing, or even flipping the sign of, its slope. This means that the networks are usually in flux, with rich dynamical possibilities.

Since there were no pre-existing principles for the exact operation and design of such networks, it was decided to allow most of their detailed properties to be genetically specified, giving the possibility of highly non-uniform dynamically complex networks. Hence, in most experiments to date nearly everything is up for grabs: the number of nodes in a network; the way they are connected; the position of the nodes on the 2D plane; the individual properties of each node controlling when (if at all) they emit a gas; which gas is emitted and how strongly; how and if nodes are connected to sensors or motors, as well as various properties of the sensors and motors themselves. About 20 variables per node are needed to describe all this. It has been shown that a well setup evolutionary search algorithm is a good tool for exploring the space of such systems, looking for interesting and useful examples that deepen our understanding of autonomous adaptive systems or provide practical engineering advantages such as robustness and reliability. The next section gives an example of using GasNets in such an evolutionary robotics setting.

3.11.2.3.2. Experimental Comparison

Various forms of GasNet have been used as robot controllers for a variety of tasks and robots [Husbands 1998; Husbands et al. 1998]. A very large number of runs of one particular experimental setup have been carried out, giving us enough data to be able to make statistically significant claims. In this series of experiments GasNets were evolved to control a robot engaged in a visually guided behaviour involving shape discrimination (the gantry robot described earlier was used). A simple robot with a fixed CCD camera as its main sensor moved in an arena as illustrated in Figure 136. Two light coloured shapes, a rectangle and a triangle, were placed against a darker background on one of the walls. The task was to reliably move to the triangle, while ignoring the rectangle, from a random initial orientation and position under highly variable lighting conditions. The relative positioning of the shapes, in terms of which was on the left and which on the right, was made randomly.

As well as network size and topology, and all the parameters controlling virtual gas diffusion and modulation, the robot visual morphology, i.e. the way in which the camera image was sampled, was also under unconstrained genetic control. This was achieved by allowing the evolutionary search algorithm to specify the number and position of *single* pixels from the camera image to use as visual inputs. The grey scale intensity value of these pixels (normalised into the range [0.0,1.0]) were fed into the network, one for each genetically specified visual input node in the net. This is illustrated in the bottom left quadrant of Figure 141. Note that this means that the evolved control systems were operating with extremely minimal vision systems, just a few single pixel values. Given the very noisy lighting conditions and the minimal visual input, the shape discrimination task becomes non-trivial. The experimental setup is shown in Figure 136. A gantry robot provided a highly controllable experimental apparatus. A ccd camera, facing down onto an angled rotatable mirror, is suspended from a gantry arm. The combined movements of the gantry and mirror make this arrangement equivalent to a standard 2 wheeled robotic platform with a camera mounted on top, such as the one shown in Figure 140. The gantry robot can be run with off-board power and computing, making long runs and data collection very easy.

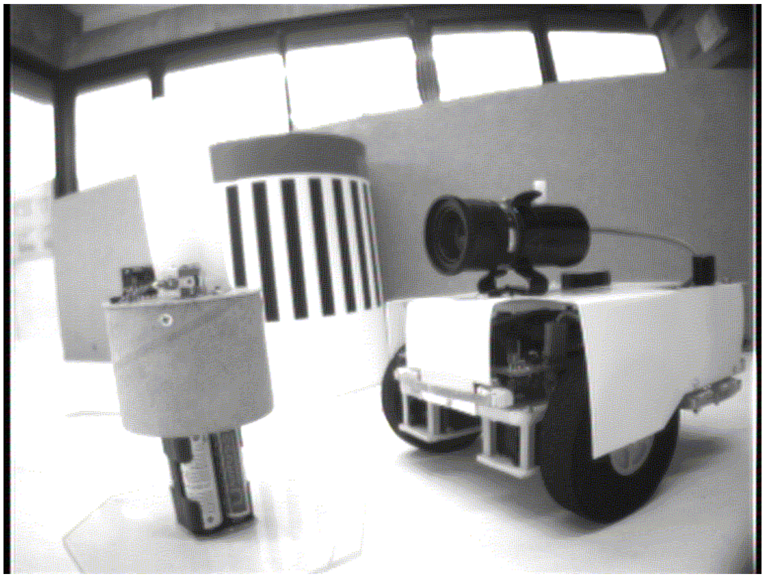

Figure 140 – A standard differential drive 2 wheeled robot with a ccd camera mounted on top.

All the evolutionary runs were carried out using a Jakobi minimal simulation of the robotic setup. The methodology behind these ultra-lean ultra-fast simulations was developed by Jakobi [1998b] to address one potential problem with evolutionary approaches to exploring classes of robotic control systems: the time taken to evaluate behaviours over many generations. Through a careful use of noise and important decisions about what not to model, a minimal simulation will run very fast but behaviours evolved in them will transfer to the real robots. For full details of the minimal simulation used for the triangle rectangle task see [Jakobi 1998b]. In the experiment described here, all successful evolved controllers crossed the reality gap: they generated the same behaviours on the real robot as in simulation. Success was defined as being able to move to the triangle and stay there 30 times in direct succession from random starting positions and orientations, under very noisy lighting conditions and irrespective of the relative positioning of the shapes on the same wall. The great advantage of using minimal simulations in this work is that we were able to perform many complete evolutionary runs and hence derive meaningful statistics.

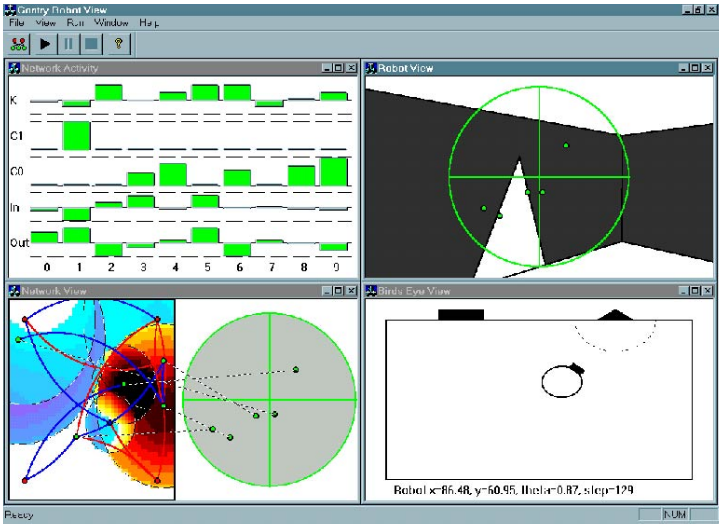

Robot x=86.48, y=60.95, theta=0.87, step=129

Figure 141 – The visualisation tool used with the minimal simulation of the shape discrimination task. The top right quadrant shows the view through the robot's camera, the bottom right gives a bird's eye view of the robot moving in the arena. The left-hand side of the screen illustrates the structure (including visual morphology) and functioning of the GasNet controlling the evolved robot. The shading in the network representation at extreme bottom left shows the gas concentrations in the network plane at the instant the snapshot was taken. The darker the shading the higher the concentration.

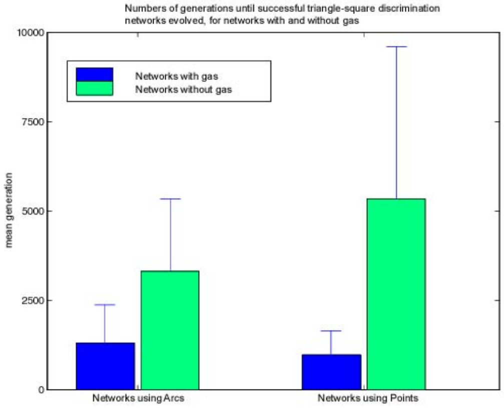

Figure 142 – The average number of generations needed to find controllers giving perfectly successful behaviour on the triangle rectangle problem. The dark columns are for networks with the gas mechanism turned on. The light columns are for networks with the gas mechanism

turned off. The figure illustrates two sets of 20 runs in each condition; the difference between the left and right-hand sets is the way in which the network connectivity was encoded.

The initial set of evolutionary GasNet experiments with this task resulted in highly robust controllers emerging about 10 times faster than in earlier runs with conventional connectionist networks [Husbands et al 1998]. Subsequent comparative runs have concentrated on identifying whether or not the virtual gas modulation mechanism was at the root of this speed up. The key result is illustrated in Figure 142. In all experiments the genotypes were strings of integers encoding the various properties of the controllers and coupled visual morphologies, a geographically distributed genetic algorithm was used with a population of 100 , a number of different mutation operators were used in tandem, including node addition and deletion operators. The fitness function was based on a weighted average of final distances to the triangle over a set of evaluations from different initial conditions and different relative positioning of the shapes. Poor scores were weighted more heavily than good scores, encouraging robustness by requiring uniformly high scores across the whole evaluation set.

It can clearly be seen in Figure 142 that controllers based on networks with the virtual gas diffusion and modulation mechanisms turned on evolve significantly faster than those that are identical in every respect (including genotypes and all the evolutionary machinery) except that the gas mechanisms are rendered inoperative. This result has been repeated under various different encoding schemes and for a wide range of mutation rates . The clear implication is that GasNets are more evolvable – their search space is more amenable to the form of evolutionary search used – than the various other forms of network explored. Obviously this could be a potentially very useful property and it is looked at in more detail in the next section.

Nearly all the successful GasNet controllers that were examined in detail exhibited surprisingly simple structures (a typical example is shown in Figure 139) relying on a very small number of visual inputs, although their internal dynamics, supported by interwoven 'chemical' and 'electrical' processes, were often intricate [Husbands et al. 1998]. A number of interesting sub-networks, such as oscillators making use of spatial aspects of the modulation and diffusion processes, were independently evolved in several runs, suggesting that they are easily found building blocks that the evolutionary process can make good use of.

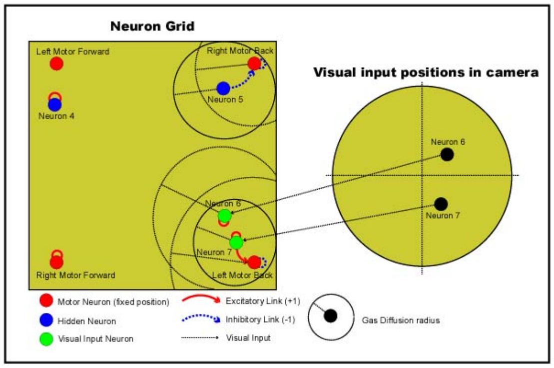

Figure 143 – A typical evolved GasNet controller for the triangle rectangle task illustrating the kind of structural simplicity often found in highly robust solutions.

Very recent work has seen the development of more sophisticated GasNets that incorporate further biological features and are even more evolvable (about a further order of magnitude faster). These have been used to successfully evolve more complex visual behaviours such as distinguishing between classes of objects and using context sensitive visual cues in navigation. Detailed theoretical work has also shed significant light on the question of why this form of network is more evolvable [Philippides et al. 2004].

Baluja (1996) presented an evolutionary method for designing a video-recognition neural controller for the NAVLAB autonomous land vehicle. The PBIL (Population-Based Incremental Learning) evolutionary algorithm was used which specifies the probabilities of having a 1 or a 0 in each bit position of a string encoding the topology and connection weights of a neural controller. This evolutionary approach performed better, on average, than standard backpropagation.

The next few sections review work on visually guided evolutionary robotics carried out at EPFL, Laussane, in the lab run by Prof. Dario Floreano.

3.11.3. Active Vision and Feature Selection at EPFL

In these experiments active vision and feature selection are co-evolved, with behaviour affecting the type of visual features gathered and the visual features affecting behaviour. Active vision is inspired by the way that mammals and insects interact visually with their environments and a motivation underlying this research is that in biological systems both sensorimotor and behavioural characteristics *constrain* the type of visual features that an organism responds to. The experiments use primitive retinal systems and simple ANNs, which requires reducing the computational complexity of visual processing by *actively selecting* parts of the visual scene for processing, thereby lessening the information load on the system by focussing only on those parts of the visual environment relevant to the current task. Feature selection involves filtering an image to enhance relevant features and remove irrelevant ones. For example, convolving an image with a Difference of Gaussians (DoG) operator can enhance edges. In biological systems the receptive fields of cells respond only to some properties of the environment; for example, centre-surround antagonist neural organizations are found in the periphery of both vertebrate and invertebrate vision systems and they only respond to edges of particular orientation.

3.11.3.1. Artificial Retina

The Artificial Retina experiments demonstrate that active vision can exploit simple features to perform complex shape discrimination. Furthermore, this discrimination is very difficult for similar systems that do not have active vision. The Artificial Retina consists of 9 square cells which move over and explore a visual image in order to perform shape discrimination (see Figure 144). The input from each cell is either the value of the top left hand pixel (sampling method) or an average of all 9 cells (averaging method). The resolution of the retina can change at run-time and is set by the system at either 15, 30 or 60 pixels per cell. The system is given 50 cycles to discriminate the shape in the visual image (either a square or a triangle), and is able to determine the direction (range [0, 359] degrees) and distance (range[0, 50] pixels) that it moves during each cycle. In each cycle the system performs the following steps:

1. selects the Artificial Retina filtering method;
2. gets the input values from the visual image;
3. computes the output values from the ANN;
4. computes its next position.

The Artificial Retina is prevented from moving outside the bounds of the visual image.

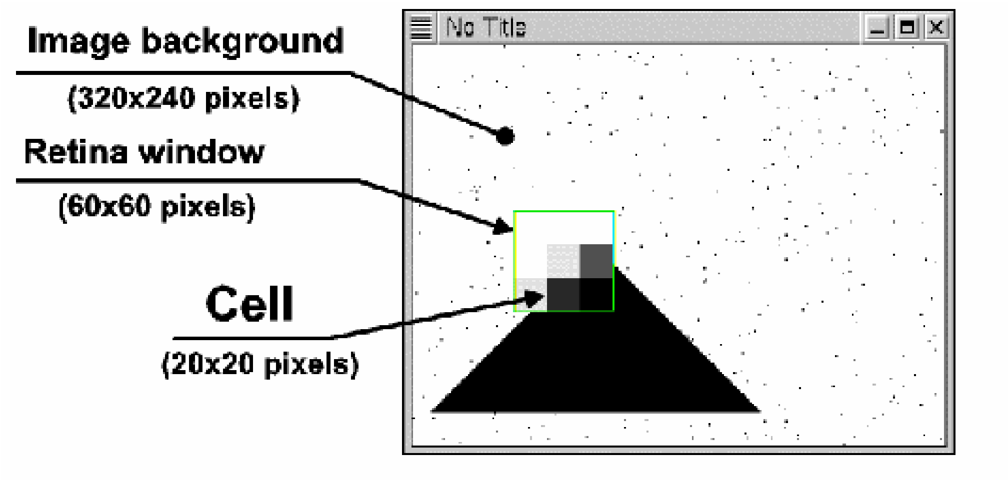

Figure 144 – The Artificial Retina is the green frame (9 cells) superimposed on a visual image, which in this case consists of a triangle.

The ANN underlying the Artificial Retina consists of 17 input units (9 inputs from the cells and 6 recurrent inputs) and 6 output units (1 specifies that the shape is a triangle, 1 specifies that the shape is a square, 1 specifies the filtering method, 1 the size of the retina window, 1 specifies how far it will move (0 to 50 pixels) and 1 the direction it will move (0 to 359 degrees) (See Figure 145).

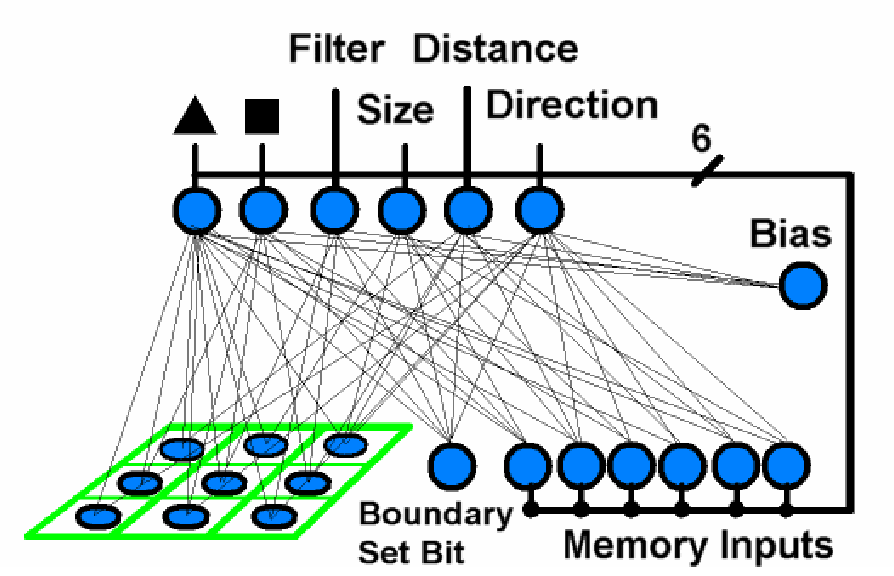

Figure 145 – The neural architecture of the Artificial Retina.

Each of the 6 output units is connected to the 17 inputs units and the strength of each of these 102 fixed strength synapses is encoded using 5 bits on the genotype, representing a connection weight in the range [-4, 3]. During each cycle the system estimates whether the figure in the visual image is a triangle or square (with a default output of triangle). Every individual is tested with 20 images each generation (10 containing triangles and 10 squares) and the fitness of an individual is proportional to number of cycles it takes to identify the feature as a triangle or a square. The triangles and squares are always black on a white background and appear in a random position each trial, with their height randomly set in the range [20, 100] pixels. At the beginning of each trial the memory units are set to zero and the retina size to 240 x 240 pixels (this resolution is only used on the first time step). A population of 100 individuals was evolved for 150 generations and the experiment was repeated 5 times.

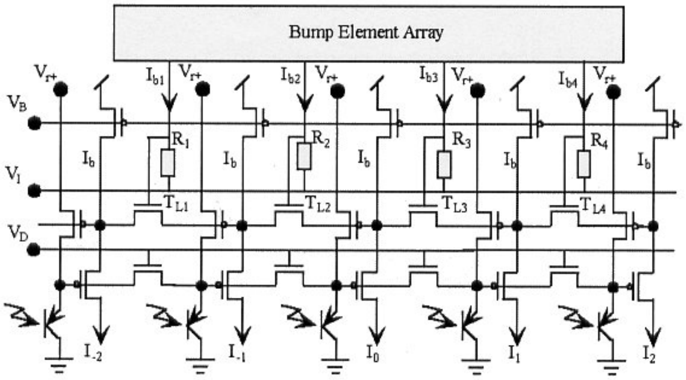

Figure 146 – Artificial retina with locally adjustable shunting inhibition – the circuit may be biased to prefer visual contrast by controlling current through the resistors using the bump element array (Maris 2001)

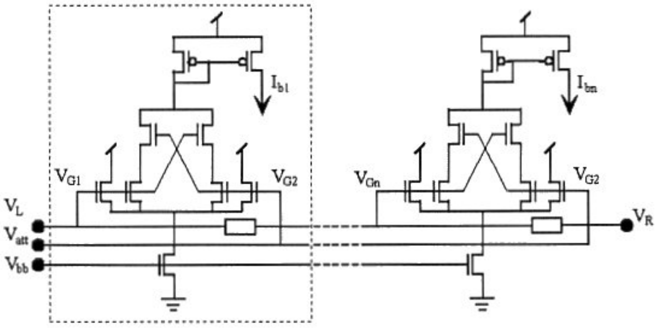

Figure 147 – Bump element array for biasing shunting inhibition in the retina – the voltage V_{att} determines the position of the attention current pattern I_{b1}-I_{bn} flowing into the retina (Maris 2001)

3.11.3.2. Results

Initially the Artificial Retina starts with a low resolution covering most of the image which then shrinks when it moves around. After the Artificial Retina has located the figure in the image, it checks to see if the right hand side is a vertical side (indicating a square). If it does not detect a vertical side it explores the image looking for a vertical side and after a few cycles it starts to look for another figure, outputting its default value of triangle as the shape in the visual image. A variation of this strategy is to scan the corners of images and determine whether they are acute or rectangular [Kato and Floreano 2001].

3.11.4. Robot Navigation

The same method described in Section 3.10.3 was used to evolve the controller for an autonomous robot with a camera that could pan (range [-100, 100] degrees] and tilt (range [-25, 25] degrees), where the task was to navigate in an arena situated in an office environment. The camera inputted a 240 x 240 pixel grey scale image to an ANN consisting of 27 input units (25 photoreceptor units and 2 proprioceptive units encoding the pan and tilt angles) and 5 output units (2 wheel motor units, 1 pan motor, 1 tilt motor and 1 sampling strategy). Each of the output units was connected to a memory unit that fed in its activation from the previous sensorimotor cycle. The 160 genetically-determined, fixed synapse strengths were encoded with 5 bits and took values in the range [-4.0, 4.0]. Each of the output neurons had a threshold value also encoded as 5 bits and in the range [-4.0, 4.0]. The activation of the output units was the net input offset by their associated threshold passed through a sigmoid function.

A population of 40 individuals was evolved for 15 generations. Each individual was tested for 2 trials of a maximum 200 sensorimotor cycles, each cycle lasting 300 ms. If the robot got to close to a wall the trial was stopped. The fitness function rewarded individuals that could move forward as quickly as possible during a trial. Evolution was carried out on real robots with each generation taking about 1.5 hours. A Koala robot was placed in a small square arena (2m x 2m, with 30cm high sides) and the fitness function rewarded controllers that moved the furthest forward without hitting an obstacle.

Successful individuals evolved a strategy of pointing the camera downwards to detect the edge between the dark floor and white walls of the arena. The controller maintained a sampling filtering strategy to enhance the contrast between the floor and the wall. The robot always moved in a clockwise trajectory around the arena. Some controllers used the camera pan to actively coordinate turning away from approaching walls; in other controllers the camera was kept panned to the left and tilted downwards and the turning of the right wheel was coordinated with the expansion of the white wall in the visual field. Both this visual looming behaviour and the edge detection behaviour are linearly-separable mappings between the visual inputs and motor outputs of the system and therefore do not require complex non-linear transformations.

3.11.4.1. Car Driving

In a this set of experiments a co-evolutionary method similar to that described previously was used to evolve controllers that could drive a simulated car over roads in the Swiss alps and it was found that active vision can exploit simple features to accomplish this. The visual scene is 600 x 400 pixels and corresponds to the windscreen of a car. The system was built by extending the Car World simulator (see Figure 148). Colour information from the simulator is converted to grey scale values and then passed to the visual neurons of an ANN which controls steering and forward speed and braking and the position of the retina on the screen and its resolution.

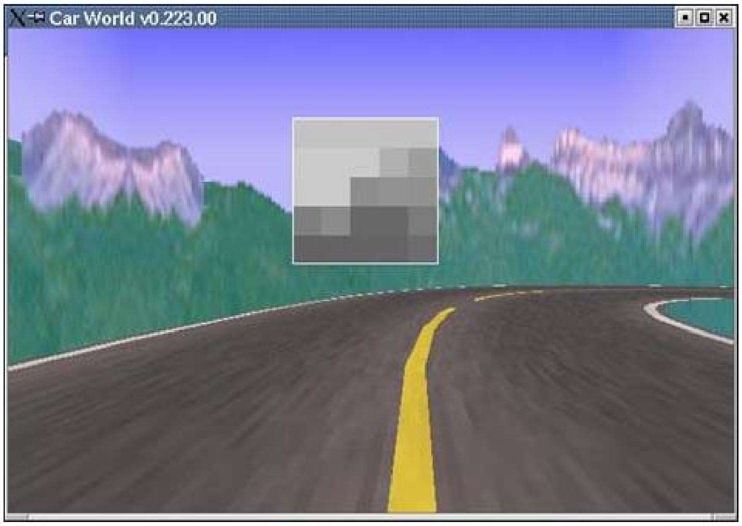

Figure 148 – Image of the Car World simulation

3.11.4.2. General Observations

All three experiments in co-evolving properties of an artificial retina and ANN controllers vary in their environmental conditions but the evolved individuals exhibit similar strategies for active vision and feature selection. The active vision systems develop sensitivity to position specific oriented edges: in the shape discrimination experiment individuals are sensitive to vertical edges on the left or right (depending on the run) of the retinal image, or inclined (60 degree) edges on the right. In the robot navigation the retina is sensitive to horizontal edges in the middle of the retinal images. In the car driving experiment individuals are sensitive to inclined (60 degree) edges on the left or the right of the retinal image. All these features are linearly separable, corners being a linear combination of two inclined edges. The system is constrained to find features of this type as the visual neurons project directly onto motor neurons and it cannot rely on more complex transformations which would require intermediate layers of neurons. The active vision system in all three experiments moves the retina across the image in order to locate and position selected features over matching receptive fields. Co-evolving active vision and feature selection enables individuals to solve position and size invariant tasks using position and size variant mechanisms.

3.11.5. Adaptive Vision-based Flying Robots

The goal of this research is to apply biomimetic and evolutionary methods to generate indoor flying robots that can perform autonomous vision-based navigation. Mainstream computer vision is generally based on a computationally expensive sequence of: pre-processing; segmentation; object extraction; and pattern recognition. It is not therefore a viable approach for robotic systems that have to act in real time. Artificial evolution is used to generate individuals with simple vision systems that are tuned to invariant relationships in their environments, usually by linking simple visual features to motor outputs.

Flying robots have a more complex environment than land-based robots and it is a bigger challenge to extend an evolutionary methodology to this task. Some of the key difficulties involved in artificially evolving flying robots are:

5. They move in 3 dimensions;
6. They cannot be tethered to a computer very easily;
7. It is harder to define and measure performance;
8. The dynamics are more complex. For example, blimps can slip sideways as they are controlled by thrust (speed derivative) and inertial and aerodynamic forces have significant effects.

At EPFL research is initially being carried out on blimps with the goal of extending the techniques to slow flying aeroplanes, such as the C4.

Figure 149 – The Blimp

The main features of the Blimp (see Figure 149) architecture are:

- Wireless connection (Bluetooth 15m range)
- Li-Poly battery (3 hour life)
- Sensors: linear camera; anemometer for fitness; MEMS gyro for yaw rotation estimation; and distance sensors for altitude measurements.

Indoor slow flyers are good testbeds for evolutionary robotics as they require fast reactions, but they are lightweight and have lower power consumption. They are flown indoors to reduce the effects of wind and other weather conditions and to allow for the visual environment to be manipulated easily.

The main features of the C4 slow flyer (see Figure 150) are:

- Wireless connection (Bluetooth)
- Weight 50g
- Wing span 82cm
- 20 minutes power autonomy
- 1.4 m/s minimum flight speed
- Vectored thrust for direction control (motor-gear-propeller articulated around the yaw axis)
- Minimum space for flying = 8 x 8m

- Faster than the Blimp and have 2 more degrees of freedom (pitch and roll)

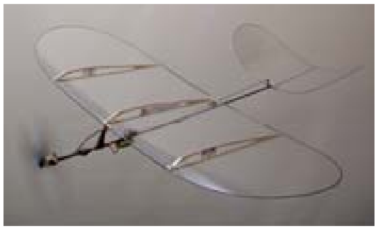

Figure 150 – C4 slow flying aeroplane

3.11.5.1. Current Research

Currently, a flight simulator is being developed for evolving controllers in simulation in order to speed up the evolutionary process. The plan is to evolve in simulation and transfer to the physical aeroplane for further incremental evolution. OpenGL and ODE are being used as they can accurately simulate 3D motion, gravity, inertia, shocks and friction. The Blimp dynamical model includes buoyancy, drag, Coriolis and added mass effects. They have so far established a good behavioural correspondence between the simulation and the physical Blimp. A model for the indoor slow flyer is currently under development.

In 2001 Floreano and his co-workers first evolved spiking neural network to control a Khepera with a simple vision system that wandered in an arena with random black and white vertical stripe patterns on its walls (see Section 6 for details of this experiment and an overview of spiking networks). The best individuals were able to move forward and avoid hitting walls. Following on from this research, the Blimp controllers have adaptive connections using spiking networks with Hebbian-like learning rules in order to facilitate transfer between simulation and reality. So far, Blimps have been evolved in simulation and successfully transferred to reality that can avoid obstacles, back away from obstacles and extricate themselves if they have collided with obstacles.

3.11.6. Blimp Experiments

The evolutionary algorithm and the spiking circuits communicate with a physical Blimp via a Bluetooth connection (15 m range). A simple linear camera is attached to the front of the Blimp's gondola and the 150 degree field of vision is mapped onto 16 photoreceptors. The activations of these receptors are convolved with a Laplace filter which detects contrast over three adjacent receptors. These 16 receptors are connected to 10 spiking neurons, with the connectivity pattern and sign of the connections determined by an evolutionary algorithm. Each neuron is represented by 27 bits, the first bit determining its sign (1, -1) and the other 26 bits indicating the presence or absence of a connection to the 10 neurons and 16 photoreceptors. All of the connections have a strength of 1. The parameters of the spiking model, which determines the response profile of the neurons, refractory periods and time delays due to connection length are fixed.

A population of 60 individuals was evolved for 20 generations, with each experiment taking 2-3 days. Each individual was tested twice on the robot for 40 seconds with the fitness function rewarding fast forward motion. The position of the robot was randomly set in between trials. The experiment was repeated 5 times.

3.11.6.1. Results

In all runs within 20 generations the best individuals could navigate around the room, using collisions with the walls as part of their strategy as these were not penalised by the fitness function (see Figure 151 below) [Zufferey and Floreano 2002; Nicoud and Zufferey 2002]

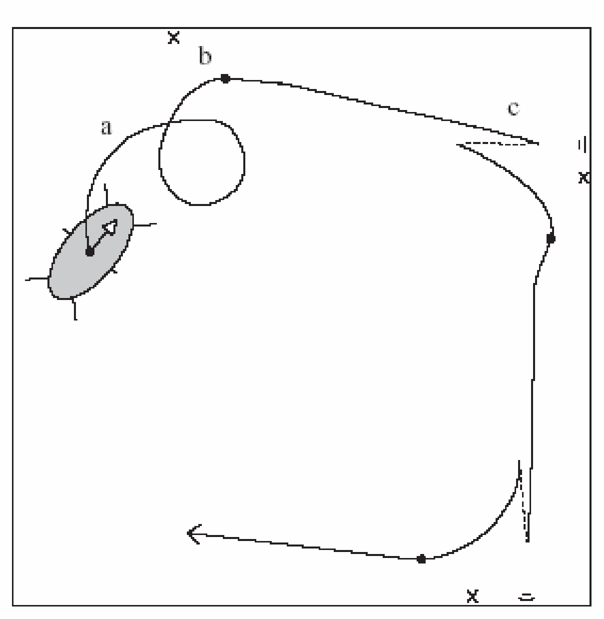

Figure 151 – Hand draw estimation of the typical path of a best individual (solid line = forward movement; dashed line = backward movement; small curves = front collision; cross and circle = place of collision with left back bumper)

3.11.7. Evolution of Adaptive Spiking Circuits for Vision-Based Behavioural Systems

The goal of this research, also carried out at EPFL, is to evolve controllers for adaptive visual navigation and to investigate assumptions about visual processing in behavioural systems. These experiments use spiking neural controllers as they are efficient at processing time-dependent sequences, they are biologically plausible and they can be implemented in embedded digital hardware.

Most biological neurons communicate by sending pulses or 'spikes' to other neurons (see Figure 152). Once a spike has been generated there is a short refractory period before the neuron can fire another short, transient signal. This research uses the Spike Response Model (SRM) which exhibits rich dynamics without being computationally expensive. Hebbian-like rules can be used in spiking networks, where they are referred to as Spike-time Dependent Plasticity (STDP). The adaptation of the synaptic strength is a function of the *time difference* between the firing of the pre- and post-synaptic neurons: if the latter fires just after the former, then it means that the first spike participated in firing the second one and the connection is reinforced.

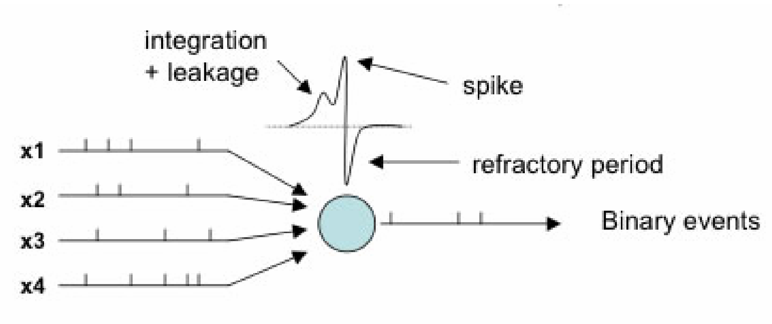

Figure 152 – Diagram of an artificial spiking neuron

These are 5 steps involved in a spiking neuron algorithm:

1. Refractory period: check if the neuron has emitted a spike within a specified period, and if it has do not update its membrane potential;
2. Compute incoming spikes: sum the spikes weighted by the signs of the emitting neurons;
3. Update the membrane potential: add the contribution of the incoming spikes to the current potential, but do not go below a minimum level;
4. Generate a spike: if the membrane potential is above a threshold, set the neuron output to 1 and the membrane potential to its minimum level; otherwise set the neuron output to 0 and do not change the membrane potential.;
5. Leakage: subtract a leakage constant from the membrane potential, but do not go below a minimum potential.

3.11.7.1. Khepera Experiments

The task was to navigate in a rectangular arena with textured walls (see Figure 153), consisting of black and white stripes of random size. Fitness was proportional to the amount of forward motion during a 40 second trial, which implicitly selected for robots that could avoid collisions. The architecture was a fully recurrent 10 spiking neuron network connected to 16 visual receptors. Each neuron was encoded with a bit string, 1 bit encoding the sign of the neuron (+/- 1) and $n + s$ bits encoding the presence or absence of connections from the n neurons and s receptors in the network.

The strength of all the connections was 1. The parameters of the spiking neuron model were fixed. A population of 60 individuals was evolved, with each controller being tested for 40 seconds on the robot.

3.11.7.2. Results

After 30 generations the best individuals could move forward and avoid walls. When the same experiment was done with the spiking units replaced by sigmoidal units none of the individuals were able to move forward and avoid collisions. The spiking networks use temporal correlations between photoreceptors and internal neurons to avoid walls and move forwards, with higher fitness individuals displaying smoother trajectories in the environment (see Figure 154).

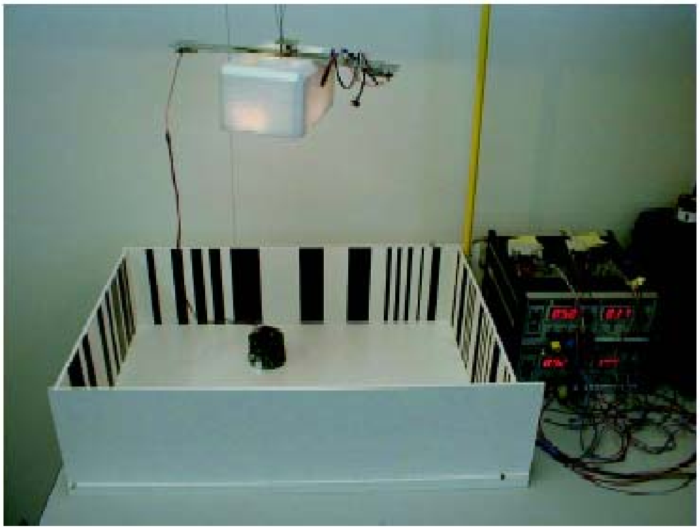

Figure 153 – Experimental set up

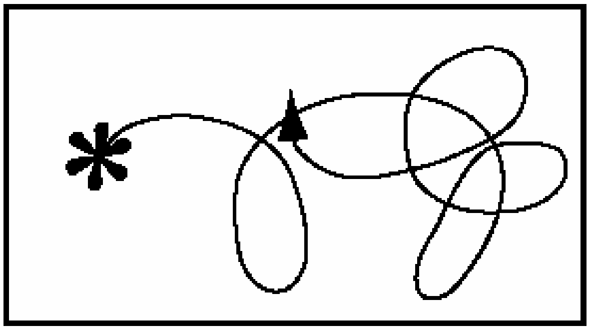

Figure 154 – Typical trajectory of a robot in the arena shown in Figure 152 controlled by a spiking neuron controller. The asterix is the starting point. The curvature of the trajectory is dependent on the pattern of black and white stripes seen by the robot.

3.11.7.3. Implementing Spiking networks in Hardware

Further experiments have mapped spiking circuits directly onto microcontrollers, representing spikes as binary events. Consequently, it is sufficient to use simple logic operations, such as NOT and AND, and instructions to move single bits to build large networks of spiking neurons that have the potential to display complex behaviour. This is in contrast to the neuron model used in the experiment described previously which requires floating point precision and non-linear functions and is therefore computationally expensive.

3.11.7.3.1. Experiments

The circuit connectivity was the same as the ANN described above and both the evolving population and the spiking circuits were evolved on a PIC (Peripheral Interface Controller) chip that controlled an Alice microrobot, which has 4 IR sensors and 2 wheel motors (see Figure 155).

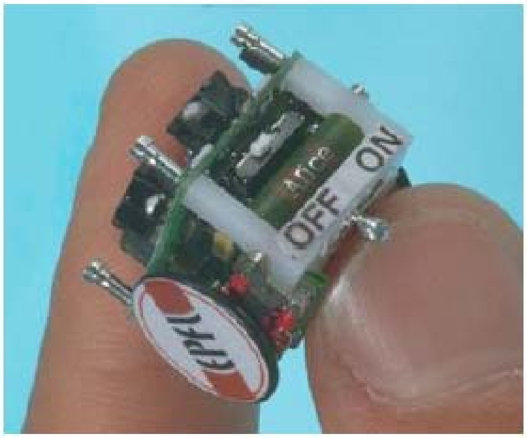

Figure 155 – An Alice microrobot

Each neuron was encoded as in Section 6.1 and an entire circuit required only 17 bytes (1 byte for the sign of the neurons, 8 bytes for the neural connections and 8 bytes for its sensory connections). The memory constraint of the microcontroller meant that only a small population of individuals could be evolved, so a steady-state GA was used. At each time step an individual was chosen at random, mutated and then had its fitness measured. If the fitness was equal to or higher than the worst individual in the population it replaced it, otherwise it was discarded.

The network consisted of 8 neurons and 8 sensor units mapped to the front left, centre and right IR sensors. Every neuron received connections from all of the sensors but the connections between the neurons and the signs of all the connections were evolved. The network was updated every 2ms and the sensor input was computed every 28 ms. The velocities of the wheel motors were set proportional to the number of spikes emitted by the motor neurons in the preceding 28ms period. The speed of the robot wheels was set as the difference in spike count between the two motor units, one neuron moving it forwards and the other backwards, with speeds in the range [-4,4].

The robot was positioned in a 25 x 18 cm arena and the task was to move forward without colliding with the walls, using a modified version of the fitness function described above. A population of 6 individuals was evolved for 3 hours with each individual being tested for 10 seconds and with 3 seconds of random movement between trials. If at any time one of the wheel speeds was set in backwards rotation the fitness of the trial was set to 0. The experiment was run 7 times and the best individuals stored offline every 3 minutes.

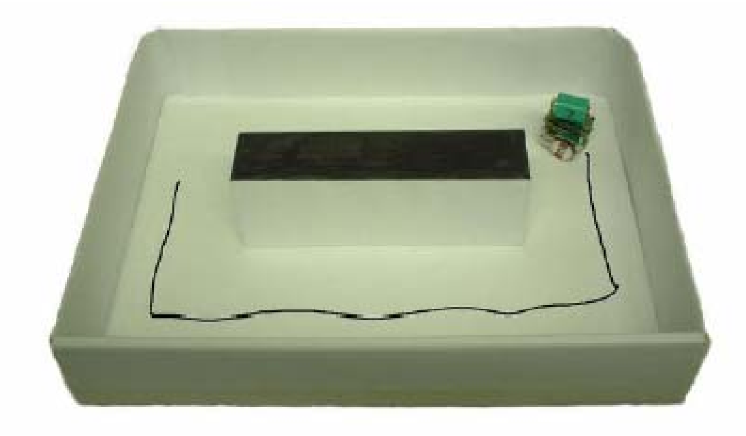

Figure 156 – The testing arena for the Alice robot. The trajectory is the distance covered in 10 seconds overlayed from video footage. The arena is 25 x 18cm.

3.11.7.3.2. Results

Within 2 hours the best individuals were able to move around the arena. Successful robots exhibit wall following behaviour; less fit individuals turn more often and therefore remain in the same area of the maze but also manage to move away from walls [Floreano, Schoerni et al. 2002; Zufferey et al. 2003].

3.11.7.4. Walking Robots

There has been a great deal of work on evolving controllers for various kinds of walking robots – a non-trivial sensorimotor coordination task. Gallagher et al (1996) describe experiments were neural networks controlling locomotion in an artificial insect were evolved in simulation and then successfully downloaded on a real 6-legged robot. In this approach, each leg was controlled by a fully interconnected network of 5 Hopfield-like continuous neurons [Hopfield 1984] each receiving a weighted sensory input from that leg's angle sensor. A genetic algorithm has been used by Galt et al (1997) to derive the optimal gait parameters for a Robug III robot, an 8-legged, pneumatically powered walking and climbing robot. The individual genotypes were encoded to represent the phase and duty factors, i.e. the coordinating parameters that represent each leg's support period and the time relationships between the legs. Controllers were thus evolved that have been proved capable of deriving walking gaits that are suitably adapted to a wide range of terrains, damage or system failures. Sensory information can be used by neural networks to provide forecast of the terrain conditions and hence improve the walking efficiency. Gomi and Ide (1997a,b) evolved the gaits of an 8-legged robot using genotypes made of 8 similarly organized sets of genes, each gene coding for leg motion characteristics such as the amount of delay after which the leg begins to move, the direction of the leg's motion, the end positions of both vertical and horizontal swings of the leg, the vertical and horizontal angular speed of the leg, etc. After a few dozen generations, a mixture of tetrapod and wave gaits is obtained. Using a developmental approach called Cellular Encoding [Gruau 1995] – which genetically encodes a grammar-tree program that controls the division of cells growing into a discrete-time dynamical recurrent neural network – Gruau & Quatramaran (1997) evolved single-leg neural controller for a walking robot which generated a smooth and fast quadrapod locomotion gait. Jakobi (1998) has successfully used his minimal simulation techniques (ultra-fast, ultra-minimal robot simulations that guarantee successful transfer of results from simulation to reality) to evolve controllers for the same 8-legged robot (shown in Figure 157). Evolution in simulation took about 2 hours only, and then transferred successfully to the real robot. Jakobi evolved modular controllers based on Beer's continuous recurrent networks to smoothly control the robot engaged in walking about its environment avoiding obstacles and seeking out goals. The robot could smoothly change gait, move backward and forward and even turn on the spot. Recently there has been successful work on evolving coupled oscillator style neural controllers for the highly unstable dynamic problem of biped walking. Reil and Husbands (2002) showed that accurate physics based simulations using physics-engine software could be used to develop controllers able to generate successful bipedal gaits. Work at Zurich University has extended this idea to various other kinds of robots and gaits, including a brachiating robot able to swing across a rope like an ape.

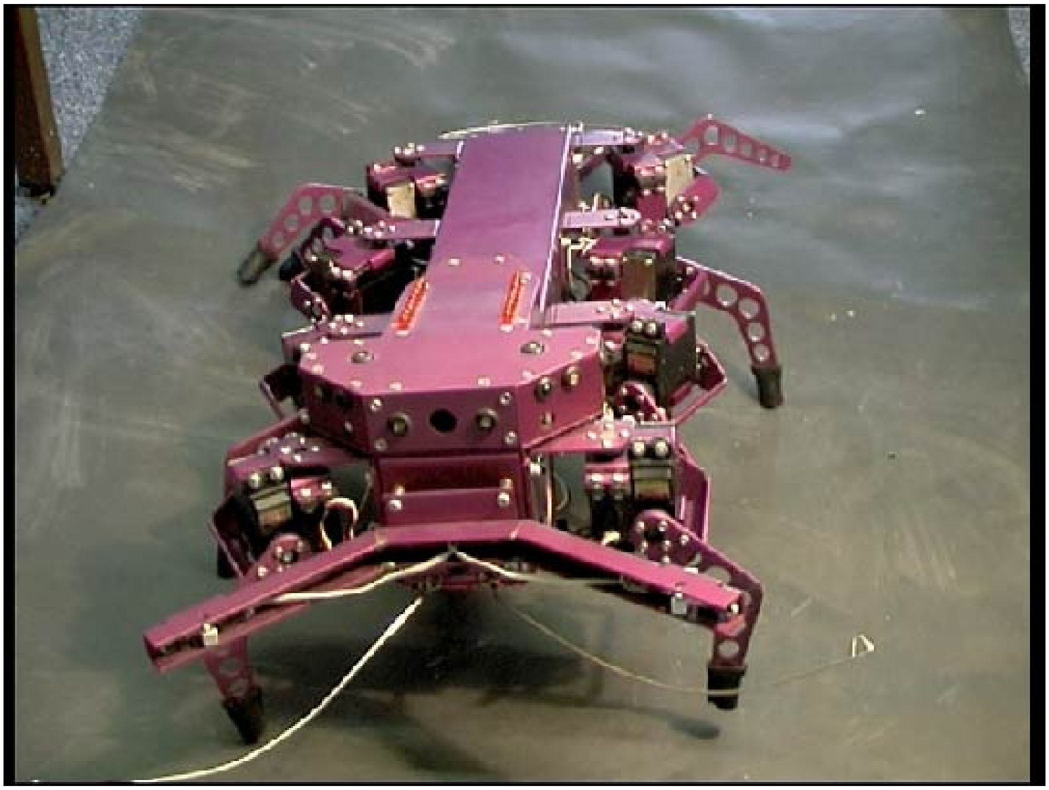

Figure 157 – 8-legged robot used in various evolutionary robotics research.

Recently McHale and Husbands [2004a, 2004b] have compared many forms of evolved neural controllers for bipedal and quadrapedal walking machines. GasNets were shown to have advantages in most circumstances.

These techniques of evolutionary robotics are applicable to other forms of locomotion as well as legged robots. However, it is important to note that these studies were conducted for relatively benign environments such as indoor mazes, etc without consideration of soil and terrain trafficability.

3.11.8. Evolution of Adaptation Rules

An important strand of work on evolutionary robotics in the Autonomous Systems Laboratory at Ecole Polytechnique Federale de Lausanne (EPFL), under the direction of Professor Dario Floreano, has been the evolution of plastic neural networks. That is networks that adapt during their operation according to rules that allow their connection strengths to changes depending on activity in the network. This exploration of evolved plastic controllers has some similarities with the Sussex work on GasNets described earlier. The form of plasticity is quite different, but the goal of evolving an adaptive system is the same. In all of the ER research at EPFL robot controllers are implemented as artificial neural networks (ANNs) and artificial evolution is used to determine the properties of the neurons, synapses and occasionally the architecture of these networks.

The goal in this research is to evolve robots with the ability to adapt online, in contrast to the more conventional approach in ER where a hard-wired solution to a particular task is evolved. The conventional approach involves evolving fixed, genetically-determined properties of a robot controller, for example the synapse strengths and signs (excitatory or inhibitory) and neuron properties of an ANN. This approach has 3 potential weaknesses:

1. if the actual environment is different to the evolutionary environment, there is no guarantee that the evolved solution will work;
2. hard-wired mechanisms rely on invariants, that is, constant relationships in the environment, or in the agent-environment interaction, and there is no guarantee that this approach will scale up to more complex tasks;
3. invariant relationships are often not obvious to an external observer, and it is therefore difficult to predict which changes in environment will lead to the failure of a hard-wired mechanism.

One way to overcome the limitations of the conventional approach is to evolve systems that have the potential to adapt online to any changes in the environment. One goal of the research at EPFL is to evolve ANN robot controllers that can continuously modify their synaptic strengths using plasticity rules encoded in the genotype. An associated goal is to evolve behaviours that are more complex than those which can be solved using hard-wired controllers by enabling robots to incrementally learn behaviours and store them. There are 4 areas of research that investigate these issues:

1. Reactive Navigation
2. Competitive Co-Evolution of Adaptive Predator-Prey Robots
3. Sequential tasks: the Light-Switching Problem
4. Adaptation to Unpredictable Changes

3.11.8.1. Methodology

In order to evolve ANN robot controllers that adapt online, the genotype represents Hebbian learning rules that determine how the synapses are changed at each time step Evolution determines which variant of four Hebbian learning rules (plain, postsynaptic, presynaptic, covariance: see Figure 158) governs each synapse in the ANN as well as its associated learning rate and sign. This is in contrast to the conventional approach to evolving ANNs where the fixed synapse strength and sign of each connection are evolved.

$$\begin{aligned}
\text{Plain Hebb:}\quad \Delta w &= (1-w)xy && (1)\\
\text{Presynaptic rule:}\quad \Delta w &= wx(-1+y)+(1-w)xy && (2)\\
\text{Postsynaptic rule:}\quad \Delta w &= wy(-1+x)+(1-w)xy && (3)\\
\text{Covariance rule:}\quad \Delta w &= \begin{cases} F(x,y)(1-w) & \text{if } F(x,y)>0 \\ F(x,y)w & \text{otherwise} \end{cases} && \\
\text{where } F(x,y) &= \tanh(4(1-|x-y|)-2) && (4)
\end{aligned}$$

Figure 158 – The form of the Hebbian learning rules

In the formulae in Figure 158, x represents the presynaptic neuron, y the postsynaptic neuron and w the strength of the synapse connecting them. The synapse strength is kept in the range [0,1] and the sign of each of the synapses is specified genetically and remains fixed. Each synapse also has a genetically encoded fixed learning rate. In each of these 4 variants of the Hebbian rules the change in synaptic change is completely determined by the activation of the pre- and postsynaptic neurons. The plain Hebb rule (1) only allows synaptic strengthening, with the greatest increase when both the pre- and post-synaptic neurons are at their maximum activation. The presynaptic rule (2) enables strengthening and also weakening of synaptic strength, when the presynaptic unit is active and the postsynaptic one is not. Conversely, the postsynaptic rule (3) causes a weakening of synaptic strength if the postsynaptic neuron is active and the presynaptic one is not, increasing the synaptic strength if both the pre- and postsynaptic neurons are activated The covariance rule (4) increases or decreases the strength of a synapse relative to the difference between the pre-and postsynaptic neurons. If the difference in activation is greater than half of the maximum activation level then the synaptic strength is reduced in proportion to this difference, otherwise it is increased in proportion to the difference.

In the adaptive synapse approach, synaptic strengths are not encoded, rather they are initially set to small random values (range [0.0, 0.1]), so individuals have to learn how to solve tasks from scratch, changing their synaptic strengths through their interactions with the environment. The general equation for change in synaptic strength is:

$$w_{t+1} = \eta \Delta w_{t+1} + w_t$$

where:

η is the learning rate
Δw *the change in weight calculated by applying one of the four Hebbian learning rules*
t *represents the current cycle or time step.*

Generally, when ANN robot controllers are evolved at EPFL, the encoding schemes illustrated in Figure 159 are used. Genetically-determined, hard-wired controllers encode each synapse with 5 bits, the first bit specifying the sign (excitatory or inhibitory) and the next 4 bits specifying the strength of the synapse.

Different experiments use different synapse strength ranges. When evolving ANNs with adaptive synapses, 1 bit specifies the sign of the synapse, the next 2 bits specify which of the 4 Hebbian learning rules is applied to the synapse and the final 2 bits the learning rate, which is generally 1 of 4 fixed values in the range [0.0, 1.0].

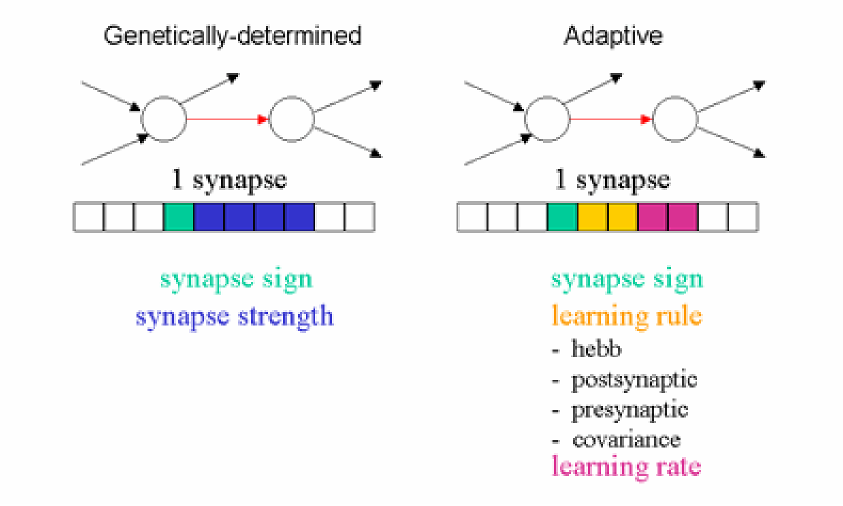

Figure 159 – Encoding schemes for genetically determined and adaptive synapses used at EPFL

At EPFL, artificial evolution is usually carried out on actual robots and sometimes in simulation, with the resulting controllers being transferred onto real robots for testing in the real world. The experiments described in this report use 3 different robotic platforms: Khepera (Figure 160), Koala (Figure 166) and Alice (Figure 155).

3.11.9. Reactive Navigation

The methodology for artificially evolving adaptive robot controllers outlined in Section 1 was first used in 1996 where it was applied to a simple navigation task where the goal was to avoid obstacles and move in as straight a line as possible within a looping maze. The fitness function used was:

$$\Phi = V(1-\sqrt{\Delta v})(1-i)$$

where:

V *is the average rotation speed of the two wheels*
*Δ*v *is the absolute value of the difference between the signed speeds of the wheels (where positive represents one direction and negative the other)*
i *is the normalised value of the infrared (IR) sensor with the highest activation*

These three components are respectively maximised by: speed; straight motion; and distance from objects. The experiment used Khepera robots and the ANN consisted of one hidden neuron and two motor neurons. All three of these neurons were connected to each of the robot's 8 IR sensors and the hidden unit, resulting in an ANN with 27 synapses. The genetic encoding also specified whether each synapse had either a *driving* or a *modulatory* effect on its postsynaptic neuron; the former determined whether the output of the unit was below or above 0.5, whereas the latter could only enhance or dampen the response, but not change its sign. The speed and direction of each of the wheels was dependent on the activation of an associated motor neuron, with activation values less than 0.5 turning the wheel in one direction and greater than 0.5 turning the wheel in the opposite direction. The genotype encoded four properties of each synapse: driving or modulatory (1 bit); its sign (1 bit); the learning rule (2 bits); and the learning rate (2 bits encoding 4 distinct values: 0, 0.33, 0.66 and 1).

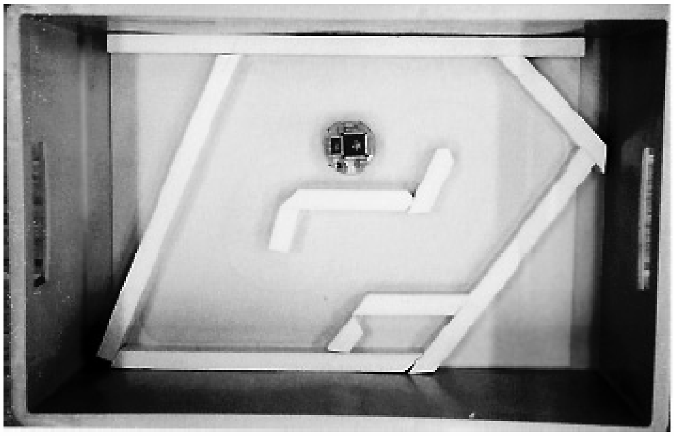

Figure 160 – Bird's-eye view of the looping maze and a Khepera robot used in the reactive navigation experiments

The population consisted of 80 individuals which were evaluated for 80 sensorimotor cycles (about 24 seconds) on a real Khepera robot every generation. The robot was randomly moved for 5 seconds between each individual trial and a generation took approximately 40 minutes. There were 5 main results [Floreano and Mondada 1996, 1998]:

1. The adaptive synapse approach took about half the evolutionary time to solve the task compared to a genetically encoded network, with the maximum fitness reached around the 50th generation.
2. All the best individuals at the end of each evolutionary run moved in the same direction around the arena, although this was not specified by the fitness function, as this resulted in the highest number of IR sensors being activated by the arena walls, thereby reducing the probability of collisions.
3. Obstacle avoidance behaviour developed within a few seconds of the robot being placed in the maze and the synapse strengths and neuron activation values randomly initialised; furthermore the behaviour was continuously refined while the individual moved around the maze;
4. Although the synaptic strengths continuously changed while the robot was moving, the overall pattern of change became stable, with any decreasing weight strength being compensated for by another weight increasing in strength;
5. The robot controllers with adaptive synapses were able to adapt to new environments, in contrast to the hardwired controllers.

3.11.10. Competitive Co-Evolution of Adaptive Predator-Prey Robots

In these experiments a predator robot tries to hit (catch) another prey robot, which tries to escape. The fitness of the prey robot is measured by how quickly it can catch the prey; the fitness of the prey is determined by how long it can escape the predator. Two Khepera robots are used in this experiment, each of which have 8 proximity (IR) sensors around their bodies. The predator also has a vision system: a 1D array of photoreceptors with a 36 degree field of vision (Figure 161, right). The prey is able to move twice as fast as the predator and also has a black protuberance that enables the predator's vision system to detect it anywhere in the test arena (Figure 161, left). When the adaptive synapses of the robots' ANN controllers are co-evolved, interesting chasing and evasion behaviours emerge. However, after many generations the co-evolutionary system displays a *limit cycle* where in some generations the predator wins and in others the prey.

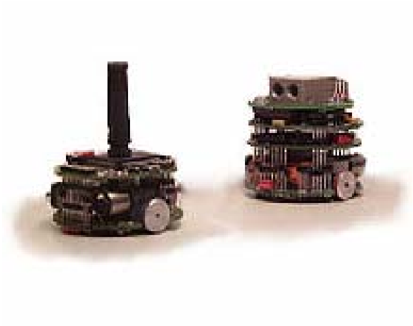

Figure 161 – Predator (right) and prey (left) Khepera. The prey robot has a black protuberance that can be detected anywhere in the arena by the predator's vision module (1D-array of photoreceptors with a 36 degree field of vision)

The ANNs consisted of 2 output neurons. In the case of the predator each of these units was connected to five photoreceptors in the vision module and the robot's 8 IR sensors; in the case of the prey each of the output units was connected to the 8 IR sensors. The output of each of these units was offset by a threshold and passed through a sigmoid function, resulting in an activation in the range [0,1]. Values below 0.5 led to the wheel turning in one direction, values above 0.5 in the opposite direction, and 0.5 in no motion. The resulting wheel speed of the prey was multiplied by 2. The properties of the synapses were encoded with 5 bits on the genotype: 1 bit for the sign and the other 4 bits for the other properties. In the fixed weight controllers these four bits encoded the strength of the synapse; in the learning condition 2 bits encoded the learning rate (0.0, 0.337, 0.667, 1.0) and 2 bits the Hebbian rule; in a third condition 2 bits encoded a synaptic strength in the range [0,0.1] and 2 bits the upper and lower bounds of a uniform noise distribution (0, +/- 0.337, +/- 0.667 and +/-1.0). In this third *adaptive noise* condition, in each sensorimotor cycle every synapse had a random noise value added to the synaptic strength (with checks to ensure it stayed in the range [0,1]. In all 3 conditions, the predator's genotype encoded 30 synapses and 2 output neuron thresholds and the prey's 20 synapses and 2 thresholds.

Each condition was run 6 times with two populations of 100 individuals which were co-evolved for 1000 generations. No statistical difference was found between any of these runs. Each individual was tested against the best individuals from the previous 10 generations to facilitate co-evolutionary stability. The individuals were always positioned in the middle of the arena and the same distance from each other, but their orientations were randomly set for each test. These lasted 500 sensorimotor cycles or until the predator and prey touched. In these experiments there were 10 sensorimotor updates per second and each run lasted between 8 to 10 hours depending on the performance of the competitors: if the predators caught the prey quickly then a test would last less than the maximum 500 sensorimotor cycles. The time to contact, *TtC,* measured in sensorimotor cycles, determined the fitness of an individual (range [0.0, 1.0]): the predator receiving 1 – TtC/500 and the prey TtC/500.

3.11.10.1. Results

There were 2 significant results in the experiments. Firstly, robot controllers with adaptive synapses always catch the prey by adapting their online behaviour to the prey behaviour. Secondly, when predators and prey can select whether to use Hebbian learning or random changes, predators always use adaptive mechanisms whereas prey almost always choose random weight changes [Floreano and Nolfi 1997a; Floreano, Nolfi and Mondad 2001].

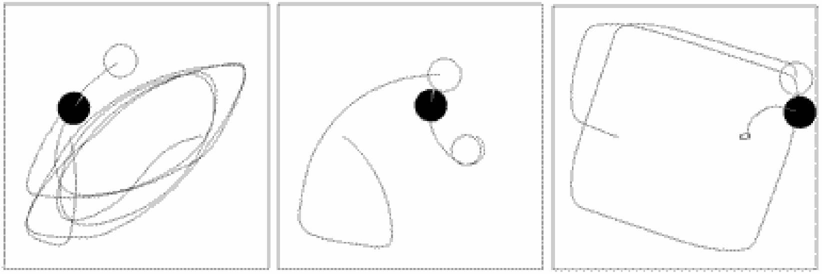

Figure 162 – Typical trajectories from three tournaments between a predator (black disk) and prey (white disk). In c the predator does not chase the prey, rather it backs up close to a wall and waits for the prey to approach, as they tend to keep close to walls. The proximity sensors of the prey do not detect the predator in time for it to steer away and the two robots collide. Situation a is earlier in evolutionary history than b which is earlier than c. Here we see a clear evolution of more refined, and quite unexpected, strategies for both the predator and the prey.

3.11.11. A Sequential Task: The Light-Switching Problem

In this experiment a Khepera robot with a vision system and a floor sensor is put in an arena with a light bulb and a light switch, marked by a black stripe on the wall and toggled depending on the vicinity of the robot to the black stripe as shown in Figure 163. The fitness of an individual is determined by the amount of time that it spends in the vicinity of the light when it is switched *on*. The light is *off* at the beginning of the experiment, so the robot must move to the switch and turn the light on. Turning on the light is not *explicitly* rewarded by the fitness function.

The ANN is fully recurrent and consists of 10 sensory units (4 IR sensors, 3 ambient light sensors, 3 photoreceptors) and 2 motor units. The activation of each unit is the sum of the incoming synaptic signals, plus any sensory inputs, passed through a sigmoid function resulting in an activation in the range [0,1]. The 144 synapses were either encoded individually as 5 bits on the genotype, or using a node encoding where all incoming synapses to each neuron had the *same* properties, resulting in a shorter length genotype. Robot controllers with genetically-determined and noisy synapses were individually encoded, whereas those with adaptive synapses were node encoded. 100 individuals were evolved for 200 generations and the experiment was repeated 10 times. Each individual is tested in the arena for 3 trials, each lasting 500 sensorimotor cycles, with each cycle lasting 100ms. At the beginning of the test the robot is placed in a random position and orientation. The initial experiments were carried out in simulation and the controllers successfully transferred to Khepera robots.

Figure 163 – Bird's-eye view of the light switching arena. The Khepera switches on the light if it passes over the black area and gains fitness by staying in the grey area while the light is on. The robot can detect ambient light and the colour of the wall but not the colour of the floor.

3.11.11.1. Results

The four main results from this experiment were:

1. adaptive individuals evolve much faster (in about half the number of generations) and evolve to a higher fitness (15% greater) than genetically-determined individuals;
2. adaptive individuals decompose their behaviour into a set of sub-behaviours: avoiding walls; phototaxis; and moving towards the light switch. In contrast, genetically-determined individuals solve the problem using a minimal strategy: looping around the arena until they switch on the light and then eventually ending up near the light;
3. Increasing the size of the neural network by 10 hidden units does not affect the results when evolving Hebbian rules for adaptive controllers, but does affect genetically-determined individuals as it prevents them from evolving minimal solutions to the task.
4. Experiments with noisy synapses showed worse performance than either adaptive or genetically-determined synapses, whether synapse or node encoded. This indicates that the effect of using Hebbian learning rules is more than random synaptic variability.

An adaptive individual was displayed in a museum for two months with visitors able to randomly initialise the weights and then watch the controller develop the sub-behaviours. Its behaviour was robust to the museum lights which differed from the environment it had been evolved in [Florean and Urzelai 1999; Florean and Urzelai 2000].

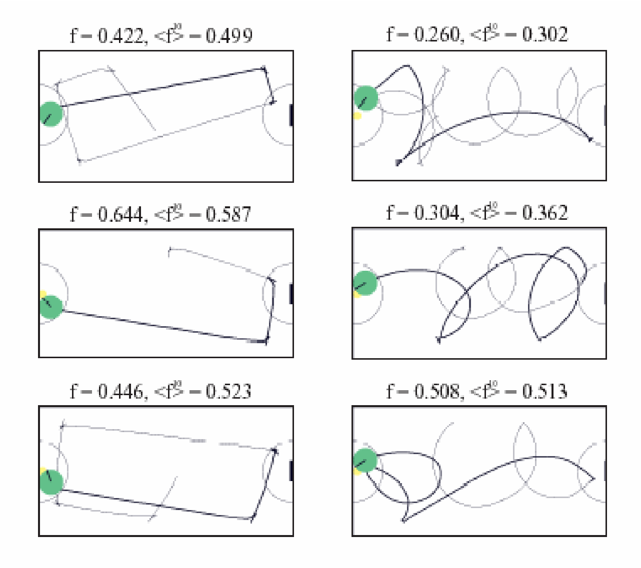

Figure 164 – Comparison of the behaviour of the three best individuals with adaptive synapses and node encoding (left) compared to the three best individuals with genetically-determined synapses and synapse encoding (right). The individuals come from the last generation of three different runs of the two conditions. The fitness value for the particular trial and the average fitness over 10 trials from different start positions and orientations is above each box. The trajectory line becomes thicker when the light is turned on. The adaptive controllers have straighter trajectories and are able to remain in the fitness area, in contrast to the loopy trajectories of the genetically-determined individuals.

3.11.12. Adaptation to Unpredictable Change

Evolving robots with adaptive controllers that have to adapt to random parameter initialisation produces systems that are robust to environmental changes *not* experienced during evolution. This evolutionary methodology prevents controllers from genetically encoding particular environmental invariant relationships in their synaptic properties and forces them to adapt online to the particular environment in which they find themselves, without any necessity for further incremental evolution. This has been shown in a number of ways:

1. When the controllers are transferred from simulation to reality the adaptive synapse controllers continue to behave with only a small reduction in fitness, reflecting a difference between the simulation and actual physics of the real world. In contrast, individuals with fixed, genetically-determined controllers often do not transfer successfully, or display a greatly impaired performance.
2. Proximity (IR) sensors respond less to dark walls than light ones and therefore dark walls seem further away. When the colour of the walls in the testing environment is changed, adaptive synapse controllers can successfully adapt whereas often genetically-determined ones hit the walls and remain stuck.

3. Changing environmental layout, such as moving the position of light and switch in the sequential switching task described above, results in genetically-determined individuals failing the task, as their trajectory is specific to a particular environment. In contrast, adaptive individuals can adjust to these changes as seen in Figure 165.
4. Changes in morphology, such as those that occur when a controller is downloaded onto a different robot (for example, going from a Khepera to Koala as in Figure 166, which are different sizes and have different layout of sensors) does not affect adaptive synapse controllers but prevents genetically-determined ones from doing the task [Urzelai and Floreano 2000a,b].

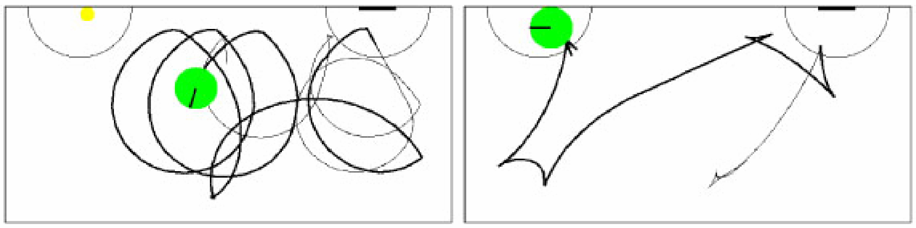

Figure 165 – Comparing the trajectories of fixed weight (left) and adaptive (right) controllers when the arena layout in the light switching task is changed. When the light switch and bulb are put in different positions to those seen during the evolutionary process the fixed, genetically-determined controllers always fail the task whereas the adaptive controllers always succeed.

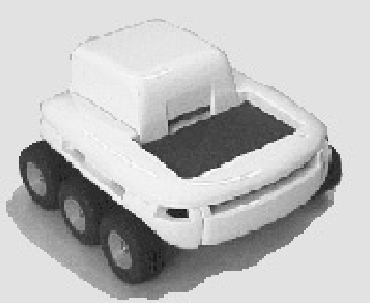

Figure 166 – The Koala robot

3.11.13. Evolution of Learning-Like Behaviours

Learning in ANNs is traditionally thought of in terms of modifications to weight strengths using techniques such as back propagation or Hebbian learning. However, Yamauchi and Beer demonstrated in 1994 that *fixed strength* continuous time recurrent neural networks (CTRNNs) could display reinforcement learning behaviour by exploiting the internal dynamics of the network. This inspired research at EPFL into robot navigation comparing CTRNNs similar to those used by Yamauchi and Beer with ANNs where the weight strengths changed according to Hebbian rules (as described in previously). These experiments demonstrate that it is not always neural activity that is responsible for behaviour or weight changes that are responsible for learning.

3.11.13.1. T-maze Experiments

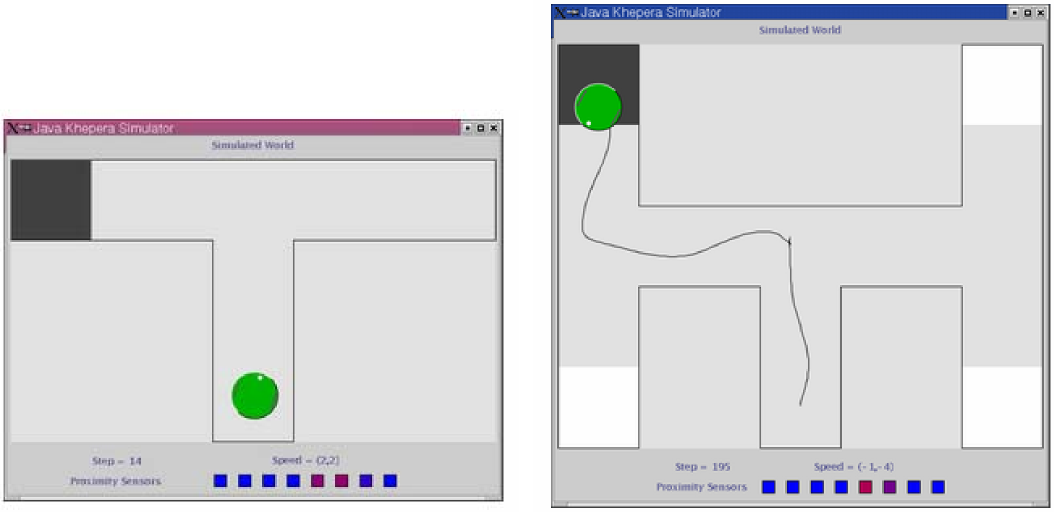

Step = 14 Speed = (2,2)

Step = 195 Speed = (-1,-4)

Figure 167 – a) T-maze and b) Double T-maze experimental set ups

In this experiment the task is to find a black area (reward zone) that is positioned in either the right or the left arm of the T-maze. The ANN consists of 4 hidden units and 2 motor output which are fully connected and also have connections to 5 sensor units (4 IR sensors and 1 floor sensor).

Each neuron's time constant and threshold is encoded using 5 bits, and the strength of each of the 66 synapses is encoded with 5 bits. The length of each controller genotype is 390 bits. The value of each sensory receptor is scaled in the range [0, 1] and the motor output is scaled in the range [-10, 10]. The activation of each neuron is calculated as follows:

$$\gamma_i(n+1) = \gamma_i(n) + \frac{\Delta t}{\tau_i}\left(-\gamma_i(n) + \sum_{j=1}^{N} w_{ij} A_j(n) + \sum_{k=1}^{S} w_{ik} I_k \right)$$

where:

N is the number of neurons
i is the index (= 1, 2, …, N)
γ_i *is the neuron state*
τ_i *is the time constant (a high value resulting in a slowly changing neuron state, a low value making the neuron more responsive to current synaptic inputs and no state in the limit)*
ω_{ij} *is the strength of the synapse from the presynaptic neuron j to the postsynaptic neuron I*
$A_j = \sigma(\gamma_j - \theta_j)$*, the activation of the presynaptic neuron, where* $\sigma(x) = 1/(1 + e^{-x})$ *is the standard logistic function and* θ_j *is a bias term*
S is the number of sensory receptors
w_{ik} *is the strength of the synapse from the presynaptic sensory receptor k to the postsynaptic neuron I*
I_k *is the activation of the sensory receptor (range [0, 1]). A forward Euler numerical integration method was used*

Each epoch consisted of 5 trials where the reward zone was kept in a constant position (either the left or the right arm of the maze), and each evaluation consisted of 4 epochs (2 epochs with the reward-zone in each location). The ANN was reset at the beginning of each epoch, but *not* between trials, thereby allowing the robot to *store knowledge* about the location of the reward zone in any particular epoch.

In the first trial of each epoch the robot was given 360 sensorimotor cycles to explore the T-maze, whereas the 4 subsequent trials were only 180 sensorimotor cycles in duration in order to encourage the robot to quickly go to reward zone. In addition a poison zone was placed in the opposite arm to the reward zone which ‘killed’ any robot moving over it, thereby ending the trial. The fitness was determined by the number of trials where an individual was located in the reward zone at the end of the trial. This experiment used a population size of 200 which was evolved for 200 generations and repeated 10 times.

3.11.13.2. Results

It was found that networks explored the T-maze in the first trial but were able to go directly to the reward-zone in the subsequent 4 trials in each epoch. Analysis determined that the state of one of the hidden units determined whether a robot turned left or right: when it was below –0.3 the robot turned left and when it was above –0.1 it turned right. This result was successfully extended to a double T-Maze task where two turns have to be remembered (see Figure 167). It was found that the networks evolved in simulation successfully transferred to real robots (see Figure 168) [Blynel 2003; Blynel & Floreano 2002, 2003].

Figure 168 – Experimental set up for testing transfer of the controller to a real Khepera robot

3.11.14. Evolutionary Hardware

The genetic methods presented above are firmly rooted into control system applications in providing the "logic" of control problems. Increasingly, however, such techniques are being applied to hardware manifestations. Bio-inspired hardware research may be partitioned into three orthogonal directions – phylogeny (P - evolution analogies such as evolvable hardware), ontogeny (O - development analogous such as embryonics) and epigenesis (E - learning analogues such as immunotronics). There are research activities that combine these, eg. evolutionary neural networks (PE). There are areas confined to electronic hardware in using genetic algorithms to evolve field programmable gate array (FPGA) architectures to solve control logic problems. FPGAs with ASICs (application specific integration circuits) are becoming increasingly commonly used for dedicated computations. Evolved hardware controllers are not programmed to follow a sequence of instructions, rather they are configured and allowed to behave in real time according to semiconductor physics. Such controllers have been shown to possess useful properties of robustness and reliability in the face of noise and various kinds of faults. Thompson (1995, 1997) used artificial evolution to design hardware circuits as on-board controllers for two-wheeled autonomous mobile robots displaying simple wall-avoidance behaviour in an empty arena. This work has been applied to Field Programmable Gate Arrays (FPGAs) which are appropriate for evolutionary applications (e.g. the Xilinx XC6200 series) – FPGAs are becoming increasingly commonly used for dedicated computations onboard spacecraft. Switches distributed throughout the chip can be set by an evolutionary algorithm to determine how each component behaves and how it connects to the wires [Thompson 1997]. Lund et al (1997) advocated the use of so-called true evolvable hardware to evolve not only a robot's control circuit, but also its body plan, which might include the types, numbers and positions of the sensors, the body size, the wheel radius, the motor time constants, etc. Evolutionary hardware has also been used for developing self-repairing circuits [Stioca et al 2002]. Lipson and Pollack at Brandeis University [Lipson and Pollack 2000] pushed this idea as far as is reasonably technologically feasible at present in an important piece of research where autonomous 'creatures' are evolved in simulation out of basic building blocks (neurons, plastic bars, actuators). The fittest individuals (defined by their ability to locomote) were then fabricated robotically using rapid manufacturing technology (plastic extrusion 3-D printing). They thus achieved autonomy of design and construction using evolution in a 'limited universe' physical simulation coupled to automatic fabrication. The highly unconventional designs thus realised proved to work well in the real world pointing the way to new ways of developing energy efficient fault tolerant machines (see Figure 169). The chief disadvantage of evolutionary hardware is that it exploits the physics of the electronic circuitry – this can makes it unportable if care is not taken to evolve for robustness to implementation media.

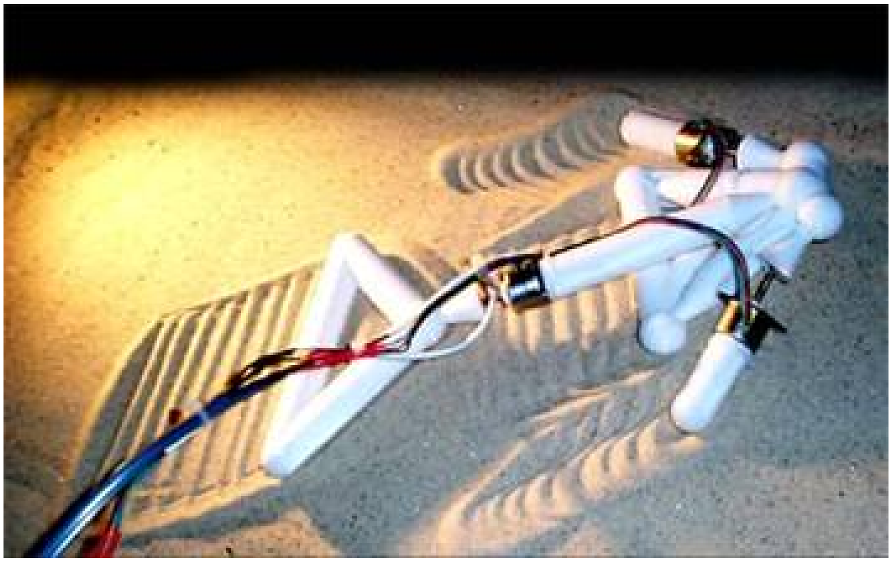

Figure 169 – A Lipson and Pollack evolved creature.

Funes & Pollack (1998) have applied evolutionary algorithms to the design of structures assembled from elementary components (LEGO bricks) in which a simulator evolves stable 3D brick structures such as tables, cranes, bridges and scaffolds within the restrictions of maximum stress torques at each joint between brick pairs. Each brick exerts an external load with a lever arm from its centre of mass to the supporting joint, forming a network of masses and forces within the structure. A genetic program allows tree representations of 3D LEGO structures which may be subject to point mutation of brick parameters and sub-tree cross-over operations at each generation. The structure to be evolved is determined by the fitness function, eg. a bridge is evolved by a Euclidean distance measure. A low-level minor fitness factor for simplicity favouring the fewest bricks prunes many of the useless bricks that inevitably occur. LEGO provides a simple and predictable modular construction tool with simple models of breakage and linkage but is limited in terms of functional robotic structures. Furthermore, this work was based on the evolution of static structures. The opportunities for evolving structure with control systems are limited so the application of biology to mechanical structures is considered in terms of materials and structural characteristics as previously outlined.

3.11.15. Multi-Robot Systems

A particularly challenging area of ER research is that involving multiple interacting robots. The robots typically provide a highly dynamic noisy environment for each other, making the coordination of their behaviours a distinctly non-trivial matter. As well as being interesting for the inherent complexities and challenges, this area of research is highly relevant to space missions where, among other advantages, reliability and robustness can be achieved through the use of multiple robots. Hence we provide a substantial dedicated section discussing and reviewing the evolution of multiple robot systems.

3.11.15.1. Cooperative Systems

The vast majority of multi-robot research addresses the question of how best to design a group of robots that work together to perform some task or achieve some global objective. Such systems are invariably described as cooperative systems, and the robots as engaging in cooperative behaviour. As Cao et al. (1997) have remarked, explicit definitions of cooperation are rare within the multi-robot literature (four such definitions are given in Table 17). Moreover, there appears to be no standard definition in the literature; for this reason, it will be useful to set out some working definitions. For the present purposes then, cooperative behaviour (i.e., cooperation) can be defined as behaviour which is directed towards achieving some common objective, and a cooperative multi-robot system can be defined as one in which the activities of all robots are directed towards achieving a common objective. Cooperative multi-robot systems can thus be contrasted with systems in which robots have been designed to have independent, partially overlapping or conflicting goals. These definitions should not be especially contentious; in particular, the definition of a cooperative multi-robot system is equivalent to Parker's (1998) definition of a 'multi-robot team' (see also Stone and Veloso 1999). Some authors include an additional requirement that cooperation must result in increased performance [Cao et al. 1997; Iocchi et al. 2001]. However, although a desired outcome, increased performance is not obviously intrinsic to cooperation. Thus, for example, if two robots collaborate to push an box from one place to another, then it seems reasonable to call their behaviour cooperative, even if a single robot could push the box just as effectively. Excluding a requirement of improved performance also serves to keep the definitions' inclusion criteria easy to apply. Classification becomes more complex if one must know whether or not performance has been improved, which can be problematic in the absence of comparative experiments.

The design of cooperative multi-robot systems brings with it a set of problems that the designer of a single robot system will not encounter. Whilst there are many circumstances in which it will be simpler to use a single robot, multi-robot systems have several properties that make them attractive for particular types of tasks or circumstances:

- **Robustness through graceful degradation.**

  Any system whose performance degrades gradually—as opposed to catastrophically—with the loss of individual parts can be said to exhibit graceful degradation. A multi-robot system can be designed so that it exhibits graceful degradation, remaining functional in the face of the loss or failure of one or more robots. This is a very useful property of a system if, for example, the task environment is remote or hazardous, or task completion is of critical importance. It is perhaps unnecessary to point out that a single-robot system degrades catastrophically in the face of the loss or failure of an individual robot.

- **Action in parallel.**

  A multi-robot system may be preferred simply in order that a task be completed quickly. Any task that can be divided into a set of sub-tasks that can be performed in parallel typically be completed faster by a group of robots than by a single robot that must perform those sub-tasks sequentially.

- **Distributed sensing and action.**

  One clear advantage of a physically distributed system over a single individual is that group members can be in different places at the same time. Multi-robot systems can offer qualitative advantages where tasks are facilitated by simultaneous, spatially distributed sensing and action. For example, distributed sensing can greatly facilitate tasks that require estimating the position or trajectory of an object or objects in the environment, or that involve visually tracking an object in a cluttered environment [Parker 2002; Pereira et al. 2003b]. Distributed action is an advantage for any task where it useful for different things to be done in the different places at the same time. Thus, for example, the manipulation or repositioning of large objects can be simplified if multiple robots able to simultaneously apply different forces at different locations [Rus et al. 1995].

- **Cost-effectiveness.**

  Several authors have suggested that a system of simple robots might prove to cheaper to construct than a single, complex robot capable of producing equivalent performance [Dudek et al. 1996; Parker 1996]. This is clearly possible. For example, a single robot might need more expensive sensors to duplicate the capabilities of a distributed system with equipped with cheap sensors. However, there has little or no comparative research into this claim.

Although the field is not yet at the stage where it is producing systems have real-world applications, a number of applications have been suggested for cooperative systems of mobile robots. These include mine clearing, toxic waste clear-up, search-and-rescue missions, material transportation, surveillance and planetary exploration [Iocchi et al. 2001; Parker 1998].

3.11.15.2. Coordination

From the perspective of control, perhaps the most central issue in the design of cooperative multi-robot systems is the problem of coordinating the behaviour of individual robots in order to achieve effective, coherent system-level behaviour. Although the concept of coordination is a central one in multi-robot research, definitions of coordination in the multi-robot literature are even rarer than definitions of cooperation (a selection are shown in Table 18). Arguably, the meaning of the word ‘coordination’ is straightforward, and follows everyday usage. Nevertheless, it will be useful to proceed with a working definition. The following definition, taken from organisation theory, and due to Malone and Crowston (1994), seems particularly appropriate in the present context:

“Coordination is managing dependencies between activities.”

Problems requiring coordination thus arise when robots are engaged in activities that are interdependent. It is for this reason that coordination is such a basic and pervasive issue in multi-robot system research: dependencies between robots’ activities will exist, at least to some degree, in all multi-robot systems. Robots activities can be interdependent in many different ways. A robot’s activity may be dependent upon other robots doing the same thing at the same time, for example when pushing a large object [Kube and Zhang 1994], or dependent upon another robot doing something different at the same time, for example one robot may need to push the object when the other steers [Brown and Jennings 1995]. The activities may also be order-dependent, as in Ijspeert et al.’s (2001) ‘stick-pulling’ experiment, in which one robot must raise a long stick partway out of a hole in order for a second robot to be able to lift it the rest of the way out. Alternatively, robots activities may be dependent on the relative positions or trajectories of others, for example, if the robots are required aggregate, disperse or move as a group [Matarić 1992]. Perhaps the most basic of dependencies for a group of mobile robots arises because two robots cannot simultaneously occupy the same physical space. In any physical system, each robot’s location or trajectory will place constraints on the set of possible locations or trajectories available to others - an issue which has been explicitly addressed in many multi-robot experiments [Balch 1999; Goldberg and Matarić 1997; Vaughan et al. 2000].

3.11.16. Cooperative Tasks

As noted above, multi-robot research has yet to progress to stage where it is producing systems with real-world applications. To date, multi-robot research has instead been conducted using simplified tasks, typically conducted in simple, relatively predictable and controlled environments. Typically, these tasks are either designed to be analogous to certain kinds of real-world applications, or useful for investigating types of cooperative strategies and behaviours that would be of general use in real-world applications. The main purpose of this subsection is to give the reader a brief overview of the types of task that have been investigated and implemented on real robots. Although many different cooperative tasks have been implemented, most are variations on a relatively small number of core task types, principally:

- Foraging
- Group movement
- Cooperative manipulation
- Robot football

3.11.16.1. Foraging

‘Foraging’ is the name given to those tasks in which robots seek out and collect objects. Foraging tasks are of interest because they present a scenario analogous to that which would be encountered with such potential multi-robot applications as mine clearing, toxic waste clear-up or search-and-rescue missions [Iocchi et al. 2001; Parker 1998]. In a typical foraging task, robots collect items that are scattered around an environment by carrying (or pushing) them to another location; a single robot is capable of transporting a single object (otherwise the task will typically be classified as cooperative manipulation, see below). Typically, the items the robots must collect are homogeneous and must be taken to a single location [Balch and Arkin 1994; Beckers et al. 1994; Goldberg and Matarić 1997; Melhuish 1999; Parker 1998]. Some authors have investigated the use objects of different types (specifically, of different colours). For example, Balch (1999) implements a variety a different foraging task in which objects of different types must be collected and taken to different locations depending on their type. Holland and Melhuish implement a task in which robots must collect two types of objects into a single grouping, but within that grouping, the objects must be sorted by type [Holland and Melhuish 1999; Melhuish et al. 1998].

Foraging tasks rarely require close coordination, as robots are primarily engaged in individually performing sub-tasks which relatively independent. This is not to say that such tasks cannot present interesting coordination issues. In a typical foraging task, robots will frequently visit the same location in order to deposit their items, and congestion problems can be significant. A number of authors have used foraging tasks to investigate strategies for coordinating robots effectively under such constraints [Balch 1999; Goldberg and Matarić 1997; Fontan and Matarić 1998].

3.11.16.2. Group Movement

Group movement tasks, as the name suggests, require robots to move together as a group whilst remaining within relative close proximity. If no specific formation is required of the moving group, then the task (and the resulting behaviour) is commonly called ‘flocking’. In contrast, a formation movement task requires robots to adopt and maintain a specific formation whilst moving (for example, a column, line, wedge, or diamond formation). A group movement task may also entail that initially dispersed robots aggregate and (if required) form themselves into the required formation. Whilst moving as a group, the robots may also be required to negotiate obstacles or walls. Several authors have implemented multi-robot systems which flock, including Matarić (1995) and Kelly and Keating (1996); the latter with groups of 7 robots and the former with groups of up to 13 robots. Interestingly, Melhuish and Welsby (2002) have recently demonstrated flocking in three dimensions using robotic dirigibles (or 'blimps'). Much of research the research into formation movement has been analytical [Axelsson et al. 2003; Egerstedt and Hu 2001] or investigated in simulation [Fredslund and Matarić 2002]. However, several authors have implemented formation movement with small groups of real robots. Most commonly, robots have been required to move in column or 'convoy' formation [Dudek et al. 1995; Pereira et al. 2003a], but are a variety of other formations have also been demonstrated [Balch and Arkin 1998; Spletzer et al. 2001].

Although group movement is not obviously analogous to any specific real-world multi-robot application, it is not difficult to see how the ability to aggregate and move as a group will be useful in many situations. The specific benefits of adopting a particular formation while moving may be less clear. Some structure in a group may be necessary because of only one robot is capable of navigation, and hence must lead [Dudek et al. 1995]. A particular formation may be necessary because of sensory constraints, for example, where robots much keep each other within a limited field of view. For similar reasons, particular formations may be more appropriate in cluttered environments. In specific applications, such as mine-clearing, there would be obvious benefits to moving in convoy, following the same path as the robot ahead. Formation movement can also be of interest because - as will become evident shortly - it presents many of the same constraints that are found in certain cooperative manipulation tasks, such as where robots must move whilst carrying an object.

Group movement tasks can require the actions of participating robots to be closely coordinated, since the appropriate trajectory or location of each robot is determined at all times by the current trajectories and locations of other robots. However, the degree of coordination required to ensure that the robots remain appropriately positioned may vary considerably between implementations, depending on factors such as sensory constraints, the existence of obstacles, and the geometry of any proscribed formation.

3.11.16.3. Cooperative Manipulation

Cooperative manipulation tasks are tasks that require two or more robots to transport, reposition or otherwise manipulate an object, and many examples of such tasks are to be found in the literature. Manipulation can take a variety of forms. The most commonly implemented form is pushing. Thus, a number of authors have implemented 'box-pushing' tasks, in which robots locate and then push a box-like object [Kube and Zhang 1994; Matarić et al. 1995; Kube and Bonabeau 1999]. A variation on box-pushing is 'caging', in which the robots surround the object and transportation is achieved by the robots maintaining relative positions whilst moving [Pereira et al. 2002a; Spletzer et al. 2001]. In some cases, the robots' goal is not to transport an object, but to reposition it, as in the task implemented by Rus et al. (1995), in which two robots were required to rotate a sofa, each pushing at opposite ends. Another form of manipulation involves carrying an object. Carrying an object will often require closer coordination of the robots actions, since an inappropriate move by one robot can cause the object to be dropped. The movements of pushing robots are clearly not so constrained. A number of authors have investigated the control of small groups of robots in the context of the constraints imposed by navigating and avoiding obstacles whilst carrying an object [Chaimowicz et al. 2001; Pereira et al. 2002b]. Although most cooperative manipulation tasks involve pushing or carrying objects, there are exceptions. As mentioned above, Ijspeert et al. (2001) implemented a task in which one robot uses a gripper to pull a stick partway out of hole, in order for another robot to be able to pull it the remainder of the way out. Donald et al. (2000) have demonstrated a particularly novel form of cooperative manipulation, in which robots use ropes to rotate and to move objects.

The ability to cooperatively manipulate objects is clearly one which has general application to many potential real-world applications of multi-robot systems. The degree coordination required in cooperative manipulation tasks varies. As should be evident, cooperative carrying and 'caging' impose similar constraints to formation movement. In cooperative carrying tasks, the constraints imposed on the formation can potentially be very strict, such that slight deviations from formation result in the object being dropped. With caging, the formation constraints may be more relaxed, such that only approximate relation positions must be maintained. Box-pushing tasks may simply require that (most) robots push the same side of the box, and all robots may not need to push at the same time [Kube and Zhang 1994].

3.11.16.4. Robot Football

Robot football (a.k.a. robot soccer) is a task-area which has received growing attention from the multi-robot community in recent years. As the name suggests, robot football is a game in which teams of robots compete to move, kick or push a ball into the other team's goal. There currently exist several organised leagues in which robot teams are able to compete [Birk et al. 2002; Kaminka et al. 2003]. Although not directly analogous to any proposed real-world application, robot football provides researchers the opportunity to design and test cooperative systems which must achieve their goal in a dynamic and truly unpredictable environments. The degree of coordination exhibited by robot football team can vary considerably, depending on the nature of the team strategies imposed by the teams' designers.

3.11.17. Homogeneity and Heterogeneity

Multi-robot systems can be usefully classified as being either homogeneous or heterogeneous. Conventionally, a homogeneous system is defined as one in which all robots have the same physical characteristics (e.g., sensors, effectors, body-plan etc.), and each is equipped with an identical controller; a heterogeneous system, by exclusion, can be defined as a system which is not homogeneous [Iocchi et al. 2001; Stone and Veloso 2000]. It also common to distinguish homogeneity and heterogeneity independently at the level of control and at the level of physically attributes. This is useful since, for example, a system which is physically homogeneous may nonetheless contain robots with different controllers. The distinction between homogeneous and heterogeneous systems is an important one from the perspective of evolutionary design, since it has a major impact on the choice of method used to generate and test systems [Quinn 2004].

3.11.17.1. The Benefits of Heterogeneity

There are a number of reasons for employing heterogeneous systems. Let us start with reasons for employing *physically* heterogeneous systems. Firstly, it will not always be practical to equip all robots with same set of sensors and effectors. If the robots are subject to size or weight constraints, it may not be possible to equip each robot with all the sensorimotor capabilities necessary for the task. In such situations, physical heterogeneity is a necessity, and the required sensors and effectors will need to be differentially distributed between the robots [Balch and Parker 2000]. Even where physical heterogeneity is not a necessity, it will often afford economic benefits. Many multi-robot tasks can usefully be decomposed into sets of different sub-tasks or roles. If a system can be organised so that robots specialise in different types of sub-tasks or role, then each robot may only need to be equipped with a subset of the sensorimotor capabilities required to complete all sub-task, or to perform all roles. In contrast, in a homogeneous system each robot would be equipped with all the sensorimotor capabilities necessary to perform all roles and sub-tasks. By reducing a system's sensor and effector redundancy, physical heterogeneity can provide a more cost-effective solution than a homogeneous alternative.

Irrespective of whether or not a system is physically homogeneous, there are advantages to robots having non-identical controllers. Unlike homogeneous systems, a heterogeneous system can rely on differences between individual control strategies to simplify the problem of coordination at the level of robot control. In many respects, these are also benefits of specialisation, but of specialisation at the level of control. Insofar as robots have particular functions, roles or areas of responsibility which are .built into - or presupposed by - their control strategies, the problem of action selection can be greatly simplified (Parker, 1996). Thus, where a homogeneous robot will face the problem of deciding which of many different sub-tasks to perform, a heterogeneous robot specialising only in certain subtasks will face a greatly simplified choice - indeed, it may be faced with no decision at all. This is particularly advantageous where the utility of a robot's actions is dependent upon the actions that others are performing at the same time. For example, in the two-robot convoy movement task implemented by Dudek et al. (1995), one robot is programmed to navigate, the other robot to follow its partner. Were the robots homogeneous, each robot would require some coordinating mechanism for deciding whether to navigate or to follow under the constraint that each robot must take the opposite decision. Dudek et al.'s robots require neither of these things. In addition to helping to reduce the need for coordinated decision-making, intrinsic differences between robots can also be used to resolve such decisions faster and more efficiently. Thus, for example, Matarić (1997) uses a pre-defined group hierarchy to determine which robot should give way to which in the context of a foraging task. Here, intrinsic differences facilitate an effective and unambiguous resolution of potential deadlock situations.

3.11.17.2. The Benefits of Homogeneity

One of the principal advantages of homogeneous system is its potential for robustness. A robust multi-robot system is one that exhibits graceful degradation in the face of the loss or failure of individual robots. In a homogeneous system, each robot is equally capable of performing any sub-task or adopting any role. Homogeneous systems are thus potentially more able to cope for the loss of individual members than heterogeneous systems are: The less that the capabilities of individual robots overlap, the more damaging the loss of an individual is likely to be [Parker 1996], and the more individuals that can be lost before the system fails completely [Ratnieks and Anderson 2000]. In the most extreme case, a system in which each robot has some unique, irreplaceable capability or role will exhibit catastrophic degradation. The loss of any individual will leave the remaining robots unable to complete their task. It should be noted that homogeneity does not, in itself, guarantee graceful degradation. To this end, a system may also require specific mechanisms or strategies to detect and compensate for the loss or failure of a robot. For example, compensating for the loss of an individual may require a reconfiguration of the system, in which roles are reallocated and sub-tasks are redistributed [Parker 1996].

The absence of specialisation affords other advantages, since any robot can be deployed, or redeployed, to any sub-task or role, to suit the current demands of the task. This has a number of advantages. Firstly, the fact that each robot can perform any subtask it encounters may result in a task being more quickly. For example, consider a robot which encounters an object, and requires the help of a second robot to move it. On average, if the system is homogeneous, an appropriate helper will have half the distance to travel than it would were the task to require two specialists, say a pusher and steerer, or a leader and a follower. Secondly, robots may be able to perform more effectively because they can exchange roles rather than having to exchange positions. Chaimowicz et al. (2001) demonstrate the usefulness of exchanging roles in a task in which two robots use a leader-follower strategy to carry an object, the follower intermittently takes over as leader in order that pair can reverse direction and more effectively negotiate obstacles.

Farinelli et al. (2002) present a robot football scenario in which the goalkeeper, upon taking possession of the ball, can move up-field whilst another robot takes over the role of goalkeeper. Finally, whereas a system of specialists represents a specific configuration of different capabilities and roles, a homogeneous system is potentially fully reconfigurable. The organisational composition of the system can be dynamically altered in line with changing environmental conditions, since the number, distribution and type of different individual roles can be changed. Examples of systems of this kind are primarily to found in the robot football literature, in which systems are capable of adopting and switching between a variety of different team formations or team strategies [Stone and Veloso 1999; Spaan and Groen 2002].

It may be worth pointing that the advantages listed above follow from a lack of specialisation, rather than from homogeneity *per se*. The same advantages can be enjoyed by heterogeneous systems to the extent that each robot's capabilities overlap. For example, the role-swapping robots used by Chaimowicz et al. (2001) in the example above, were actually physically heterogeneous robots. However, since each was capable both of following and of leading, an exchange of roles was as possible as it would be in a homogeneous system.

3.11.18. The Distribution of Control

Table 17 – The Four Main Types of Control Architecture Organisation.

Architecture	Distribution of Control
Centralised	A single agent (computer or robot) is responsible for controlling all robots.
Distributed	All robots are autonomous and system control is distributed (more or less equally) between robots.
Hierarchical	System control is globally decentralized, but locally centralised through a ”chain of command”.
Hybrid	Robots have a degree of local autonomy, but higher-level system coordination is achieved through centralised or hierarchical control.

With respect to control, a multi-robot system can be organised in a variety of ways. Four main types of organisation are typically distinguished, these are: *centralised*, *distributed*, *hierarchical* and *hybrid* control systems [Cao et al. 1997]. If control is centralised, a single agent is responsible for controlling all robots. This agent may be one of the robots or, as is more commonly the case, an off-board computer. Any system in which there is no single controlling agent is classed as a decentralised control system. One form of decentralised control is distributed control, in which all robots are autonomous and hence system control is distributed (more or less equally) between robots. Distributed control has been the dominant control paradigm in multi-robot systems research—at least with respect to physically implemented systems. Another, less common, form of decentralised control is hierarchical control, in which system control is globally decentralised but locally centralised through a ”chain of command”. Finally, there are also hybrid centralised/decentralised control architectures, in which robots are accorded a degree of local autonomy, but higher-level system coordination is achieved through centralised or hierarchical control. For the present purposes, it will be sufficient to discuss the advantages and disadvantages of a distributed control relative to those of centralised control. Distributed and centralised control are two most common architectural paradigms in multi-robot system research. In addition, each of the four architectural types above can be distinguished by the presence or absence of centralised control structures, (whether these are local or global structures, augmented by partial autonomy).

Since distributed and centralised control are two most common architectural paradigms, the relative merits of each have been discussed by a number of authors. In such discussions, the relative merits of distributed and centralised control are often treated as if equivalent to those of local and global control (see below). Thus, it is commonly stated that one of the main advantages of centralised control is the availability of global information and the centralisation of decision making [Cao et al. 1997; Lumelsky and Harinarayan 1997]. A centralised controller may utilise all of the sensory information available to the system, and can use this information to formulate global plans and ensure efficient and effective coordination of the system. However, global planning, global coordination and global access to information are all possible with distributed control systems given sufficient communication bandwidth and computational power. Thus, it seems difficult to argue that this is an inherent advantage of centralised control. However, where such global strategies are employed, a centralised architecture does offers some benefits. Primarily, the centralised approach will tend to be the more minimal control architecture. Firstly, it requires only a single complex controller, in contrast, the equivalent distributed system would require many complex controllers (i.e. one for each robot). Secondly, global strategies run in parallel on a distributed system may require addition mechanisms to ensure synchronicity. In addition, if the centralised controller is an off-board computer, it may well have more computing power than an on-board robot controller.

The main advantage of a distributed control system lies in its potential for robustness and fault tolerance. Firstly, centralised systems are completely vulnerable to the loss or failure of their centralised controller; a single fault has the potential to render the entire system useless. In contrast, the equal distribution of control in a distributed system gives it the potential to be robust to the loss or failure of any part of the system. Secondly, centralised systems are vulnerable due to their reliance on regular ‘ethereal’ communication between the centralised controller and the robots. A robot that ceases to be able to communicate with its controller will cease to be able to make any coherent contribution to the performance of system’s task. Any global disruption of communication the system’s communication capabilities will have severe consequences. In contrast, in a distributed system, communication failures do not necessarily prevent a robot from operating coherently. Deprived of global information, an autonomous robot can still operate using local information. Thus, where distributed systems do exploit ethereal communication, they have the potential be more robust to local or global communication failures than a centralised system. Other advantages commonly attributed to distributed architectures are more properly attributable to local control; it is these that will be addressed next.

3.11.19. <u>Local and Global Control</u>

In addition to distinguishing architectures based on the distribution of control, it is also common to distinguish them by the extent to a robot has access local or global information. Parker (1993) characterises local control as follows:

> Local control laws [...] guide an agent’s actions based on the proximate environment of that agent. Such information is derived from the agent’s sensory capabilities, and thus reflects the state of the agent of the world near the agent.

Global control laws, by exclusion, guide an agent’s actions based on what lies beyond the agent’s proximate environment. For example, a global control architecture may exploit information provided by other robots via radio/Ethernet communication (e.g. Østergaard et al., 2002; Chaimowicz et al., 2001), or provided by off-board sensors, such as an overhead camera or by external locational systems, such as global positioning systems (GPS) or radio beacon triangulation [Balch and Arkin 1998; Goldberg and Matarić 1997]. This may additionally be supplemented non-local information stored in internal maps or world models that are either provided in advance, or dynamically constructed by the controller.

Global control elements are a necessary requirement of a centralised control architecture, they are not in a distributed control architecture. Nevertheless, access to global information and global knowledge clearly does have the potential to improve system coordination and performance. What are the benefits of restricting a system to local sensing and control? There are a number of reasons:

- **Implementing local control strategies can help to reduce the total amount of equipment required by the system.**

  Equipping each robot with, for example, radio or GPS capabilities may not be an attractive option if unit cost, power consumption, size or weight are important considerations, and most of these will invariably be important considerations in space applications. In such cases, it may be preferable to do without performance gains due global control capabilities unless they are crucial to the completion of the task. However, local control does not necessarily guarantee a minimal configuration of equipment, since there may be trade-offs between local and global capabilities: for example, robots equipped with radio communication may require fewer, or less sophisticated, local sensors.

- **Local control has greater potential for robustness and fault-tolerance.**

  As discussed above, an important benefit of a distributed control architecture is its potential to be robust to the loss or failure of any part of the system, and hence to degrade gracefully. This potential for robustness can be exploited by minimising the global impact of local failure. In this respect, global control strategies can reduce robustness and fault-tolerance in two ways. Firstly, by relying on externally provided information. If robots' control strategies rely on information transmitted by other robots, the loss of an individual robot or the failure of its communication equipment can significant system-wide consequences. Additionally, the effects of a local malfunction (such a failed sensor) can be greatly magnified if this leads to the affected robot to transmit inaccurate information. Similarly, the potential for local failures to a catastrophic impact on system performance increases if all robots are reliant upon the same, single external piece of equipment, such as a communications server, GPS., or radio beacons. Secondly, robustness and fault-tolerance can be reduced by relying on stored environmental information, such as global models and maps. Such strategies are vulnerable both to error and to out-of-date information. Both of which are more likely in a fast-changing environment. In general, a strategy based on responding to the immediately perceivable environment should be more robust to unexpected changes in the environment [Brooks 1992; Werger 1999].

- **Local control can minimise a robot's computational requirements, as a controller's task is restricted to responding to its local environment.**

  "Processing [...] global information", as Parker (1993) points out, "requires time and resources, both of which are limited in real-world applications". Controllers in locally controlled systems will be able to react faster and will require less computational resources than systems utilising global strategies. The computational complexity of the task facing a controller with access to global information can be significantly greater than for one responding only to local sensory input. The situation worsens with scale. As Matarić, (1997) points out, given a group of *n* robots, each of which may be in one of *s* states, the global state space is of size *sn*, whereas, the local state space is simple *s*, irrespective of the size of the system. Computational requirements will be further increased if a controller must also build (or maintain) a model of its environment, and significantly so if the information is used for planning and thus requires a branching search algorithm. Thus, local control will tend to minimise the computational requirements of an individual controller. Moreover, these requirements are essentially independent of the size of the system, global control, in contrast, scales very badly.

Of course, local control also has it drawbacks. Local control will rarely be sufficient to produce optimal solutions to tasks [Parker 1993], indeed some tasks will be very difficult, if not impossible, to perform without some element of global control. In addition, effective local control strategies can be very difficult to design. Where a system is reliant upon local, distributed control, coherent system-level behaviours must be coordinated through the local interactions of individual robots [Matarić 1992]. As many authors have noted, the design of local control laws which combine to yield "emergent" global behaviour can be both conceptually and analytically challenging [Matarić 1992; Kube and Bonabeau 1999].

3.11.20. Why Evolve Cooperative Multi-robot Systems?

Arguments in favour of evolutionary design of robotic systems were given earlier in this document. Recall that the general argument proceeded from the potential complexity of the task of hand-designing robot control systems. In particular:

- Autonomous robots are closely and recursively coupled to an environment, making the dynamics of the robot-environment interaction difficult to predict. However, an understanding of this dynamic is integral to most stages of the development of a controller by conventional design processes.
- Conventional design approaches proceed by decomposing complex problems into a set of independent modules, each of lesser complexity. However, the component parts of a robot controller can exhibit a significant level of interdependence, influencing one another either directly, or indirectly as a result of interactions with a shared environment. Such interdependence makes effective decomposition problematic.

In the light of the above, it was argued that artificial evolution is a potentially useful design tool for two main reasons:

- The evolutionary design process is not dependent upon the predictability of system dynamics, not does it require prior decomposition of complex design problems into a set of more manageable sub-problems. The designer's role is essentially reduced to providing a suitable means by which to evaluate potential solutions, and to specifying the set of building blocks from which those solutions will be constructed.
- Since an evolutionary approach is not constrained to finding solutions that are easily decomposed, or even easily understood, it offers the possibility of exploring regions of a design space that conventional design approaches would typically be constrained to ignore.

The same arguments do, of course, apply to the design of control systems for multi-robot systems. However, it is possible to make an even stronger case for evolutionary design in a multi-robot context, particularly in the case of specific kinds of multi-robot systems that are likely to be interest to the space industry. The potential benefits of evolutionary design seem most at their greatest in the context of the design of multi-robot systems with local, distributed control, where robots have minimal, imprecise and noisy effectors and sensors. As discussed above, systems of this kind afford a number of potential benefits in the context of space robotics. These include robustness to individual robot failure, fault tolerance, scalability, and the potential for minimising power consumption, size and weight. Whilst offering these benefits, systems can present a significant challenge to the conventional approaches to design. In particular, system dynamics can be extremely difficult to predict, and the design problem can be extremely resistant to decomposition.

Where a system is reliant upon local, distributed control, the main challenge facing a designer is to ensure that coordinated and coherent system-level behaviours arise as a result of the local interactions of individual robots [Matarić 1992]. This is challenging because of the potential difficulties in predicting the 'emergent' global behaviour of a locally controlled system, and secondly, because the behaviour of each individual must - at least to some extent - be considered in the context of its effects on the system as whole. The more closely coupled a system of robots, the harder it will be to consider the behaviour of each robot in isolation from the behaviour of the system as a whole (making decompositional design problematic), and the harder it will become to predict global system behaviour as a function of local control rules. Ensuring effective coordination of a closely coupled system entails considering the actions of each robot in the context of the potential effects they will have on the other robots, and hence their consequences for the behaviour of the system as a whole. The consequences of a robot's action will often be difficult to predict. Any given may trigger a chain of reactions (e.g. A affects B, B affects C), and this chain may be recursive, either directly (A affects B, B affects A) or indirectly (A affects B, B affects C, C affects A). Since each robot is acting concurrently, it will often be necessary to consider multiple, intertwined chains of cause and effect operating in parallel—a task made harder by the existence of sensory and motor noise, and other sources of environmental stochasticity. A designer's task is further complicated if sensors provide only limited and ambiguous information about the robots' environment. With limited sensing capabilities, current input will rarely be sufficient to disambiguate environmental state since a range of potentially very different environmental conditions may be commensurate with the same set of sensory input. It is thus be necessary consider how a robot should respond to current sensory input given that the environment may be in one of the number of different states. This means predicting the impact of the robots' response on global system behaviour in each of those possible environmental states in order to ensure that the response will be appropriate to all of them.

The previous paragraph raised the issue of the difficulty of designing individual robot behaviour in isolation from global system-level behaviour. However, the existence of inter-robot dependencies also raises problems for the decompositional design of individual control systems. Recall that conventional approaches to controller design are modular. With respect to local control, the most commonly used architecture is some form of behaviour-based architecture [Balch and Arkin 1994; Brooks 1991; Matarić 1997]. On this approach, the controller is decomposed into set of (relatively) independent behaviour modules, each equivalent to some basic behavioural competency (e.g., collision avoidance, wandering, following another robot, photo-taxis etc.). These modules operate in parallel subject to some arbitration mechanism. Clearly, if each component can be designed independently, this can reduce the complexity of the design task. In making the general case for the evolutionary design of (single) robot systems, it was pointed out that such components can rarely be designed (completely) independently of one another. In particular, if the control module is one that is reliant on close coordination of sensing and action, it will be very difficult to design it independently of any other modules that must run in parallel with it. In the context of locally-controlled multi-robot systems the problems of control-module interdependence are exacerbated since a designer must consider additional levels of potential interdependence. If a control module is one which has any influence on a robot's interaction with another robot (or is influenced by such an interaction), then the design of that module will be not be independent the design of other robots' control systems, nor of the system as whole. More generally, any decompositional approach to controller design will face the problems of interdependence on at least three levels:

- inter-module interdependence, within a single controller.
- interdependence between a module and the controllers of other robots.
- interdependence between a module and the overall system behaviour.

The design of systems with local, distributed control, comprising robots with minimal effectors and sensors can present a significantly complex task. The behaviour of such systems can be difficult to predict, since they can exhibit complex dynamics. In addition, systems of this type have the potential to be extremely resistant to decompositional design. Independencies exist at a number of levels, with system-level behaviour resulting from both inter- and intra-robot interactions. For this reason, artificial evolution has the potential be an extremely useful design tool for such systems. It should not, however, be assumed that the benefits of artificial evolution of limited only to these types of systems. Systems that incorporate global control elements or accurate sensors can still produce dynamics that are difficult to predict, and be subject to a significant degree of interdependency.

3.11.21. Review of Evolved Systems

There are still relatively few examples of evolved multi-robots systems in the Evolutionary Robotics literature. The number of examples of evolved cooperative multi-robots systems is even smaller. For this reason, the review that follows does not focus simply on cooperative multi-robots systems, but its scope is expanded in two directions. Firstly, the review encompasses simulated cooperative systems. Secondly, non-cooperative systems are included in review of real (i.e., physically instantiated) multi-robot systems.

3.11.21.1. *Simulated Cooperative Systems*

Many researchers have used artificial evolution to the design of control systems for cooperative systems of ‘simulated robots’. However, it is important to point out that in such cases the word ‘simulation’ is typically used extremely broadly. The majority of such work with has utilised simple, abstract simulations, in which no particular robot is modelled and the problems of sensorimotor coordination and control are reduced to minimum [Agah and Bekey 1996; Barfoot and D’Eleuterio 1999; Bongard 2000b; Collins et al. 2000; Luke and Spector 1996; Unemi and Nagayoshi 1997; Wu et al. 1999]. This work is not uninteresting, but it is of limited relevance in the current context. The number of researchers who have evolved cooperative multi-robot systems using realistic simulations of physical robots is surprising small; the list that follows appears to be exhaustive. In each case, researchers were concerned with the design of a distributed control system and - with one exception - with the design of homogeneous systems.

The most common task type addressed has been group movement (i.e., variations on ‘flocking’), accounting for five of the six cases listed.

The first example is due to Martinoli (1999), who evolved controllers for groups of 3, 5 and 10 simulated Khepera mini-robots. (The wheeled Khepera mini-robot has been introduced in earlier sections of this report). The robots were required to perform a foraging task in a walled arena, collecting items that were represented by black circles on the arena floor, a robot collected an item by driving over it, whereupon the item disappeared from the arena. The robots were physically homogeneous and equipped with active infra-red (IR) proximity sensors, which allowed them to sense walls and other robots, two downward-facing ‘floor sensors’ with which to sense the collectable items, and two motor driven wheels. In the context of a comparative experiment, both homogeneous and heterogeneous (control) systems were successfully evolved.

Quinn (2000, 2001, 2004) evolved controllers capable of performing a group movement task. This was undertaken for groups of two and for groups of three simulated Khepera mini-robots. Each robot was equipped with 8 IR sensors only, giving it a very limited sensor range and capabilities. The specific nature of the task was as follows: The robots were initially placed at random positions and orientations under the constraint that they were within sensor-range of one another. The robots were then required to move a certain distance from their starting point, whilst remaining within sensor range one another and avoiding collisions. Quinn sought to evolve homogenous control solutions to this task, as part of a comparison of two different methods for generation of robot groups. Successful controllers were evolved using both methods for both two and three-robot groups. The task was solved by robots first moving into - and then travelling in – a specific formation. Interestingly, solutions to these tasks also entailed different robots adopting different roles (for example, leader and follower in the two-robot case). Since the system was homogenous, roles had assigned dynamically, with the robots having to 'decide' through interaction which robot should perform which role.

Baldassarre et al. (2002, 2003b) implemented another group movement task using simulated Khepera mini-robots. Their system was also homogeneous, comprising four robots, each equipped with 8 infra-red and 8 ambient light sensors, four directional microphones, a speaker emitting continuous sound (not under control of the robot), and two motor driven wheels. The robots were initially scattered in a rectangular walled arena that contained a light source placed at one end. Their task was to aggregate (using phonotaxis) and then move as a group towards the light source. Baldassarre et al. evolved a number of different systems capable of performing this task, with different systems flocking in different formations, some static (with respect to relative robot positions) and some dynamic. In the case of some of the systems that they evolved, they observed that different robots made functionally distinct contributions either to group formation, or to group movement, in this respect, then, they identify different roles. It is interesting to note that since robots were homogeneous, and the robot's control systems were reactive, these differences in behaviour were a function of differences in external state only.

The three remaining examples have come out of the recent multi-institutional SWARMBOTS project. In each case, the robots simulated were 's-bots', a model recently developed for use in the SWARMBOTS project. The simulation used was constructed using a rigid body dynamics simulator package, yielding relatively accurate physics. One feature of these wheeled robots is that they have the capacity of forming physical links with other robots. Both of the experiments described below address the problem of group movement by a system of physically linked robots.

The first example is due Baldassarre and colleagues [Baldassarre et al. 2003a]. Baldassarre et al. consider a task in which the robots are already physically linked, and must move as a group in an obstacle-free environment. Each robot has a turret that can rotate independently of the robot's chassis, and the links between the robots are attached to these turrets. In this experiment, the robots' only sensory input comes from traction sensors that detect the direction and intensity of the force that the turret exerts on the chassis, the robots' controllers must then provide output to two motor-driven wheels. Using extremely simple feed-forward neural network controllers, they evolved a control system using four robots linked in a linear formation, and showed that the solution was robust to initial conditions (i.e., different combinations of turret/ chassis orientations). Interestingly, they were able to show that the solution generalised well to different groups size (from 3 to 10 robots), and also to a variety of formations other than the linear formation for which they had evolved the controller. Although evolved in an empty arena, the linked groups of robots were also able to negotiate a walled arena containing obstacles. The force exerted by a robot impeded by a wall or obstacle was sufficient to cause the linked group to change its direction of travel and take successful avoidance action.

The second example is due to Trianni and colleagues [Trianni 2003; Trianni, Nolfi and Dorigo 2004] who re-implemented and then extended the linked group movement experiment due to Baldassarre et al. (2003a). His extension entailed placing the robots in a walled arena, the floor of which contained a four square holes. The robots were additionally equipped four downward-facing proximity sensors; these enabled the robots to sense the presence of holes (prior to falling into them). The proximity sensors can also detect the presence of walls, since the walls were sloped. Collisions with walls could also be detected, albeit indirectly, with the robots' traction sensors. They were able to evolve controllers for a homogeneous system of four robots, linked in a linear formation, which were capable of navigating within the arena with a reasonable degree of success. To some extent, the control solutions generalised to other formations, but not with same degree of success observed with the original experiments in empty environments.

3.11.21.2. Non-Cooperative Systems

At least four separate examples of evolved non-cooperative multi-robot systems can be found in the literature. The first of these is system evolved by Nolfi and Floreano (1998); this was an explicitly competitive system comprising two robots with fully conflicting goals. The other three examples are due to Watson, Ficici, and Pollack [Ficici et al. 1999; Watson et al. 1999], to Simões and Dimond (1999, 2001), and to Usui and Arita (2003). These three sets of authors evolved systems of robots, each which pursued independent goals. Although these goals may partially overlap – as they do in the case of Simões' experiments – the robots lack a common, global objective. For this reason, these are not cooperative systems under the definition adopted in this report. In what follows, I will review each of these four non-cooperative systems in turn.

Nolfi and Floreano's predator-prey system appears to be the first published example of the application of evolutionary techniques to the design of controllers for multiple physical robots. The authors were interested in the co-evolutionary dynamics of competing species, and in particular, in the possibility that competition between co-evolving populations could lead to evolution of increasingly complex competitive strategies as a result of evolutionary 'arms races'. They sought to investigate this by evolving a population of 'predator' robots to compete against members of a population of 'prey' robots in pair-wise contests. After performing a number of experiments in simulation [Floreano and Nolfi 1997a,b], they performed an experiment using real robots [Nolfi and Floreano 1998; Floreano et al. 1998].

The robots used in Nolfi and Floreano's experiments were Khepera mini-robots. The robots' contest took place in a walled arena, with the predator robot's task being to "catch" the prey robot (by making physical contact with it), and the prey robot's task being to evade capture for as long as possible. Each robot also equipped with eight infrared (IR) proximity sensors distributed around its body, each with an effective range of 3-4cm. The predator was additionally equipped with a vision turret (a forwards-facing 64 pixel linear CCD camera with a 36 degree field of view); thus, so long as it was facing in the right direction, the predator was capable of sensing the prey from any location in the arena. However, although the prey's sensing capabilities were more limited, it was capable of a maximum speed twice that of the prey. Each robot was controlled by a simple feed-forward neural network in which there was a weighted connection from each sensor input to each of two motor outputs. The prey's network had eight inputs (one for each IR sensor), the predator's network had an additional five inputs for vision (the CCD camera input was mapped onto five binary 'photo-receptors'). In both cases, a controller's genotype comprised a set of encoded network connection weights. Although the evolutionary algorithm itself ran on an off-board computer, all fitness evaluations were carried out on the physical robots by downloading an individual controller from each species onto each robot. Each pairing was evaluated for 36 seconds, or until the prey was caught. The prey robot timed the evaluation using an on-board clock, and this information was used by the off-board computer to calculate performance scores for each controller.

The three remaining non-cooperative systems are due to Simões and Dimond (1999, 2001), Watson, Ficici and Pollack [Ficici et al. 1999; Watson et al. 1999] and to Usui and Arita (2003). All three use a similar type of approach, which Simões and Dimond term an 'embedded evolutionary system', Watson Ficici and Pollack call 'embodied evolution', and Usui and Arita term 'situated and embodied evolution'. To avoid too much repetition, it will be useful to give a basic outline of this approach before reviewing the individual systems. Whilst each set of authors has adopted a slightly different implementation, the basic method has the following features:

- A group of physical robots operate in a shared environment.
- Certain parameters of each robot's control system are denoted as 'genetic material' and a set of such parameters comprise a genotype. These parameters might be, for example, a set of classifier rules, or a set of connection weights for a neural network controller.
- An evolutionary population of genotypes is distributed across the physical robots. In some implementation there is a single genotype per robot, in others each robot stores a number of genotypes, one of which specifies the robots current controller.
- At the beginning of the experiment, all genotypes are randomly initialised.
- Each robot is capable of monitoring its own performance, and is thus able to assess its current fitness.
- Genetic material is periodically exchanged between robots over communication channels. In such exchanges, fitter robots are more likely to be donors, less fit robots are more likely to be the recipients. Exchanged material may also be subject to mutation.

Watson, Ficici and Pollack's experiments were primarily motivated by the problem of evolving controllers for single-robot systems. They were concerned with the potentially prohibitive amount of time required to test evolving controllers using a physical. They pointed out this could be significantly reduced by evaluating controllers in parallel rather than in series, and set out to demonstrate such a system. To this end, the authors evolved a system of eight custom-built 'Tupperbot' robots [Ficici et al. 1999; Watson et al. 1999]. Each robot was equipped with two light sensors and two motor-driven wheels, and was controlled by feed-forward network comprising two binary inputs, two motor outputs and four connection weights (one connection from each input to each motor output). One input was constantly on (serving as a bias), the other was on if the left sensor activation was higher than the right, and off otherwise. Only the four connection weights were evolved; each weight could take one of 16 integers values (giving a search space of 164 possible network configurations). At the beginning of the experiment, each robot's weights were randomly initialised, and the robots were placed at random locations in a walled arena. In the centre of the arena was a single, omni-directional light source. The robots were required to evolve phototaxis (i.e., the ability to steer towards the central light-source). This was the robots only task. Once a robot had reached the light-source, its evolving controller was temporarily overridden by a hard-wired controller. This controller then steered the robot to a random location elsewhere in the arena from which it could continue phototaxis. Genetic material was exchanged through local infra-red communication: Robots broadcast single connection weight values – 'genes' – which could be received by other robots within a certain proximity; the broadcast genes could be mutated by an offset of 1. Robots broadcast their genes probabilistically, at a rate proportional to their 'energy level', and a robot receiving a broadcast connection weight accepted the weight with a probability inversely proportional to its own energy level. A robot's energy level approximately reflected its ability to perform the task - levels were increased each time a robot reached the light source, but also decayed exponentially over time. Thus, the more regularly a robot reached the light source, the higher its (average) energy level would be.

Watson, Ficici and Pollack were consistently able to evolve robots capable of performing photo-taxis. They report that the evolved solutions were comparable in performance to controllers they had designed themselves.

Simões and Dimond (1999, 2001) employed a similar approach to that described above, this time evolving controllers for robots able to move around in an environment whilst avoiding obstacles. Their system comprised five robots. Each robot was a two-wheel differential-drive rover, circular, around 20cm in diameter, and equipped with 8 collision sensors (bumpers) and 8 peripheral active infrared proximity sensors. Robots were also equipped with AM radios; these were used to exchange genetic material. The robots' task environment was a 4 x 4m walled arena containing within it a number of additional walls and a variety of obstacles. At intervals during the experiment, the obstacles and walls within the arena were reorganized. Each robot was controlled by a neural network controller, which mapped the sixteen sensory inputs onto two motor outputs. The network weights along with several other parameters (such as maximum wheel speed) constituted each robot's genotype. At the beginning of the experiment, the robots' controllers were randomised, and the robots randomly placed in the arena. A robot was able to assess its own performance (i.e., fitness) by monitoring its own wheel speeds (to monitor movement) and bumper sensors (to monitor collisions).

The nature of the embedded evolutionary algorithm used by the Simões and Dimond differed from that employed by Watson, Ficici and Pollack in two main respects. Where the latter's algorithm was asynchronous and local, Simões and Dimond's algorithm was synchronous and global. Robots engaged in periods of self-assessment (i.e., fitness monitoring) interspersed with a periodic, synchronised, global exchange of genetic material. At each 'breeding event', the fittest robot was left unchanged, and the remaining four robots received genetic material from a partner that was chosen at biased random (biased towards the fittest). Using this approach, Simões and Dimond were consistently able to successfully evolve robots capable of moving and avoiding collisions in arena containing obstacles.

The third of the embedded evolutionary algorithm experiments is due to Usui and Arita (2003). Usui and Arita employ a system of six Khepera mini-robots, each utilizing 6 active IR proximity sensors (the two rear-most sensors were not used) and controlling two motor-driven wheels. Robots were controlled by a simple feed-forward network, mapping the 6 inputs onto two motor outputs. A genotype comprised the set of weights for such a network. Each robot was equipped with an infra-red emitter-receiver facilitating a local exchange of genetic material. The robots' task was similar to that used by Simões and Dimond. Placed in a walled obstacle-free arena, each robot must seek to maximise the time it spends moving, whilst minimizing the amount of time it spends in close proximity to walls and to other robots. Each robot was therefore able monitor its own performance by monitoring its wheel speeds and the activation of its proximity sensors.

The embedded evolutionary algorithm used by Usui and Arita differed from both to the two reviewed above. Each robot contained an independent genetic algorithm; that is, a robot contained a number of genotypes, which it was engaged in serially evaluating, instantiating each in turn and then assigning it a fitness score. New genotypes were generated periodically through selection and mutation, and replaced low-fitness genotypes. At the same time, robots transmitted genotypes between themselves. The transmission of genetic material between robots occurred as follows: At intervals, each robot broadcast a single genotype. The broadcasts were asynchronous, but each robot broadcast with the same frequency. For each broadcast, a genotype was chosen at random, biased towards genotypes with the highest fitness. In this way, fitter genotypes tended to migrate between robots.

3.11.21.3. Cooperative Systems

To date there appear to have been only two examples of evolved cooperative multi-robot systems. One is due to Quinn and colleagues [Quinn et al. 2002, 2003], and the other due to Nelson and colleagues [Nelson et al. 2002, 2004]. These are both systems in which robots have been evolved to work together to perform a given task, and where it is system-level performance rather than individual performance that has been optimised. In contrast to experiments described above, in both sets of experiments controllers were evolved in simulation before being successfully transferred onto real robots.

Quinn and colleagues were interested in the evolution of minimally equipped multi-robot systems capable of closely coordinated behaviour. They argued that systems operating under such constraints presented a challenge for more conventional approaches to design, and could therefore potentially benefit from the use of an evolutionary approach. The task they chose was the group movement task that Quinn had already explored in simulation (reviewed above).

Three custom-built robots were used, each built to the same specification. Each robot had a rectangular body, was equipped with two motor-driven foam-rubber wheels which provided locomotion through differential drive (an unpowered castor wheel ensured stability). Each robot was also equipped with four active IR sensors, each comprising a paired emitter and receiver, two facing forwards and two facing to the rear. The limitations imposed by such a minimal set of sensors was exacerbated by the existence of inter-robot sensory interference – in addition to sensing reflected IR (the normal function of an IR sensor), robots also directly sensed IR emitted by other robots. The robots were controlled by artificial neural networks. Each neural network comprised four sensor input nodes (one for each sensor), four motor output nodes (two for each motor), and some number of neurons. These were connected together by some number of directional, excitatory and inhibitory weighted links. The networks had no explicit layers, so any neuron could connect to any other, including itself, and could also connect to any of the sensory or motor nodes. The thresholds, weights and decay parameters, and the size and connectivity of the networks were under evolutionary control. Each group of three robots that was evaluated was specified by a single genotype. The genotype was first decoded to produce a controller, and then ‘cloned’ to produce a set of three identical controllers. This group of three controllers were then evaluated in simulation and the genotype was assigned a fitness score accordingly. To facilitate transfer from simulation to reality, Jakobi’s (1998) minimal simulation methodology was employed.

Quinn and colleagues were able to evolve several different solutions to the group movement task, and to successfully transfer some of these to the real robots. As in the earlier experiments with simulated Khepera robots, the task was solved by robots first moving into - and then travelling in – a specific formation. Quinn and colleagues analysed the best evolved solution, showing that it relied upon each robot adopting a different role, and upon the robots collectively ‘deciding’ which robot would perform each role. The authors claimed that evolved system would have proved extremely difficult to have designed by hand.

The second example of an evolved cooperative system is due to Nelson and colleagues (Nelson et al. 2002, 2004). They sought to evolved cooperative teams or robots to compete in an adversarial game with other teams. A team comprised a pair of robots. The robot platform used was the `EvBot’ - a custom-built multi-robot system platform developed at North Carolina State University (Galeotti et al 2002). An EvBot robot is a rover utilizing twin-tracked locomotion and that can be equipped with a variety of sensor modalities. In this particular experiment, each robot was equipped only with a video camera. The game played by two teams of robots was ‘capture the flag’. The two teams were placed in a walled arena containing additional internal walls (i.e., a maze). Each team was assigned a base. The game was won if a robot from one team captured (i.e., touched) the base of the other team. As in Quinn and colleagues’ experiment, each team was generated by ‘cloning’ a single genotype. Genotypes were evolved in a single population, and during evolution, genotype fitnesses were assessed by pairing genotypes and conducting games between the two encoded teams.

Robots were controlled by artificial neural networks. Similarly to Quinn et al.’s implementation, network size and connectivity were under evolutionary control, as were network weights and neuron thresholds and activation function parameters. The robots’ sensory (video) input was pre-processed before being passed to the neural network. Aided by a colour coded environment, the vision system could distinguish five object types: wall, own/other team robot, own/other goal. For each detectable object type, range and angle values were passed to the neural network, which then specified motor output for the two wheels.

Nelson and colleagues were able to evolve teams that could play the 'capture the flag' game. They found that evolved teams played better than teams of robots with random controllers (i.e., those producing a random sequence of movements). They were, however, able to hand-design a controller that was superior to the evolved controller.

3.11.21.4. Detailed Examples

This section offers detailed descriptions of two of the experiments that were introduced in the previous section. The experiments that will be presented are those due to Baldassarre et al. (2003) and Quinn et al. (2002, 2003). Both provide examples of the use of artificial evolution to design locally controlled, distributed systems of minimally equipped, cooperating robots, capable of producing closely coordinated behaviour.

The goal of Baldassarre's experiment [Baldassarre et al. 2003] was to evolve controllers for a homogeneous system of 4 simulated, linked robots. Starting from random initial orientations, the robots were required travel as far as possible from their initial location, a requirement that entailed the robots negotiating and then maintaining a common direction of travel. The only sensor input available to each robot was from a "traction sensor", this gave each robot an indication of the intensity and direction of the force due to the robot or robots to which they were connected. Baldassarre et al. were able to find solutions to this task using an extremely minimal control architecture: a reactive, feed-forward network with no hidden layers.

3.11.21.4.1. The Robots

The type of robot modelled in Baldassarre et al.'s experiment is the "s-bot", a design under development by the EU-funded *Swarm-bots* project (http://www.swarm-bots.org). The Swarm-bots project is concerned with the construction of a self-assembling and self-organising robot colony made up of a large number (i.e. 30-35) of s-bot robots. Figure 170 shows a graphic visualization of the robot, and a photograph of two prototypes. Each s-bot is a fully autonomous mobile robot with the capacity to sense its local environment, locomote, and grasp objects. Each robot also has the capacity to communicate with others by means of visual and audio communication. Most relevantly in the present context, the robots can become physically joined to one another by means of either rigid or flexible gripping arms. By linking together, a group of robots have the potential to perform tasks that a single robot could not achieve alone, such as transporting heavy objects (for example, by 'caging') and traversing rough terrain (for example, by bridging gaps, as can be seen in Figure 170).

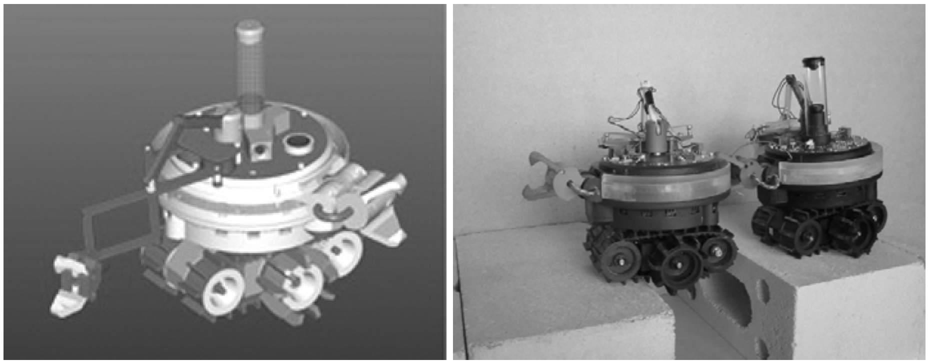

Figure 170 – The "s-bot" robot, under development by the Swarm-bots project Left: A graphic visualisation of the robot. Right: Photograph of two proto-type "s-bot" linked together, traversing a gap (http://www.swarm-bots.org/).

The main features of the s-bot design are as follows:

- Size: main body diameter: 116 mm; height: 100 mm.
- All-terrain mobility; differential drive locomotion combining tracks and wheels.
- Rotation of the turret (i.e. main body) with respect to the chassis (i.e. wheel/track base).
- Force sensor between chassis and main body.
- Torque sensor on wheels and body rotation.
- A 3-axis accelerometer.
- One degree of freedom rigid arm with gripper, allowing the robot to link to other robots as well as to grip objects.
- Three degrees of freedom flexible arm with gripper, as above, allowing the robot to link to other robots as well as to grip objects.
- Optical barriers on the grippers.
- 15 IR proximity sensors located around the circumference of the robot.
- 8 ambient light sensors located around the circumference of the robot.
- 4 downward facing IR proximity sensors located on the underside of the robot.
- 24 (8x3) colour LEDs located around robot's body (for local inter-robot visual communication.).
- Omni-directional colour CCD camera.
- One speaker and four microphones. (For transmission and directional reception of audio communication, in addition to more general audio sensing.).
- Humidity sensors.
- Temperature sensors.

The majority of the s-bots features listed above were not incorporated in the model used in Baldassarre et al.'s experiment. Each robot controlled one set of effectors: those powering the robots' differential drive locomotion. In addition, each robot had a single sensor modality, a "traction sensor", located at the point at which its turret (i.e. main body) joins its chassis. This sensor returns the angle and intensity of the force exerted by the turret on the chassis, which is the combined result of movement of the robot itself and of the robot(s) to which it is connected. An interesting consequence of this is that the turret of each s-bot physically integrates the forces applied by the robot itself and by the other robots; a robot's sensor thus gives information on the average force being applied to the turret and consequently on the average direction in which the robots to which it is connected are attempting to move. Rigid and flexible links were modelled, but the process of linking and unlinking was not - the robots were assumed to be linked from the outset. In addition, links were passive, that is, the robots had no control over their gripping arms. Baldassarre et al. constructed their simulation by using SDK Vortex™, a rigid body dynamics simulator. This simulator served to reproduce the dynamics due to inertia, friction and collisions between physical bodies with a reasonable degree of accuracy.

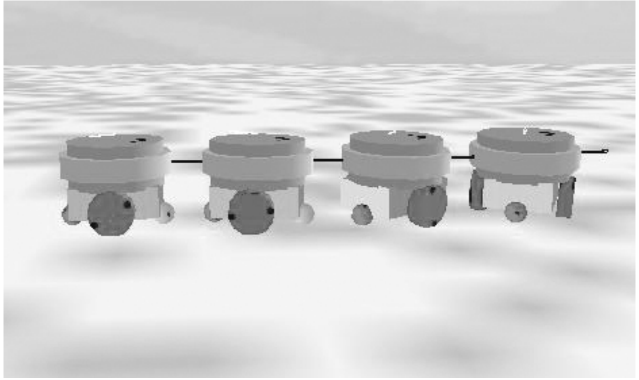

Figure 171 – Four linked robots in an empty environment, as rendered by the physics simulator (Baldassarre et. al. 2003a).

3.11.21.4.2. Neural Networks and Encoding Scheme

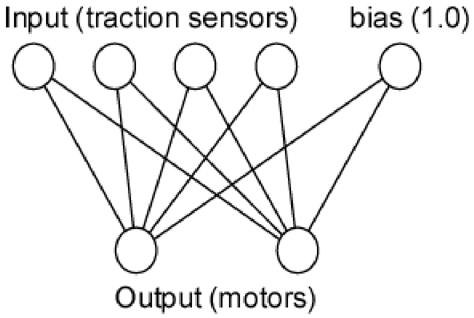

Figure 172 – Neural network controller: A simple feed-forward network with no hidden layers.

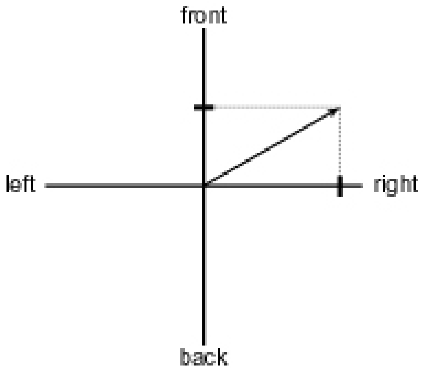

Figure 173 – Mapping direction and intensity of traction force onto four network input values.

Each robot was controlled by a simple feed-forward network with no hidden layers, as shown in Figure 172. The network comprised five input nodes, four for sensor input and one bias unit, and two motor output nodes. Each input was connected to each output by a weighted connection. Each of the sensor inputs represents the intensity of the traction sensor in one of four different directions relative to the robot's chassis, as illustrated in Figure 173. The activation values of the inputs are normalised to the range [0:1]. The bias input provides a constant input of 1. Connection weights were in the range [-10:10]. The activation function of the output (motor) nodes was the standard logistic function, and then mapped linearly onto the range [-10:10] to set the wheel speeds.

Only the 10 connection weights of the neural network controller were placed under evolutionary control. A simple fixed-size binary encoding was used. Each connection weight was represented by 8 bits which, when decoded, mapped onto the range [-10:10]. The total length of a genotype was thus 80 bits.

3.11.21.4.3. Evaluation and Fitness Function.

In this experiment, an encoded potential solution - or `genotype' by analogy with natural systems - specified the parameters of a single neural network controller. A team was generated by decoding a genotype and then making four identical copies of the decoded controller - one for each robot. This method, commonly known as "cloning", serves to ensure that the team have identical controllers, and are thus homogeneous. Once decoded, a team was subject to a subject of 5 trails, each lasting 150 time-steps. In each trial, the team received a score according to the fitness function that is set out below. The fitness score assigned to the genotype specifying the team was then simply the mean trial score. At the beginning of each trial, each robot's orientation was set at random, thereby subjecting teams to a range of different initial conditions. However, since different sets of initial conditions can present a team with different challenges. In order to present a "level playing field", five sets of randomly generated starting orientations were generated at the beginning of each generation of the genetic algorithm. At each generation, each team was then evaluated the same five set of starting orientations.

The task set by Baldassarre et al., was simply for the robots to move as far as possible from their initial position. The fitness function was designed to reflect this requirement. The score assigned to a team in a single trial was simply: d/D_{max}, where d was the Euclidean distance between the location of the team's centroid (or centre-point) at the beginning of trial and the location of the team's centroid at the end of the trial. D_{max} is the distance a single (i.e., unattached) robot is capable of travelling in 150 time-steps when moving at full speed - dividing by D_{max} therefore served to normalise the fitness score.

3.11.21.4.4. The Evolutionary Algorithm

Baldassarre et al used a simple generational genetic algorithm (G.A.). An evolutionary population contained 100 genotypes. The initial population was randomly generated, with each bit in each genotype being set on or off with equal probability. At the end of each generation, (i.e. after all individuals had been evaluated), genotypes were ranked by score. The 80 lowest scoring genotypes were then discarded, and the next generation was "bred" from the remaining 20 genotypes. Each of these 20 "parent" genotypes was copied 5 times, producing 5 "offspring" genotypes. The resulting 100 offspring genotypes were then subject to mutation, and formed the population which would be evaluated in the next generation of the genetic algorithm. Mutation consisted in flipping each bit each bit in each offspring genotype with a probability of 0.03.

3.11.21.4.5. Results

Baldassarre et al. conducted 10 separate evolutionary runs, each lasting 100 generations of the genetic algorithm. To get a good estimated of the performance of the evolved solutions, they evaluating the best controller from each run over a series of 100 separate trials. The mean score for each run was, in descending order: 0.947, 0.946, 0.945, 0.943, 0.934, 0.931, 0.923, 0.860, 0.839, and 0.765. (Recall that a score of 1.0 was impossible, since this would mean that the team could move as far a single robot, unlinked robot).

The teams evolved in the ten runs produced qualitatively similar behaviour. At the beginning of a trial, the robots pull in different directions; they then begin to orient their chassis in the direction that the majority of robots are pulling and begin to move in that direction, compensating for mismatches in orientation that arise whilst the team are moving. To an extent, this process can be seen from Figure 174, which shows the orientation of the chassis of a team of four robots over the course of two independent trials. The basic strategy underlying the robots behaviour can be summarised as follows:

- In the absence of any traction sensor input, a robot moved ahead at full speed.
- Given traction sensor input, a robot tended to turn so as to negate that force. Thus, if the force came from the left, the robot turned left; if it came from the right, the robot turned right. This resulted in the robots matching each other's orientations.
- Robots did not turn in response to low traction force coming from the opposite direction to which they were travelling, but did so in response to large forces. This helped resolve situations in which, for example, three robots were headed north, whilst one was headed south. In such a case, all robots would experience a force opposing their direction of travel, but only the single south-facing robot would be subject to sufficient force to ensure that it rotated.

In addition to the above, the authors also note that the robots' response to sensor input was non-linear with respect to the intensity and direction of force, and that these non-linearities appear to have played an important functional role in ensuring that the robots consistently rapidly achieved (and maintained) a common direction.

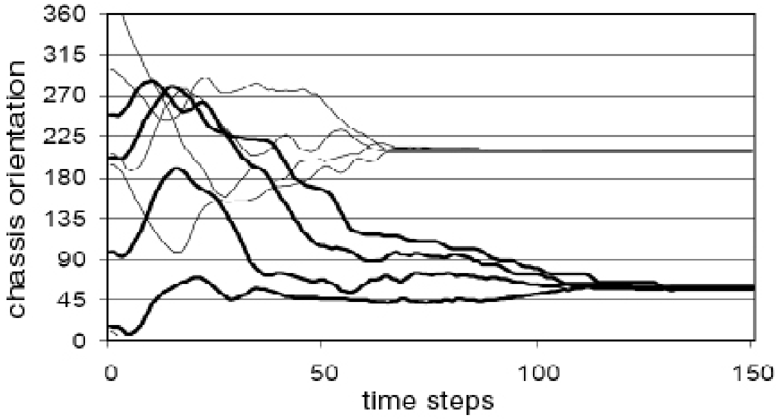

Figure 174 – The orientations of the chassis of each of the four robots over the course of two independent trials, each with a different set of starting orientations (Baldassarre et. al. 2003a).

3.11.21.4.6. Generalisation

Baldassarre et al. also investigated how well their evolved controllers would perform in situations other than those for which they had been evolved. They investigated three main variations on the original task specification:

- **Variation in the team's formation:** The authors investigated different team sizes, different team formations, and the use of flexible links between the robots.
- **Collision avoidance:** The authors investigated the ability of formations or robots to negotiate obstacles and walls.
- **Object transportation:** The authors investigated the ability of the robots to transport and object to which they were attached and an object which was enclosed by the formation.

The evolved controllers were found to generalise well in each case. Each will now be briefly reviewed.

3.11.21.4.7. Varying the Robots' Formation.

The original controller was evolved to control a system of four robots, connected in a line formation using rigid connections between robots. Baldassarre et al. investigated the following variations on the original formation set-up:

- **Variation in team size:** The authors investigated teams of between 2 and 10 robots connected in line formation.
- **Different formation topologies:** In addition to the line formation, the authors also investigated the performance of teams with two different formation topologies both

comprising eight robots: a `star' formation and a ring formation. Examples of each of these formation topologies can be seen in Figure 175.

- **Flexible links between robots.** During evolution, robots had been connected by rigid joints. The authors also investigated the consequences of using of flexible links with each of three formations described above (i.e. line, ring, star). A flexible link comprised two segments connected by a hinged joint.

The authors observed that the evolved strategy generalised well to each of the cases above. In each case, the robots retained their ability to negotiate a single common direction of travel and then move in this direction in a coordinated fashion.

3.11.21.4.8. Obstacle Avoidance

Using the original controllers, Baldassarre et al. tested the ability of teams of robots to negotiate walls and obstacles. They used two test environments: a walled arena and a walled arena with obstacles. They tested each of three formations listed above, using both rigid and flexible connections between each. In each case, formations of linked robots were found to be capable of negotiating walls and obstacles. This can be explained by considering that when a robot collides with an obstacle (or wall), the robot will be subject to a traction force in the opposite direction of the robot's direction of movement (this force is due the robot's turret pushing against the encountered obstacle). This would be sufficient to cause a single, unconnected robot to alter its direction. In addition, the other (connected) robots will also be subject to a traction force in the opposite direction their direction of movement, by virtue of the colliding robot's collision. It is interesting to note that when the robots were connected by flexible links, the shape of their formation became sufficiently deformed that the group could pass through relatively confined spaces. This can be seen in the case of the ring formation shown in Figure 175 (ii).

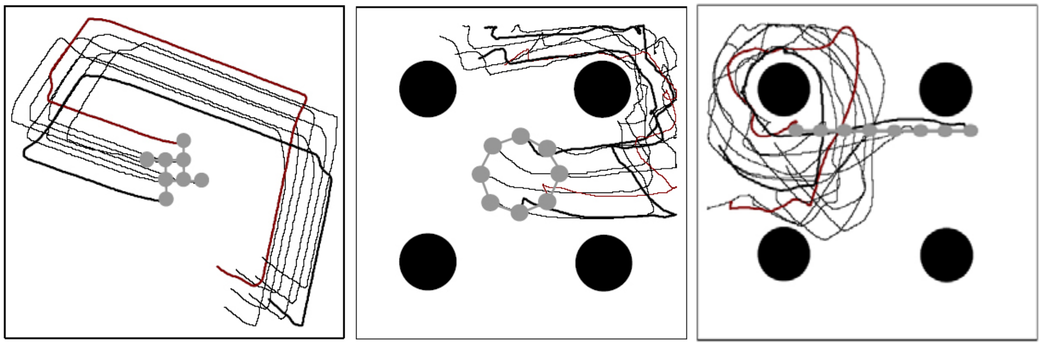

Figure 175 – Negotiating obstacles with different formations. Each example shows traces of the robots' movements as they move around in a walled arena. (i) A 'star' formation of 8 robots connected by rigid links in a walled arena. (ii) A `ring' formation of 8 robots connected by flexible links in a walled arena containing 4 circular obstacles. (iii) A 'line' formation of 8 robots connected by flexible links in a walled arena with 4 circular obstacles (Baldassarre et. al. 2003a).

3.11.21.4.9. Object Transportation

Baldassarre et al. also found that s-bots connected to an object, or connected so as to form a closed structure around an object, were capable of coordinated movement whilst pulling/pushing the object. shows the first of these two cases, in which each robot is connected (only) to the object. The robots' success with this task variation can be attributed to the fact that the object itself acts as a rigid link between the robots. The authors note that the robots were only successful when the weight of the object was relatively low. Heavier objects generate sufficient inertial force to disrupt team movement. This is primarily because transportation was reliant upon some of the robots *pushing* the object, and thus being subject to traction forces in the opposite direction to their direction of movement. Recall from above that the robots did not change direction in response to small forces in the opposite direction to their direction of movement (and hence can push light objects), but did do so in response to larger forces (and hence could not push heavy objects).

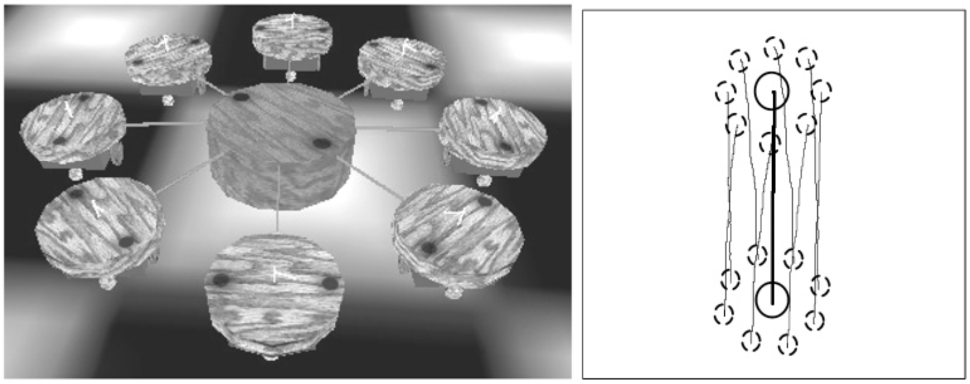

Figure 176 – A group of eight robots transport an object. Left: The manner in which the robots were connected to the object. Right: A trace of the movement of the robots and the object (Baldassarre et. al. 2003a).

3.11.21.5. Homogeneous System of Robots

The goal of Quinn et al.'s experiments was to evolve controllers for a homogeneous system of 3 real robots. The wheeled robots were minimally equipped with active IR sensors that provided extremely limited and ambiguous information about the location of other robots. Starting within sensor from random initial orientations and positions, the robots were travel a distance of 1m from their initial location whilst avoiding collisions and remaining within sensor range of one another. Solutions to this task involved robots adopting different roles. The authors claim that, by virtue of the robots' minimal sensors and the close coupling of the robots, designing controllers by more conventional means would have presented a significant challenge.

3.11.21.5.1. The Robots

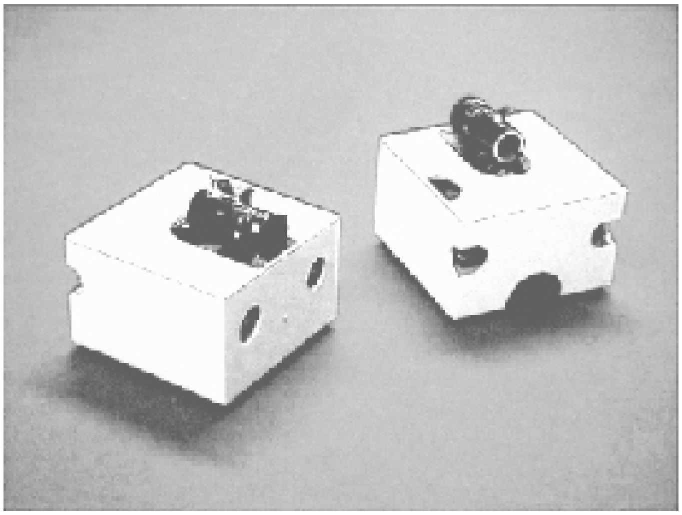

Figure 177 – Two of the robots used in Quinn et al's experiments. Each robot was equipped with two front-facing and two rear-facing 4 active IR sensors, and two motor-driven foam-rubber wheels. The cameras shown in the photo were not used in the experiments (Quinn et al. 2003).

Quinn et al. used a system of three robots that had been designed and constructed at Sussex University. Two of the robots are shown in Figure 177. All robots were built to the same specification. Each robot's body is 16.75 cm wide by 16.75 cm long by 11 cm high (this excludes the additional height of the camera). Two motor-driven wheels, made of foam-rubber, are arranged one on either side of the robot and provide locomotion through differential drive; the robots have a top speed of just over 6cm/s. An un-powered castor wheel, placed rear-centre, ensures stability. In the experiments, a robot's only source of sensory input comes from its four active infrared sensors, each comprising a paired infrared emitter and receiver. Each robot has two infrared sensors at the front and two at rear, as illustrated in Figure 178. Although the robots are also equipped with a 64 pixel linear CCD array camera (seen in the photograph), a 360 degree electronic compass, and wheel rotation sensors (i.e. shaft encoders), the controllers evolved by Quinn et al. were prevented from making use of these devices.

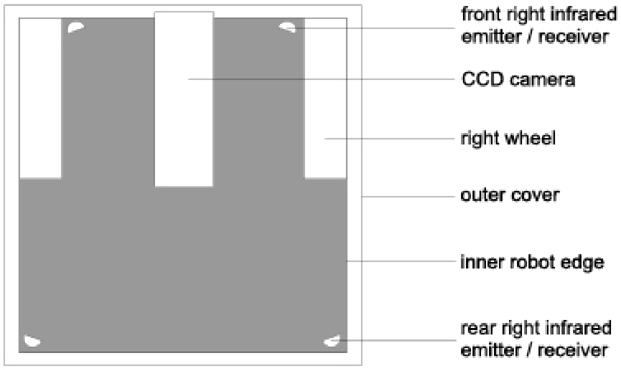

Figure 178 – Plan view of a robot (Quinn et al. 2003).

The robots were controlled by a host computer, with each robot sending its sensor readings to, and receiving it's effector activations from this machine via a radio communication system. Each robot had a 80C537-based micro-controller board for low-level control. The host computer was responsible for running the controller for each robot, updating each controller's inputs with the sensor readings from the appropriate robot, and transmitting the controller's output to the robot. Bandwidth constraints allowed for each robot to be updated at a rate of approximately 5 Hz. Note that radio communication occurred only between a controller (on the host computer) and its physical robot, there were no communication channels between robots or between controllers. In addition, although the physical instantiation of the robots was implemented as a host/slave system, conceptually the robots can be considered as independent, autonomous agents by virtue of the logical division of control into distinct and self-contained controllers on the host machine.

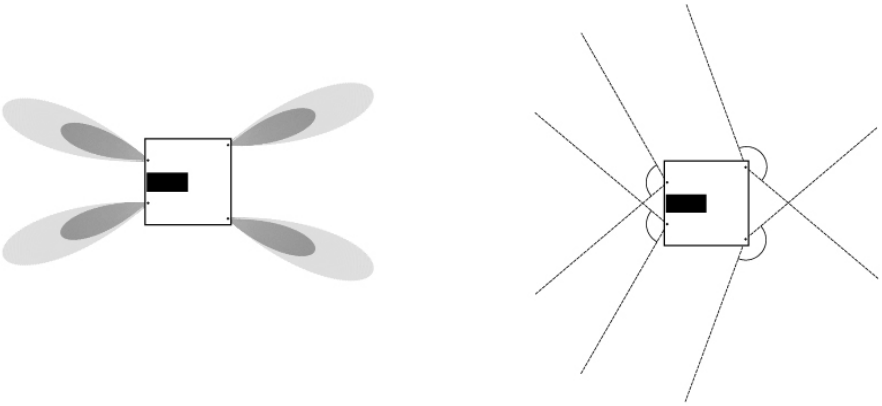

Figure 179 – An illustration of the limitations of the robots' IR active sensors. Left: An approximate illustration of the extent to which reflected IR can be sensed (dark grey area), and

the extent to which IR beam is perceptible to other robots (light grey area). Right: The angles from which a robot can perceive the IR emissions of others (Quinn et al. 2003).

One of the main factors that made the robots' task challenging was their extremely minimal sensors. It will be useful to describe the limitations of the active infrared sensors, especially those specific to a multi-robot scenario. An active infrared sensor comprises a paired infrared emitter and receiver. Its normal function is to emit an infrared beam and then measure the amount of infrared that reflects back from nearby objects. In this way, the robots can use their sensors to detect other robots up to a maximum of about 18cm (i.e. just over one body length away). The dark grey beams in the left-hand panel of Figure 179 approximately indicate the limited areas in which a robot can detect other robots in this manner. IR sensors are sometimes referred to as proximity sensors, however this is somewhat misleading. Whilst the sensor reading due to reflected IR is a non-linear, inverse function of the distance to the object detected, it is also a function of the angle at which the emitted beam strikes the surface of the object, and of the proportion of the beam that strikes that object. It is because an IR sensor reading combines these three factors into a single value that, even in normal function, sensor readings are ambiguous.

The ambiguity of IR sensors is significantly increased this experiment, due to the use of multiple robots. This is because the robots' sensors interfere with one another. Since each robot is constantly emitting IR, a robot's IR emissions can also be directly sensed by other robots. The light grey beams in the left-hand panel of Figure 179 indicate the approximate area in which a robot's infrared emissions may be directly detected by other robots. The maximum range at which emissions can be detected is approximately 30cm—almost twice the range at which a robot can detect an object by reflected IR. The right-hand panel of Figure 179 illustrates the range of angles at which a robot can receive the IR emissions of other robots. The sensor value due to receiving another's IR emissions is also the combined function of a number of factors: It will depend on the distance between the robots, the angle at which the emitted beam strikes the other robot's receiver, and which portion of the beam strikes the receiver (IR is significantly more intense at the centre of the beam than at the edges). Readings due to direct IR are thus ambiguous for the same reasons as reflected IR. However, ambiguity is compounded by the fact that, to the robot, readings due to reflected IR are indistinguishable from those due to the reception of IR emissions of other robots. Moreover, a sensor reading may be the result of a combination of both reflected and direct IR and it may be due to one or both of the other robots.

From the above, it should be clear that the robots are subject to significant sensory constraints. Sensors provide very little direct information about a robot's surroundings. Any given set of sensor input can be the result of any one of large number of significantly different circumstances. Furthermore, outside the limited range of their IR sensors, robots have no indication of each other's position. Any robot straying more than two body-lengths from its team-mates will cease to have any indication of their location (a situation which can also occur at much closer distances).

3.11.21.5.2. Task Specification

The robots were required to perform a coordinated movement task, as follows: The three robots are placed in an obstacle free environment in some random configuration, such that each robot is within sensor range of the others. Thereafter, the robots are required to move, as a group, a certain distance away from their initial position. The robots were not required to adopt any particular formation whilst moving as group, only to remain within sensor range of one another, and to avoid collisions. During evolution, groups of robots were evaluated on their ability to move their group centroid 1 m within the space of three simulated minutes. However, the authors expected that a team capable of this would be able to sustain formation movement for much longer periods. Since the robots start from initial random configurations, it was anticipated that successful completion of the task would entail two phases. The first entailing the team organizing itself into a formation, and the second entailing the team moving while maintaining that formation.

3.11.21.5.3. Robot Simulation

Controllers were evolved in simulation before being transferred onto the real robots. As discussed earlier in this document, a major problem with evolving in simulation is that robots may become adapted to inaccurate features of the simulation that are not present in the real world. However, building completely accurate simulation models of the robots and their interactions would be an onerous, and potentially impossible, task; moreover, it would be unlikely that such a simulation would have significant speed advantages over evolving in the real world. To address this problem, Quinn et al. employed Jakobi's minimal simulation methodology (Jakobi 1998). They built a relatively crude, fast-running simulation model of the robots and their interactions, based on a relatively small set of measurements. The parameters of this model were systematically varied, within certain ranges, between each evaluation of a team. Parameters included, for example, the orientation of each robot's sensors, the manner in which a robot's position was affected by motor output, the degree of sensor and motor noise, and the effects of sensory occlusion and IR interference. While it was generally either difficult or time-consuming to measure parameters needed for the simulation with great accuracy, it was relatively easy to specify a range within which each of the parameters lay, even if that range was wide. Varying parameters within these ranges meant that a robot capable of adapting to the simulation would be adapted to a wide range of possible robot-environment dynamics, including those of the real world. In addition to compensating for inaccuracies in their measurements, the authors also used variation was used in the same way to compensate for inaccuracies in their modelling, estimating the error due to these inaccuracies and adjusting parameter ranges to compensate. Importantly, this approach allowed them to sacrifice accuracy for speed and employ cheap, inaccurate modelling where more accurate modelling would have incurred significant computational costs. A full, detailed description of their implementation can be found in a technical report [Quinn et al. 2002].

3.11.21.5.4. Evaluation and Fitness Function

In this experiment, a genotype specified the parameters of a single neural network controller. A team was generated by decoding a genotype and then making three identical copies of the decoded controller – one for each robot. This ensured homogeneous controllers. The team was then evaluated in simulation. An evaluation comprised multiple trials, each from a different starting position (as will be discussed shortly). In each trial, the team received a score according to the fitness function that is set out below. The fitness score assigned to the genotype specifying the team was then simply the mean trial score. At the beginning of an evolutionary run the number of trials per team was set at 60. However, once the population had begun to find reasonable strategies (i.e., once scores began to exceed 70% of the maximum possible), the number of trials was increased to 100.

Different starting positions could potentially present different challenges. For this reason, in any given generation of the genetic algorithm, every team was evaluated over the same set of starting positions. At each generation, a set of *N* starting positions were randomly generated, where *N* was the number of trials per team at that stage of the run. The method of generating a starting position is shown in Figure 180. Each team was evaluated for one trial from each of the starting positions in this set. An additional source of potential variation between teams was due to the variation in simulation parameters that were introduced as part of the minimal simulation methodology (e.g. wheel speeds, sensor positions, etc.). To counteract this, a full set of simulation parameters, randomly generated within the appropriate ranges, were associated with each of the different starting positions. These simulation parameter sets were also generated anew at each generation of the evolutionary algorithm. In exposing each team to the same set of evaluation conditions as its competitors, Quinn et al. sought to ensure that, insofar as was possible, variation in score between teams reflected variation in their ability, not variation in their evaluation conditions.

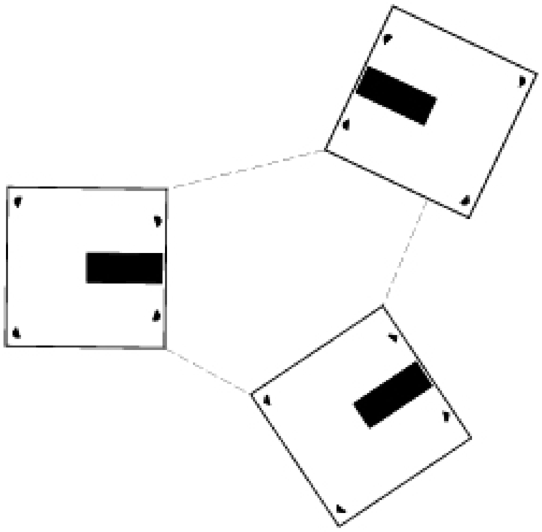

Figure 180 – An example starting position. Starting positions were randomly generated by the following method: Each robots' orientation was set randomly in the range [0:2π], and the (shortest) distance between the edges of each robot and its nearest neighbour was set randomly in the range [10 cm : 22 cm].

Reflecting the task description, the evaluation function was designed to assess the ability of the team to increase its distance from its starting position, while avoiding collisions and staying within sensor range. It therefore consisted of three main components: First, at each time step of the trial, the team is rewarded for any gains in distance. Second, this reward is multiplied by a dispersal scalar, reducing the fitness increment when one or more robots are outside of sensor range. Third, at the end of a trial, the team's accumulated score is reduced in proportion to the number of collisions that have occurred during that trial. The simulation, like the real robots, was updated at 5 Hz, thus each trial lasted a maximum of 900 simulation time steps (three simulated minutes). A trial was terminated early if (i) the team achieved the required distance, or (ii) the team exceeded the maximum number of allowed collisions.

More specifically, a team's trial score was given by:

$$P\left(\sum_{t=1}^{T}\left[f(d_t;D_{t-1})(1-\tanh(st-20))\right]\right)$$

Here P is a collision-penalty scalar in the range [0:5:1], given by $P = 1 - c/(2c_{max})$, where:

c *is the number of collisions between robots*
c_{max} *= 20 is the maximum number of collisions allowed*

Simulation time steps are indexed by t, and T is the index of the final time step of the trial. The distance gain component is given by the function f. This measures any gain that the team make on their previous best distance from their initial location: A team's location is taken to be the centre-point (or centroid) of the group. If d_t is the Euclidean distance between the group's location at time step t and its location at time step $t = 0$, D_{t-1} is the largest value that d_t has attained prior to time step t, and D_{max} is the required distance (i.e. 100 cm), then the function f is defined as:

$$f(d_t , D_{t-1}) = D_t - D_{t-1} \quad \textbf{if } D_{t-1} < d_t < D_{max}$$

$$f(d_t , D_{t-1}) = 0 \quad \textbf{otherwise}.$$

The final component of a team's trial score, the scalar s_t, is a measure of a team's dispersal beyond sensor range at time step t. If each robot is within sensor range of at least one other, then $s_t = 0$. Otherwise, the two shortest lines that can connect all three robots are found, and s_t is the distance by which the longest of these exceeds sensor range. In other words, the team is penalized for its most wayward member. The *tanh* function ensures that, as the robots begin to disperse, the team's score increment falls away sharply. However, the gradient of the *tanh* curve falls off as the distance between the robots increases, ensuring that increases in distance will still receive some minimal reward, even when the robots are far apart.

3.11.21.5.5. Neural networks

The robots were controlled by artificial neural networks. The thresholds, weights and decay parameters, and the size and connectivity of these networks were under evolutionary control. The networks were encoded by means of a topological encoding scheme, designed to minimise the disruption of network topology resulting from the addition and deletion of genetic material. A neural network comprised four sensor input nodes, four motor output nodes, and some number of neurons. These were connected together by some number of directional, excitatory and inhibitory weighted links. The networks had no explicit layers, so any neuron could connect to any other, including itself, and could also connect to any of the sensory or motor nodes.

The neurons used were loosely based on model spiking neurons. At any time step, the output, O_t, of a neuron is given by:

$$O_t = 1 \quad \textbf{if } m_t > T$$

$$O_t = 0 \quad \textbf{otherwise}$$

where:

T *is the neuron's threshold*

Here m_t is analogous to membrane potential in a biological neuron; it is a function of a neuron's weighted, summed input(s) integrated over time, such that:

$$m_t = am_{t-1} \sum_{n=0}^{N} w_n i_n \quad \text{if } O_{t-1} = 0$$

$$m_t = bm_{t-1} \sum_{n=0}^{N} w_n i_n \quad \text{if } O_{t-1} = 1$$

where:

> a *and* b *are decay constants (constrained to the range [0:1])*
> w_n *designates the weight of the connection from the* n^{th} *input (*i_n*) that scales that input (unconstrained)*

For certain parameter settings this neuron will behave like a simple spiking neuron, accumulating membrane potential, firing and then discharging (i.e. with $a > b$ and $T > 0$). However, the neurons also exhibit a range of other temporal dynamics under different parameter settings.

Each sensor input node outputs a real value in the range [0:0:1:0], which is simple linear scaling of the reading taken from its associated sensor. Motor outputs consist of a `forward' and a `reverse' node for each motor. The output, M_{out}, of each motor node is a simple threshold function of its summed weighted inputs:

$$M_{\text{out}} = 1 \quad \textbf{if} \;\; \sum_{n=0}^{N} w_n i_n > 0$$

$$M_{\text{out}} = 0 \quad \textbf{otherwise}$$

The final output of each of the two motors was attained by subtracting that motor's reverse node output from its forward node output. This gave three possible values for each motor output: {-1, 0, 1}.

3.11.21.5.6. Encoding Scheme

Networks were encoded by a topological encoding scheme devised by Quinn [Quinn et al. 2002; Quinn 2004]. A main concern in the design of this encoding scheme was that topology was not determined by the position at which individual neurons and connections were encoded on the genotype, as it is in standard fixed-sized encoding schemes. Such position-dependent encoding schemes are particularly vulnerable to disruption by the insertion and deletion of genetic material. The encoding scheme used employed a position-independent encoding, whereby each encoded neuron was associated with an identity tag, and these tags were used to specify connections between neurons, in a manner detailed in the following paragraph.

The procedure for encoding and decoding networks was as follows: Each neural network controller was encoded in a `genotype', this consists of a list of `genes'. Each gene encoded the parameters for an individual neuron and its associated connections. The form of each gene was $\{X, T, a, b, L_{in}, L_{out}\}$, where T, a and b are the (real-valued) threshold and decay parameters, and Lin and Lout are lists of input and of output connections respectively. X is an integer that serves as an `identity tag'. This is assigned to a gene upon its creation (genes are created either in the first G.A. generation when a population is created, or when a new gene is added to a genotype as a result of a macro-mutation). The identity tag (i.d.) assigned to a new gene is initially unique - no other gene will have the same i.d. at the time it is first assigned. The i.d. is not subject to mutation, but it is inherited. Hence, if two genes (on two different genotypes) have the same i.d., it will be because both genes have a common ancestor. Sensor and motor nodes also have unique identity tags. The identity tags of the neurons, and the sensor and motor nodes, are used to determine network connectivity; each encoded connection specifies a source or destination by reference to these tags.

Each element in the connection lists, L_{in} and L_{out}, took the form, $\{Z,w\}$, where w is a (real-valued) connection weight, and Z is an integer, specifying an identity tag of the node (sensor, motor, or neuron) to which it should connect. In decoding the network, the source (or destination) of each connection in the list is determined by matching the tag it specifies with the i.d. of a neuron in the network, or that of a sensor or motor node. A connection may specify the tag of the neuron from which it originates, resulting in a recurrent connection being made. The integer Z is inherited and is subject to mutation (see below).

3.11.21.5.7. The Evolutionary Algorithm

A simple, generational, genetic algorithm (GA) was employed for this experiment. An evolutionary population contained 50 genotypes. In the initial population, each genotype encoded a randomly generated network with three neurons. Each neuron was randomly assigned between zero and seven input connections, and between one and seven output connections. Weights and thresholds were initially set in the range [-5:5] but were thereafter unconstrained. At the end of each generation, (i.e. after all individuals had been evaluated), genotypes were ranked by score. The 10 lowest scoring individuals were discarded and the remainder used to generate a new population. The two highest-scoring individuals (the `elite') were copied unchanged into the new population; thereafter, genotypes were selected randomly with a probability inversely proportional to their rank. 60% per cent of new genotypes were produced by recombination; mutation operators were applied to all genotypes except the elite.

3.11.21.5.8. Mutation Operators

Genotypes were subject to both macro-mutation (i.e. structural changes) and micro-mutation (i.e. perturbation of real-valued parameters). Micro-mutation entailed that a random Gaussian offset was applied, with a small probability, to all real-valued parameters encoded in the genotype, such that the expected number of micro-mutations per genotype was 2. The mean of the Gaussian was zero and its standard deviation was 0.33 of that parameter's range (in the case of decay parameters) or its initialisation range (in the case of weights and thresholds). Three types of macro-mutation were employed. First, a new neuron could be added to, or a randomly chosen neuron deleted from, the genotype. New neurons were initialised as described above, except that a new neuron could not have more than two input and two output connections (this served to reduce the probability that such an addition would be disruptive). The probability of neuron addition was set at 0.004 per genotype, and that of deletion was 0.01. Second, a new connection could be added to, or a randomly chosen connection deleted from, a randomly chosen neuron with the respective probabilities of 0.02 and 0.04 per genotype. Finally, a randomly chosen connection could be chosen and reconnected elsewhere, this occurred with a probability of 0.04. Deletion operators were applied with a higher probability than the addition operators, in order to ensure a steady (rather than rapid) rate of genotype growth.

3.11.21.5.9. Recombination Operator

The recombination (or 'cross-over') operator creates two new `offspring' individuals by combining encoded variables from two 'parent' genotypes. Since the authors used a variable-size encoding, the addition and deletion of genetic material meant that two genes at the same position on their respective genotypes may not be correlated at the phenotypic (i.e. network) level. If crossover were simply based on genotype position, as it is with fixed-length encoding schemes, it would often be highly disruptive, with unrelated variables being crossed-over (Jakobi and Quinn, 1998). To avoid this problem, the recombination operator utilised the gene's i.d. tags, rather than their position on the genotype, in order to maintain the structural integrity of the resulting genotypes. The i.d. tag is used by the recombination operator to pair genes with a common ancestor, and thereby helps to ensure that genes with similar phenotypic functionality were crossed. Crossover takes two parent genotypes, **P1** and **P2**, and generates two offspring, **O1** and **O2** in the following manner: Each of offspring initialised as a copy of a different parents (e.g., **O1** = **P1** and **O2** = **P2**). The offspring genes were then paired by identity number; any gene that **O1** has in common with **O2** is crossed (i.e. swapped) with a probability of 0.5. Any remaining un-paired genes are not crossed.

3.11.21.5.10. Results

Quinn et al. undertook a total of 10 evolutionary runs, four of which were terminated at an early stage because they seemed unpromising. The remaining 6 runs were left to run for between 3000 and 5000 generations of the evolutionary algorithm, being terminated once they ceased to show signs of any further improvement. The best recorded fitness score in each of these six runs exceeded 95 out of a possible 100 (recall that fitness scores are the mean score achieved over 100 trials). However, the existence of between-generation variation in starting positions and simulation parameters meant that the fitness scores tend to produce an overestimate of controllers' ability to perform in simulation. To achieve a better estimate, Quinn et al took the best controller from each of the six runs and averaged its performance over 5000 simulated trials. The mean, standard deviation, and quartile scores for the best controller of each run were as shown in Table 18.

Table 18 – Evaluation Scores Achieved by the Best Controller

Mean	Standard Deviation	Upper Quartile	Median	Lower Quartile
99.71	3.601	100.0	100.0	100.0
99.22	6.392	100.0	100.0	100.0
98.50	6.583	100.0	100.0	99.97
95.04	11.66	100.0	100.0	96.25
93.01	18.72	100.0	100.0	95.0
89.78	24.78	100.0	100.0	96.8

The best controller from each of the six runs was tested on the real robots and evaluated in simulation for 5000 trials. It was reported that in each case the behaviour observed in simulation was qualitatively reproduced in reality, that is, the same kinds of behaviour observed in simulation were observed on the real robots.

In the case of the highest scoring controller, the authors conducted a sequence of 100 consecutive trials with the real robots, each from a random starting position (generated by the procedure described above). The team successfully completed all trials. There was thus no evidence of any degradation of performance as a result of transferring the controllers to real robots. Transfer from simulation to reality appears to have been completely successful. Video footage of this team can be found on the Web page of one of the authors (http://www.cogs.susx.ac.uk/matthewq). The teams from the remaining five runs were not been tested in such a systematic fashion. The authors report that controllers from the second- and third-highest-scoring runs also transferred with no apparent degradation of performance. However, in the remaining three runs there was a noticeable degradation in the performance of the controllers when they were instantiated on the real robots.

The less successful transfer of controllers from lower-scoring runs was not surprising, as Quinn et al point out. In fact, it is an expected consequence of using Jakobi's minimal simulation methodology (Jakobi 1998). Recall that Quinn et al implemented a fast-running, inaccurate simulation. Inaccuracies in the simulation were covered up by exposing controllers to significant variation in the simulation parameters during evolution. Until controllers had adapted to the full range of variation in simulation parameters, there was no guarantee that they would transfer effectively to real robots. Thus, the less consistently controllers can perform a task in simulation, the less likely they will be able to do so on real robots, and the less likely a successful transfer is.

3.11.21.5.11. Behaviour of an Evolved System

Having achieved a successful transfer from simulation to reality, Quinn et al., went on to investigate the behaviour of the most successful of the six teams (i.e., the controller which averaged 99.7 in simulation). The team's behaviour had two main stages. Firstly, from their initial random positions, the robots organised themselves into a line formation. Secondly, once in formation, the robots move away as a group, maintaining their formation whilst moving. They showed that formation movement by the team is dependent upon each robot performing a functionally distinct role, and described the process by which these roles become allocated in the absence of any intrinsic differences between the robots. In describing the behaviour of this team, it will be useful to start with the second main phase, formation movement, before describing the process by which the robots organised themselves into the formation.

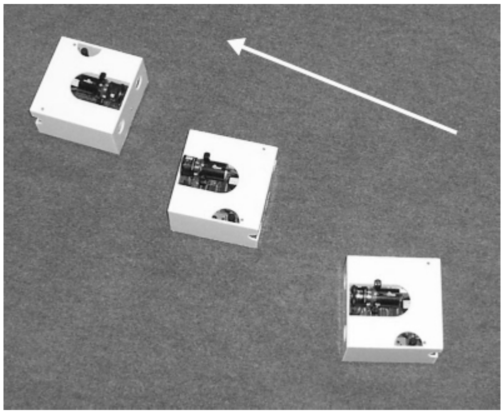

Figure 181 – Video still of the robots moving in formation. The arrow indicates the direction of overall movement. Note that the front-most robot travels in reverse, the two remaining robots move forwards (Quinn et al. 2003)

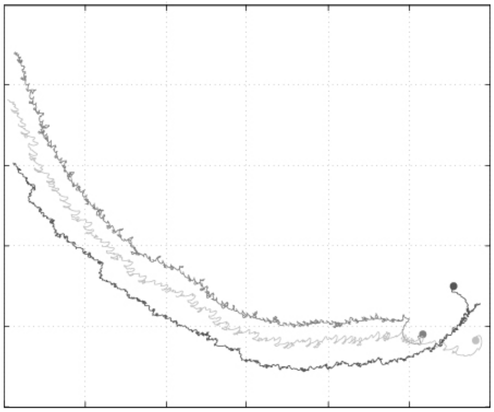

Figure 182 – An example trajectory. A trace of the position of each robot, recorded over a 5 minute period. Grid divisions are at 50cm intervals. Robots start bottom right (their initial positions indicated by dots). The data was generated in simulation (Quinn et al. 2003)

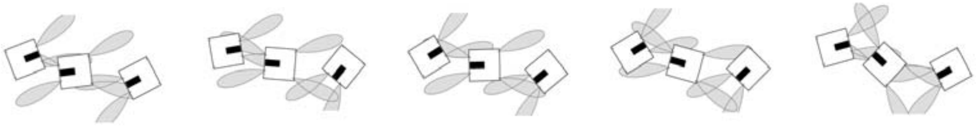

Figure 183 – An illustration of relative robot positions during formation movement. The time sequence above illustrates the relative positions of the robots during formation movement at intervals over a brief (4 second) period. Whilst moving in formation, the robots maintain contact with one another primarily by means of directly sensing each others' active IR beams; these beams are shown in grey in the diagram above (Quinn et al. 2003)

The team travelled in a line formation, as shown in the video still in Figure 181. The lead robot travels in reverse, while the middle and rear robots travel facing forwards. When travelling in formation, the team move at just over 1 cm/s, a relatively slow speed compared with the 6 cm s¡1 maximum speed of which an individual robot is capable. Whilst moving, each robot swings clockwise and anticlockwise while maintaining its position. The authors comment that, watching the video footage sped up, the team's locomotion appears almost snake-like. The sequence of diagrams in Figure 183 goes some way toward illustrating this aspect of the team's locomotion. Note from these diagrams that the robots rely almost entirely on the direct perception of each other's IR beams (i.e. sensory interference) in order to coordinate their movement. Despite the oscillating angular displacement of the robots, the robots' formation was extremely robust. The authors reported that, both in simulation and in reality, the robots appear to be able to maintain their formation indefinitely, despite having only been evolved to travel a distance of 1m. In the case of the real robots, the authors report that limitations in the size of their testing area precluded testing whether the formation would break down over large distances, since the robots cannot travel much further than 10 metres before encountering a wall. However, no formation had been observed to break down in the absence of an obstacle.

Formation movement was dependent upon each robot performing a different role, that is, each robot made a different contribution to the maintenance of the formation and its continued movement. This was shown by independently removing the front and rear robot from the formation. It was found that:

- The rear robot had no significant effect on the other two robots’ ability to maintain formation, but was crucial to sustaining locomotion. If the rear robot was removed from a full formation, the locomotion of the remaining pair ceased, there was no further significant displacement of their position. However, the pair maintained the same configuration as when in full formation, continuing to oscillate. If the rear robot was replaced, the group moved away once more.
- The middle robot responds to the presence of the rear robot by moving forwards. (This follows from the fact that, the absence of the rear robot, the remaining pair cease to travel.)
- In the absence of the front robot, the configuration adopted by the middle and rear robot in the formation was unstable. If the front robot was removed from the full formation, the middle robot swung around toward the rear robot and the two robots formed an opposed pair (adopting the same configuration as the front and middle robot in a full formation).

The authors also investigated the effect of reorganizing the robots’ formation, and found that each robot maintained its role within the formation is solely by virtue of the spatial organisation of the formation, rather than any long-term differences in internal state. If the middle robot was quickly picked up and rotated by 180 degrees, the formation was maintained and the team start to move in the opposite direction, with the robots which were previously front and rear adopting the roles appropriate to their new positions in the formation. If the rear robot is removed from the formation and appropriately placed behind the front robot, formation again moved off in the opposite direction, with each robot performing the role appropriate to its new position.

The authors then went on to investigate how the robots’ roles were initially allocated; in other words, how the robots achieve their formation positions from their random initial positions. Recall the system is homogeneous, and therefore the robots required some way of dynamically deciding which robot would adopt which role, and ensuring that each role was filled by only one robot. At a basic level of description, the process occurs in two (typically overlapping) stages. The first phase ends when two robots form an opposed, face-to-face pairing. This serves to differentiate the group. The excluded robot’s role is now determined—it will become the rear robot in the formation. The paired robots will become the front and middle robots in the formation, however, which of the two robots will take which role remains undecided at this point. The second stage entails the excluded robot ‘attaching’ itself to the rear of one of the paired robots. This serves to differentiate the paired robots, determining which will become the front and which the middle robot.

An explanation of this process requires some understanding of how a robot responded to IR sensor input. The following is a simplified description:

- In the absence of any sensory input, a robot moved in a small clockwise forwards circle.
- Activation of either (or both) of the rear sensors (in the absence of significant front sensor input) caused the robot to turn more tightly in a clockwise direction.
- A robot was generally ‘attracted’ to any source of front sensor input. It would rotate anticlockwise in response to any front left input and clockwise in response to front right input.
- Sufficient simultaneous activation of both front sensors causes the robot to reverse (thus avoiding collisions).

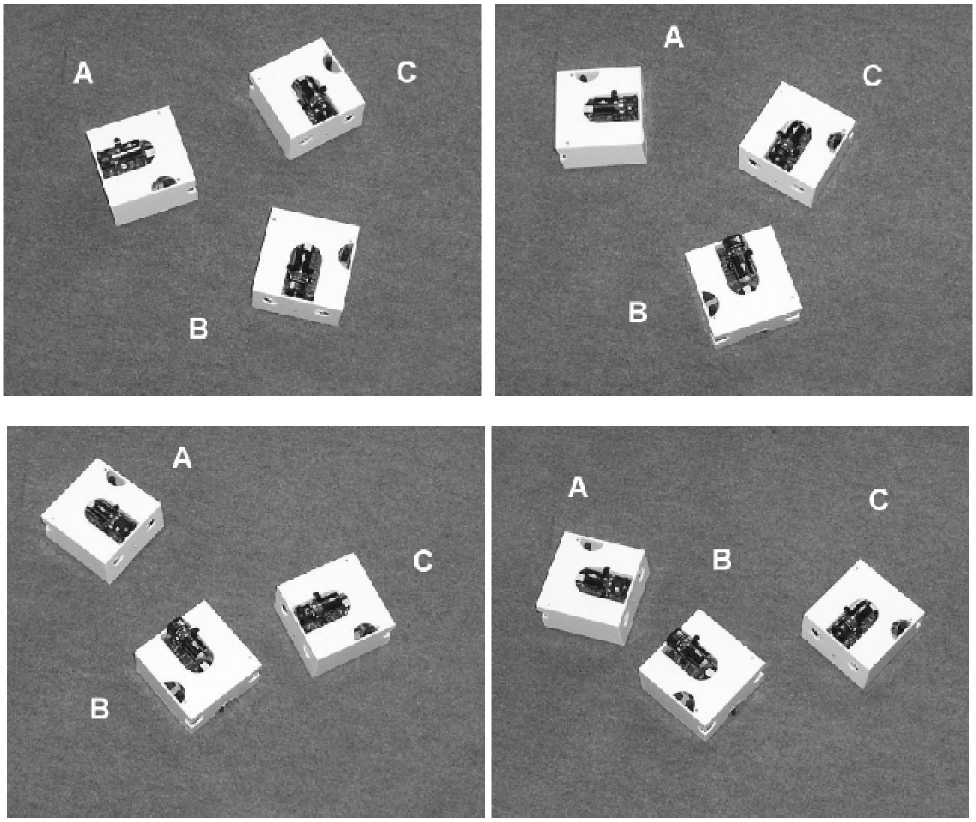

Figure 184 – An example of the robots moving into formation position. (i) The robot's initial positions. Initially, robot C is 'attracted' B's rear sensors, causing B to turn tightly, A circles away, clockwise. (ii) robots B and C begin to form a pair as A circles round towards them (iii) A disrupts the pair formation of B and C, subsequently pairing with B. (iv) C becomes attracted to B's rear sensors and begins to move into position. Shortly after this, the robots achieve their final formation (Quinn et al. 2003)

In the light of the above, the process of forming a formation can broadly be explained as follows: From its initial position, each robot will begin to circle clockwise until its senses another robot. A robot will typically first encounter either the front or rear IR beams of another robot (direct IR), or one of its side panels (reflected IR). (Recall that a robot can sense both IR reflected off the body of another robot and the IR beam of another robot, the latter being perceptible from twice the distance than the former.) The robot will become 'attracted' to whatever source of frontal IR input it encounters. A robot 'attracted' to the side of another robot will simply be ignored as it cannot be sensed. A robot attracted to the rear IR beams of another robot will in turn activate that robot's rear sensors, causing it to turn sideways on. If however a robot becomes attracted to the front IR beams of another, it will in turn activate the front sensors of that robot as it approaches, both robots will turn to face each other—mutually attracted. Prior to the arrival of the third robot, the facing pair maintain a stable, opposed formation, this serves as 'holding pattern', in which the pair await arrival of the remaining team member before they begin to move off. The remaining robot will subsequently become attracted to rear sensors of one of the pair, bringing the formation into completion. (Once in this paired formation, the robots cease to respond to rear sensor input by turning.) The process of achieving formation is not always quite as simple as the above description might imply. The pairing process may have to be resolved between three robots (as for example, in panels ii and iii of Figure 184 where one robot may disrupt the pair-forming of the other two). However, the explanation given above should be sufficient understand the basic dynamics of the process by which the robots achieve formation.

3.11.22. Evolution of Cooperation and Labour Division in Artificial Ants

This multi-robot research at EPFL is motivated by the fact that ants constitute around 15% of the biomass in most terrestrial environments and are therefore very successful biological systems. Two key features of this success seem to be their ability to cooperate and their efficient division of labour.

3.11.22.1.1. Experiments

This research investigates the effects of relatedness, levels of selection and group size on the evolution of cooperation and labour division in artificial ant colonies. Experiments are carried out in simulation, but the intention is to test the model on groups of Alice microrobots (see Figure 155).

Each ant in the simulation is encoded on the genotype as 3 thresholds (5 bits per threshold) which determine which of a limited number of behaviours an ant performs at a given time step in the simulation given the activity of neighbouring ants. The task is to forage food items, that is, locate them and return them to the nest within 5 simulation cycles. Small food items can be foraged by individual ants, whereas larger ones require ants to co-operatively forage. The performance of the colony is averaged over 20 trials.

The 4 experimental conditions that investigated the link between co-operative foraging and genetic relatedness were:

1. Heterogeneous colonies with individual level selection;
2. Heterogeneous colonies with colony level selection;
3. Homogeneous colonies with individual level selection;
4. Homogeneous colonies with colony level selection.

A population of 400 individuals was split into 20 colonies of 20 ants and evolved for 100 generations. Homogeneous colonies consist of clones with the same genotype, whereas heterogeneous ones consist of individuals with different genotypes. In individual level selection, individuals are selected from different colonies on the basis of their performance and bred to form 20 new colonies; in colony level selection, all the individuals in the highest fitness colonies are bred to form 20 new colonies.

In physics-based simulations it was found that homogenous groups outperform heterogeneous group in tasks that require cooperation and that colony level selection leads to the best performance [Perez-Uribe et al. 2003].

3.12. BIOMIMETIC SENSORS & SIGNAL TRANSDUCTION

Sensors are a necessary component to any closed loop control system. They provide feedback data on the status of the environment. Animals are adaptable and robust, capable of learning sophisticated navigation strategies and are supported by a rich array of sensors [Hallam & Hayes 2001]. In particular, biological organisms are characterised by distributed sensing capabilities which provide the basis for protection through reflex action. The basis for most biomimetic sensors is microtechnology, particularly for cantilevered sensors. For harsh environments, Si substrates are limited to 250^{o}C so wide bandgap semiconductors such as SiC with a bandgap of 2.2 eV are more appropriate for high temperature operation on planets such as Mercury or Venus. SiC has a high elastic modulus (of 448), high Moh's hardness (of 9), high wear resistance (of 9.15) and chemical inertness to acids making it ideal for mechanical components. SiC does not melt but sublimes at 1800^{o}C and can be grown as thin films by atmospheric or low pressure chemical vapour deposition on SiO_2 for SOI (silicon-on-insulator) fabrication.

Hair-like mechanical sensors may be based on pressure sensing piezoelectric cantilevers – such sensors are ubiquitous in the biological world (Ayers 2000). The vestibular sense is evolutionarily the most ancient sensory modality which in fish and crustaceans is implemented by the statocyst which a chamber containing a massive particle (usually calcium); the chamber is lined with hair receptors which sense the position of the particle and so sense orientation. As an illustration, the animals can be confused if the calcium is replaced by a ferromagnetic one, and the animal placed in a controlled magnetic field! The function of hair cells in the human ear for both hearing and vestibular control are well-characterised.

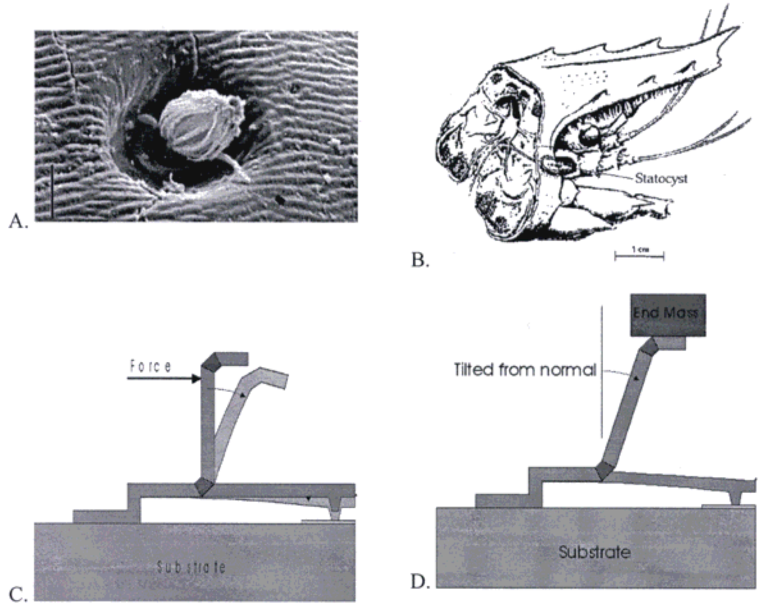

Figure 185 – a) Lobster water current receptor; b) Lobster statocyst balance organ; c) MEMS hair sensor; d) Inclinometer

The primary sensory cortex areas of the human brain have afferent connections from the sensory nuclei of the thalamus.

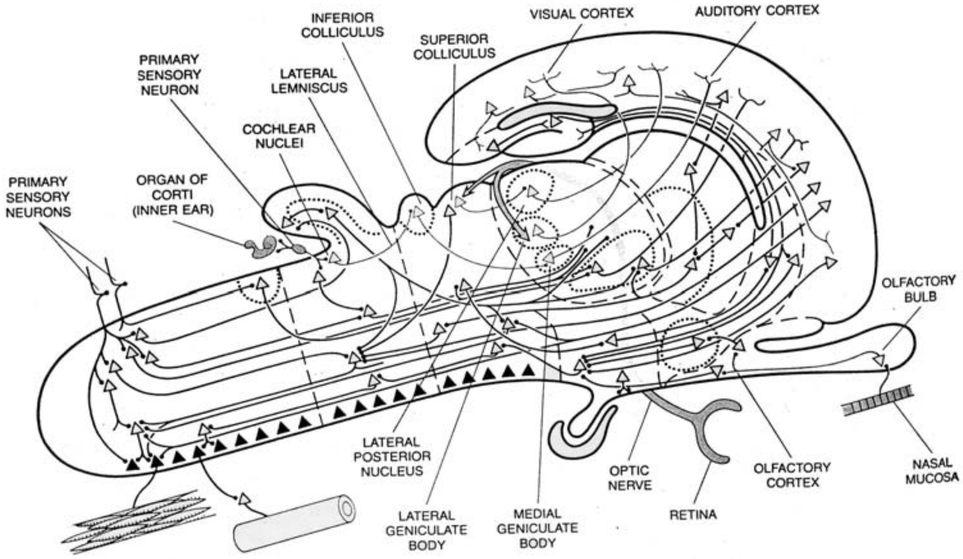

Figure 186 – Sensory (auditory, visual and olfactory) cortical pathways in the human brain: auditory information passes from cochlear nuclei of hindbrain through inferior colliculus of midbrain, mediate geniculate body of thalamus to auditory cortex; visual information passes from retina through two channels, one via lateral geniculate body of thalamus to visual cortex, the other via superior colliculus of midbrain to thalamus to visual cortex; alfactory informnation passes from nasal receptors directly to olfactory bulb of the cerebral cortex.

At the back of the brain on the surface of the occipital lobes lies the primary visual cortex (area 17). In the photoreceptor, light causes conformation changes in rhodopsin molecules which cause the release of neurotransmitters. This causes the closure of Na+ channels and decreases its output of neurotransmitter. Visual data comes from the lateral geniculate nucleus of the thalamus. The primary auditory cortex resides in the temporal lobes while olfaction is focussed on the forward underside of the frontal lobes. Auditory hair cells and photoreceptors of the eye do not fire action potentials but signal the magnitude of the receptor potential to adjacent neurons via chemical synapses. The auditory cortex receives inputs from the medial geniculate nucleus of the thalamus. The angular gyrus between the visual and auditory cortex mediates between them. The olfactory system is the only system whose primary sensory neurons lie at the surface of the body without any transduction element – all other primary sensory neurons lie in ganglia under the surface of the body. Fibres project into the olfactory bulb which terminate in the olfactory cortex. The auditory system begins with the primary sensory neurons terminating on secondary sensory neurons of the cochlear nucleus in the hindbrain which convey fibres to the medial geniculate body of the thalamus which projects to the auditory cortex.

Perception is fundamentally a categorisation process and relies on specific input features (cues) to perform such inferences of identification [Bruner 1957]. Clues to the nature of perception may be obtained from illusion effects such as the Miller-Lyer line. The line with extended arrowheads looks longer than the equal length line with arrowheads inclined inwards. Perception is thus an active process. From a purely signal processing view, transduction alone does not explain this effect: Magnitude, $M = kI^{0.3}$ where I=signal intensity. Perception is a decision process which utilises discriminatory cues of the input signals. This is an unconscious inference process of selecting the correct category which defines a set of specifications to the input. The most primitive categorisation process is to distinguish between features from the background and the capacity to distinguish events in one modality from another. Indeed, categorisation may be fundamental to "cause" and "effect" thinking as this requires the categorisation and conservation of a process' identity. The primitive categorisation is followed by a bracketing process of continual narrowing of the categorisation using additional cues in a trial-and-check process for selective gating of categories by inhibiting unlikely alternatives. Central to categorisation is the notion of accessibility, the speed at which stimulus input is coded into categories. This implies that perception is an active filtering process in that accessibility is dependent on perceptual readiness. It minimises the "surprise" element of the environment by matching category accessibility with event probabilities and likelihood and it maximises the attainment of desired objects or events. The higher the likelihood of an occurrence of an event, the greater the readiness to its identification. The more accessible a category, the less the stimulus input required for its activation. The likelihood of sensory events determines its perceptual readiness and accessibility. Such probabilities of occurrence of events are learned through experience. The relation between classes of events is not an autocorrelation function but is altered by learning. This is essentially model-building. These probability estimates of event occurrences are subjective in that strongly desired events exhibit higher probability estimates than those of less desired events (eg. gambler's fallacy), i.e. there is a strong motivational influence. The categorisation process implies that there is a match process between the stimulus input and stored memory traces. Hence, it is impossible to divorce perception from cognition in any well-defined fashion.

3.12.1. Vision

There is a need for some form of distance sensing. The primary sensory modality for most animals which provides this is typically vision which generally imposes a high computational load requirement for extensive image processing. Animals devote a significant fraction of their nervous system to vision and other sensory processes – the sensory cortex is estimated at 25% of the cerebral cortex in humans. Biological visual processing is based on feature detection involving three different regions of the brain: the retina, the lateral geniculate nucleus and the visual cortex. An object is recognised by the simultaneous appearance of a number of visual features in their correct relative spatial positions. The emphasis of such distance sensing is for generalised navigation for foraging, taxis, habituation, obstacle avoidance and exploration rather than prescriptive behaviours such as wall-following. Vision has evolved independently in many different phyla including some primitive organisms which often exhibit photoreceptors (eg. annelid worms).

Template-matching of sensory data with stored computational models in image processing is common [Binford 1982]. A fundamental problem is that the number of possible scenes far exceeds the number possible sensory measurements that can be made so simplification is necessary. For a polyhedron with *n* faces, there are *2n* views yet vision systems should be viewer-independent.

Table 19 – Subsystems Supporting Human Vision (Kosslyn 1994)

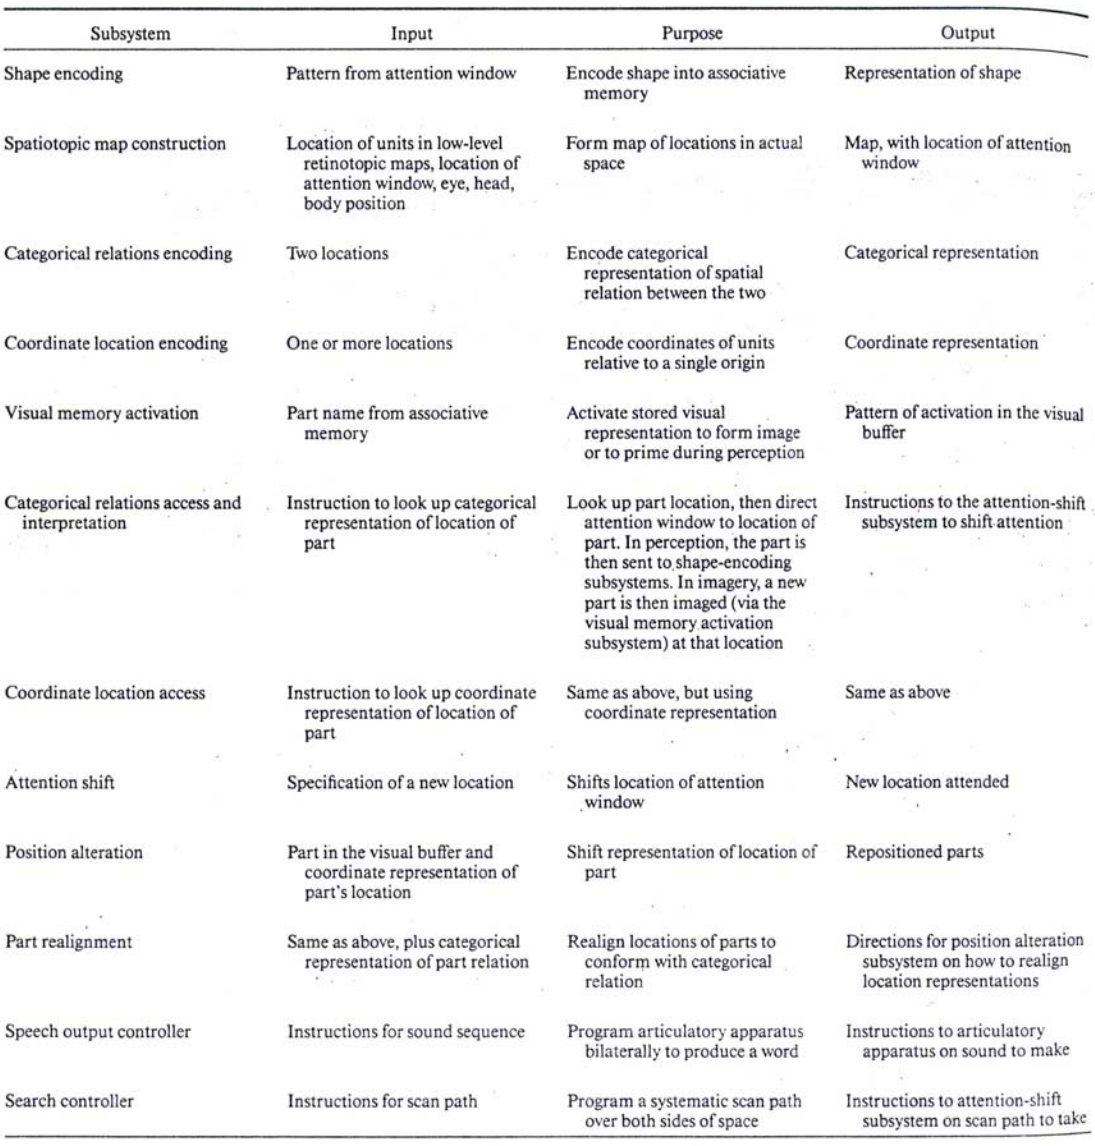

Subsystem	Input	Purpose	Output
Shape encoding	Pattern from attention window	Encode shape into associative memory	Representation of shape
Spatiotopic map construction	Location of units in low-level retinotopic maps, location of attention window, eye, head, body position	Form map of locations in actual space	Map, with location of attention window
Categorical relations encoding	Two locations	Encode categorical representation of spatial relation between the two	Categorical representation
Coordinate location encoding	One or more locations	Encode coordinates of units relative to a single origin	Coordinate representation
Visual memory activation	Part name from associative memory	Activate stored visual representation to form image or to prime during perception	Pattern of activation in the visual buffer
Categorical relations access and interpretation	Instruction to look up categorical representation of location of part	Look up part location, then direct attention window to location of part. In perception, the part is then sent to shape-encoding subsystems. In imagery, a new part is then imaged (via the visual memory activation subsystem) at that location	Instructions to the attention-shift subsystem to shift attention
Coordinate location access	Instruction to look up coordinate representation of location of part	Same as above, but using coordinate representation	Same as above
Attention shift	Specification of a new location	Shifts location of attention window	New location attended
Position alteration	Part in the visual buffer and coordinate representation of part's location	Shift representation of location of part	Repositioned parts
Part realignment	Same as above, plus categorical representation of part relation	Realign locations of parts to conform with categorical relation	Directions for position alteration subsystem on how to realign location representations
Speech output controller	Instructions for sound sequence	Program articulatory apparatus bilaterally to produce a word	Instructions to articulatory apparatus on sound to make
Search controller	Instructions for scan path	Program a systematic scan path over both sides of space	Instructions to attention-shift subsystem on scan path to take

Partial image matching may be achieved by evaluating similarity rather than congruence. A statistical measure of a distance metric is insufficient as the number of possible scenes far exceed the number of measurements, so structural simplification is required. A prototype object is characterised by a geometric function, i.e.. three dimensional form. Indeed, geometrical properties causally define function. For the Martian surface, such models must be based on ellipsoidal fitting – generalised cylinders are generally useful only for fitting to artificially-constructed objects with complex structures for part/whole representations. Image processing requires edge detection, edge smoothing, segmentation into spectrally uniform regions using histogrammic methods, the extraction of objects using B-splines, and finally matching with a model with a similarity measure of feature vector sets. Typically, image acquisition time is the most significant temporal latency in the control of robotic devices. However, the Martian environment does not suffer from poor reflectance characteristics from water or misclassification of dense vegetation. Recognition is the final stage of image processing which requires both sensory data and *a priori* knowledge.

An edge detection operator such as the Sobel, Hough or Canny edge operator may be applied (the Canny edge operator can extract occluding boundaries). A 5x5 convolution mask can be applied to suppress noise and smooth the image. The Laplacian operator $I_{xx} = \frac{\partial^2 I}{\partial x^2}$ and $I_{yy} = \frac{\partial^2 I}{\partial y^2}$ can be approximated locally by the following 3x3 convolution mask:

$$\begin{vmatrix} -1 & -2 & -1 \\ -2 & 12 & -2 \\ -1 & -2 & -1 \end{vmatrix}$$

The image may be processed thus: $A = w \times \left(\frac{\nabla I}{\nabla I^T}\right)$ where w=Gaussian smoothing mask

The eigenvalues of A are proportional to the principal curvature so can be used to predict corners of A where the eigenvalues are high. This comprises 90% of image computation, so it would be best performed in parallel. Step edges in range data almost always correspond to depth discontinuities, i.e.. surface boundaries.

A visual discovery algorithm to extract different image intensities from the background based on the human visual cortical receptive field model using multi-scale Gabor filters has been implemented [Burl & Lucchetti 2000]. This has the advantage of detecting visually different objects from the background without *a priori* models. It is particularly suited to microscopic imaging. They used Gabor filters in conjunction with centre-surround and corner-sensitive filters. The centre-surround filter identified bright or dark localised regions that differ from their surroundings.

Vision involves sensing of electromagnetic radiation which, although it can occur across the electromagnetic spectrum, is restricted to infrared/ultraviolet ranges in biological organisms. A receptor cell containing the light-sensitive pigment rhodopsin detects light. The most common configuration is two eyes for binocular depth perception though not necessarily (spiders possess eight eyes but commonly two of the eyes are more developed) – carnivores employ binocular vision by mounting their eyes to the front of the head with a baseline but herbivores generally mount their eyes on the side of the head for wide field scanning. Insects generally possess compound eyes with multiple receptors of constant, low resolution. Experiments indicate that they employ optic flow for vision with high frequency resolution (~300 Hz compared to human sight which can detect frequencies of 20-30 Hz) [Srinivasan 2002]. This type of vision based on image flow is computationally much simpler than mammalian vision. Insect vision offers a generalised navigation reference system based on polarisation angle of solar flux for almost any rocky planet. Their eyes are static and do not move relative to the body as they are too close for stereo-imaging [Neumann 2002; Harrison 2002; Dickenson et al. 2002]. Chordates however have moveable eyes across two rotation axes which provides for foveated active vision. This allows enhanced resolution across the fovea but reduced resolution in the periphery to reduce to data flow to the brain. Animal behaviour varies in complexity, eg. a frog's behaviour is very simple – its reaction differs according to the visual stimuli it receives [Lettvin et al. 1959]. A dark object moving across the frog's eye will elicit a food catching response while a dark object moving above eye level elicits predatory avoidance behaviour. It appears from studies in frog vision that detection of motion is evolutionarily more primitive than detection of static objects [Lettvin et al. 1959]. Visual processes in frogs are independent of general illumination and involved image processing of:

- sharp edges and contrasts;
- edge convex curvature of objects;
- movement of edges;
- dimming caused by object extent.

Each fibre group from the retina mapped into parallel sheets in the frog's tectum, one sheet and group for each of four operations. The visual processing hardware connections were non-random and was geared towards "bug perception" for catching food items of a given size and motion.

An image is formed by the lens on the retina that is inverted and reversed on the retina. The first layer is that of the photoreceptive retina which detects light signals and transforms them into electrical signals. Photons are absorbed by the photopigment rhodopsin in the membranes of photoreceptor rod or cone cells of the retina to produce an electrical output. The retina comprises of 1.2 x 10^8 cone and rod cells. Individual neurons do not transmit large amounts of symbolic information. Input units are connected to the outside environment each of which responds to a narrow range of parameter values (receptive field). These input neurons act as selectively tuned filters to particular signal patterns. Different units simultaneously influence each other and mutually constrain each other as the output from other units affect the receiver unit according to the connection weights. Cells near the fovea have small receptive fields for precise determination while cells away from the fovea have larger receptive fields (particularly for night vision) [Hubel & Weisel 1979]. The second layer comprises two concentric layers of 1.2 x 10^6 retinal ganglion cells – on-centre and off-centre arrays with a centre-surround arrangement. The retinal cells sends impulses to the bipolar retinal cells and retinal ganglion cells. The retina performs preprocessing of the visual information. The retinal ganglion cells enhance relative illumination differences for greater contrast. Each cell is divided into circular symmetric excitatory centre and an inhibitory surround generating concentric receptive fields with illumination contrast comparison. Response fields overlap many neighbouring units whose peak sensitivities vary from neuron to neuron systematically. Local feature extraction is enabled by centre-surround oriented intensity discontinuities. The gains of these cells decrease exponentially with distance from the centre to act as filters to locate boundaries of objects in the visual field by comparing illumination levels between neighbouring regions of the visual field. Multiple layer parallel feed-forward networks with Hebb-type synaptic weight modification can generate the major features of the mammalian primary visual cortex. Feature analysing centre-surround cell and orientation-selective cell layers can emerge through self-organisation and become progressively organised in orientation columns through the network under random weight inputs, i.e.. without environmental input to the network [Linsker 1986 a,b,c]. Layer A comprised photoreceptor cells, Layer B of retinal cells, Layers C to f represented circularly symmetric cells (retinal ganglion) and layer G represented orientation-specific cells of the primary cortex. When connected with excitatory lateral connections, layer G yielded columnar registration of its orientation-specific cells. The ~10^6 axons of the retinal ganglion cells converge into the fibre optic bundle and output along the optic nerve. Signals representing each half of the visual field from each eye cross over to the opposite hemisphere of the brain at the optic chiasm. From there, the majority of the retinal ganglion cells terminate on the neurons of the left and right lateral geniculate nucleus of the thalamus either side of the brain stem. Some project into the superior collicus of the midbrain before projecting into the posterior lateral nucleus of the thalamus, and then onto the visual cortex including those areas (V1 to V3) which are devoted to extracting edge and line features. Some of the fibres bypass the visual cortex projecting into the oculomotor nucleus and then onto the iris – this is the basis of the pupillary light reflex. There are additional reflexes of the eye with a similar architecture – the optokinetic reflex which comprises afferent pathways from the retina to the visual cortex onto the superior collicus from which fibres project into the motor nuclei of cranial nerves 3, 4 and 5. The thalamus is involved in the focus of attention and acts as a relay station for all sensory information except olfaction. It transforms all sensory signals into a code suitable for the cerebral cortex. The thalamus sends the signal along the lateral geniculate axons to the striate or visual cortex (Brockmann's area 17) on the occipital lobe. Hence, there are two visual channels ascending to the cerebral cortex. The visual cortex comprises several layers II, III, IV, V and VI which are subdivided into a large number of vertically arranged columns of neurons. The retinal connections via the lateral geniculate nucleus connect into the visual cortex (at area IV). The retinal connections to the visual cortex are topographically arranged, bearing direct correspondence to the receptive fields such that neighbouring cells have neighbouring receptive fields. The cortical neurons differentially respond to directed line and edge segments at specific orientations. Simple cells in the upper part of layer IV respond to specific positions and orientations of lines while circularly symmetric cells reside in the lower part of layer IV. Layers II, III, V and VI comprise of complex cells which can also detect orientation-specific lines but are less position-specific. Layer III responds only to line orientation. Layer V responds to motion and direction. Layers I and II coordinate these parallel systems. This is the first step towards extracting features. The visual cortex outputs to areas deep in the brain and to nearby parts of the cortex. All thalamocortical projections are reciprocated – Layer VI of the visual cortex projects feedback fibres back from the visual cortex to the lateral geniculate nucleus, Layer V to the superior collicus in the midbrain, Layers II and III to other parts of the cortex, for combination of the receptive fields into more complex fields for higher levels of image recepresentation. Similarly, the auditory cortex projects fibres back to the medial geniculate body, the somatosensory cortex projects fibres back to the ventral nucleus. Hence, the state of the cortex influences processing in the thalamus through feedback mechanisms. From the sensory cortex, fibres project to adjoining sequences of the association cortex which represent the largest fraction of the cortex for more advanced processing including the fusion of auditory and visual information. The final destination of such processing lies in the hippocampus and amygdala. The amygdala is a nucleus near the tip of the hippocampus which receives input from the olfactory tract and outputs to the hypothalamus. There are two cortical regions controlling gaze in each cortical hemispheres – the premotor cortex controls scanning and roving eye movements; the visual cortex controls object tracking, pursuit and saccade reflexes. Observer motion is compensated for through vestibular and proprioceptive feedback from the cerebellum.

CCD cameras based on pixellated pin-hole optics with a lens and retina are biomimetic in design. Vision processing traditionally involves feature extraction through a series of processing stages – Laplacian-of-Gaussian filtering, edge detection, region-growing, image segmentation, and model-matching with stored prototypes. The biological weighting function w(x,t) is a spatial-temporal function which in vision processing may be represented by a Gaussian, differences of Gaussian or Gabor function. In optical flow imaging, the receptive field defined by the excitation of the neuron given a stimulus is given by:

$$y = \int w(x).v(x).dx$$

where:

v(x)=image flow velocity input stimulus
w(x)=expected egomotion velocity field

In addition, stereopsis attempts to find a function δ for which condition holds:

$$I_l(x + \tfrac{1}{2}\delta(x)) = I_r(x - \tfrac{1}{2}\delta(x))$$

where:

δ(x)=disparity map at point x

In polar coordinates, for an agent moving at speed $\begin{pmatrix} v \\ w \end{pmatrix}$, the rate of change of position is given by [Mallot 1997]:

$$\begin{pmatrix} \dot{r} \\ \dot{\phi} \end{pmatrix} = \begin{pmatrix} -v\cos\phi \\ \frac{v}{r}\sin\phi + w \end{pmatrix}$$

These processes are computationally intensive. The chief difficulty in vision research is in image processing for the extraction of lines and edges for object recognition, particularly in natural scenes. Many insects such as bees employ near complete 360^o imaging of the environment which has been emulated in Franz & Mallot's (2000) design of a conical mirror imager to provide omnidirectional ring-shaped imaging field without camera movement.

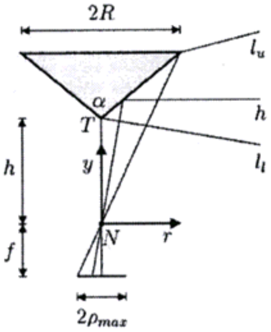

camera parameters:
R upper cone radius
α cone angle
h distance of T and N
f focal length
$\rho_{max} = fR/(h + R\cot\frac{\alpha}{2})$

projection in cylinder coordinates, (r, y),
for $r \geq 0$, $r\cot\alpha \leq y - h \leq r\cot\frac{\alpha}{2}$:
$(r, y) \mapsto \rho$

$$\rho = -f\frac{-r\cos\alpha + (y-h)\sin\alpha}{r\sin\alpha + (y-h)\cos\alpha + h}$$

Figure 187 – Imaging geometry for a conical mirror imager (Franz & Mallot 2000)

Of great interest is active (foveated) vision which emulates foveated vision with automatic pan and tilt feedback control to provide gaze stabilisation, smooth pursuit, optokinetic nystagmus (OKN) and vestibular-ocular reflexes (VOR) [Ballard 1991; Landolt 2002]. We are interested in particular in foveated vision which reduces the image processing requirements of robotic vision systems. Active vision makes the observer an active participant in the world to be analysed to simplify processing. Saccades are rapid jumps in eye position (limited to ~6 saccades/s) which re-position the foveated gaze to different parts of the visual field. Smooth pursuit scans moving objects and is limited to ~100^o/s. Control of foveated visual gaze through the vestibular-ocular reflex (VOR) reduces the requirements for image processing. Organisms are not passive receivers of sensory data but actively explore and manipulate the environment and this requires the imposition of constraints to provide environmental context. Only a small number of features in the environment that are directly relevant to the task are registered, vastly decreasing the amount of image processing required as receptive field registration rather recognition is required. A frog's behaviour is very simple – its reaction differs according to the visual stimuli it receives [Lettvin et al. 1959]. A dark object moving across the frog's eye will elicit a food catching response while a dark object moving above eye level elicits predatory avoidance behaviour. The vestibular-ocular reflex (VOR) to track moving objects with smooth pursuit. The human eye is characterised by a large concentration of ~10^5 optical fibres in $\pm 0.5 - 1^o$ foveal region of the retina along the optic axis with high resolution. Away from the fovea the distribution of optical fibres becomes more diffuse and so of lower resolution through the majority of the visual field. Active control of attention and the non-uniform resolution provides orders of magnitude improvement in image processing requirements, hence the element density of pixels of the

CCD may be patterned after the human eye with high resolution FOV of 1-2^{o} and low peripheral resolution view of ~100^{o}. Gaze shifting and gaze holding movements of the eye provide a means to rapidly aim this narrow FOV. Fast saccadic eye movements move the fovea between different targets in the visual field while active gaze control keeps the fovea on a specific target. The retina and the regions of the brain that control the 6 muscles of the eyeball in pairs corresponding to the 3 axes of rotation are connected in the optokinetic response neurons in a negative feedback loop. There is also a feed-forward vestibulo-ocular reflex which uses the vestibular apparatus of the inner ear to detect movement of the head which responds much faster than the optokinetic feedback loop. A similar kind of control system architecture would be required for artificial gazing. Ballard (1991) advocated an animate vision paradigm based on this function. It would require an invariant object centred frame of reference (eg. the local inertial frame). Two cameras with foveated vision which are able to pan and tilt in synchronicity, each with independent saccading ability at 3 Hz and convergence control in the yaw axis could be used. A vestibular-ocular reflex with smooth pursuit of targets would provide robust static and moving target tracking. At any instant only a small number of features in the environment that are relevant to the particular task are registered and new features are incrementally registered only when required. This approach to vision would dramatically decrease the requirement for data transmission at any instant to the ground, but it is not clear if this approach would be suitable for teleoperation or even virtual reality methods as such gaze shifting and gaze holding would require tight slaving of the remote cameras to the teleoperator eye movements which average saccade movements 4-5 times per second – the remote camera inertia and articulation motors could not match those of the human eyeball. Active vision is required if visual fixation is to be maintained onto a moving target – human vision uses a number of mechanisms such as saccades, smooth pursuit and optokinetic reflex (OKR). A high resolution fovea with a centre-surround variation in resolution offers considerable reduction in image processing requirements.

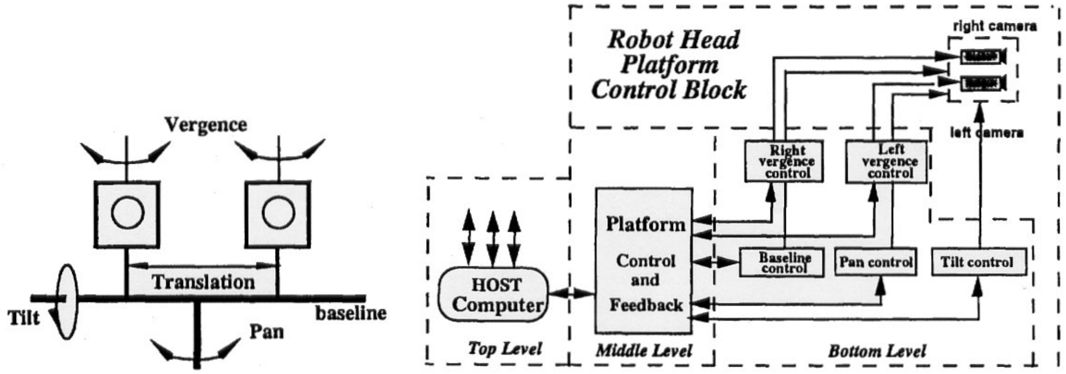

Figure 188 – a) Vergence, translation, pan and tilt degrees of freedom in an active vision system; b) vergence control system for active vision (Marefat & Wu 1996)

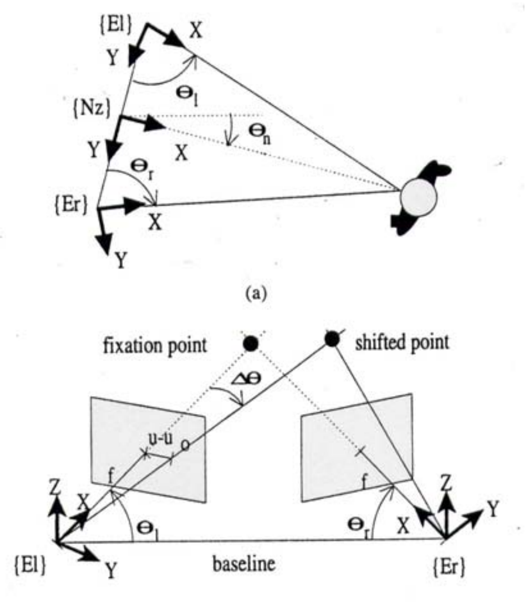

Figure 189 – Camera vergence angle control (Dias et al. 1998)

Optical flow field methods provide the means for motion-detection [Giaccone 1998]. As a robotic rover is a moving vehicle, detection of motion allows the determination of the motion of the vehicle. Optical flow methods generate a vector field representing the motion of the pixels of an image. Associated with each pixel is a velocity vector and the vector field of all pixels of the image form the optical flowfield. Pre-smoothing the images with a Gaussian filter to reduce aliasing introiduces blurring. Continuity of the motion field is ensured by applying the constant brightness constraint which states that a moving pixel's brightness is constant, or at least almost constant: $I_t(x) - I_{t-1}(x) + \nabla I^T u_t(x) = \varepsilon$ where ε is a small threshold. It is assumed that apparent 2D motion of objects in an image sequence is represented as an affine model:

$$u(x) = \begin{pmatrix} u(x) \\ v(x) \end{pmatrix} = \begin{pmatrix} a_0 + a_1 x + a_2 y \\ a_3 + a_4 x + a_5 y \end{pmatrix} = \begin{pmatrix} 1 & x & y & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & x & y \end{pmatrix} \begin{pmatrix} a_0 \\ \ldots \\ a_5 \end{pmatrix}$$

The parameters a are coefficients of affine functions of pixel coordinates. This equation is then embedded into the constant brightness constraint. It is not clear if motion-detection is necessary as the speed of the rover will be low, though it does add further capabilities in support of autonomous navigation. Optic flow is based on insect visual navigation such as the fly.

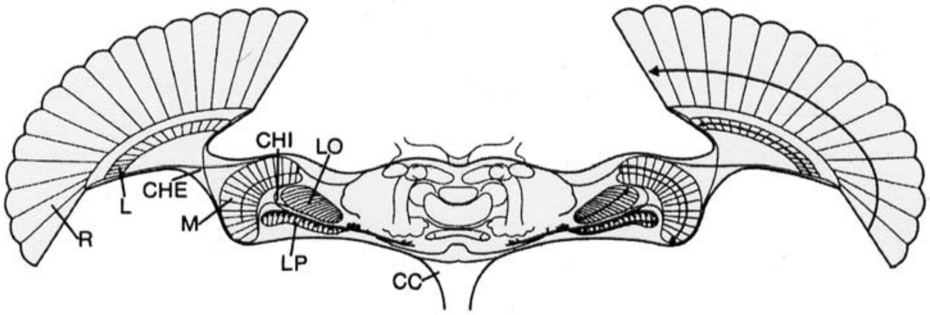

Figure 190 – Fly brain with large compound eye – retina R, lamina L, medulla M, lobula LO, lobula plate LP and cervical connective CC are connected thus: L and M through external chiasm and M, LO, LP through internal chiasm (Huber et al. 1999)

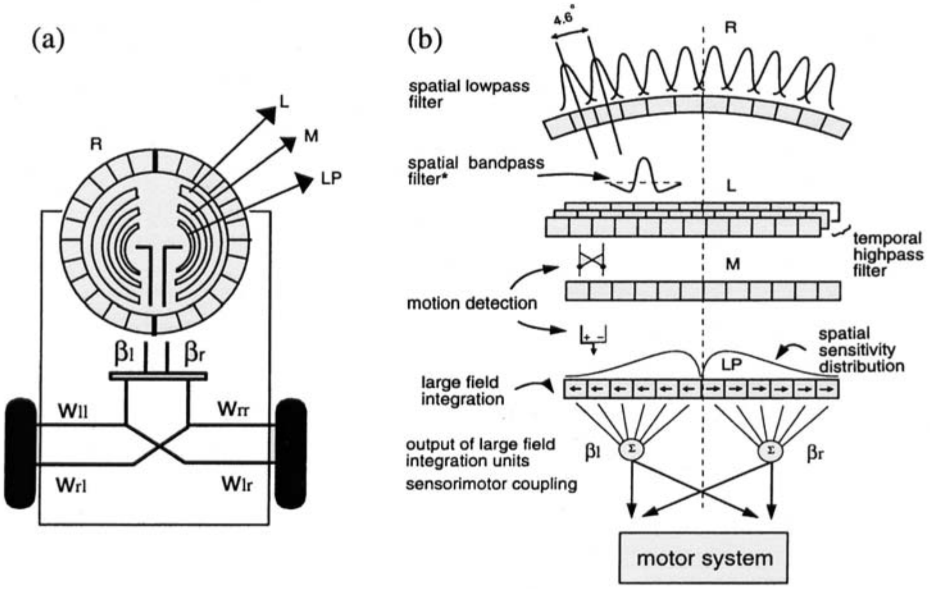

Figure 191 – a) Simplified model of three layers of fly visual system including lamina, medulla, lobula plate and transmission weights w_x coupling the visual outputs β_x to the motor system; b) model of visuomotor controller with same functional processes (Huber et al. 1999)

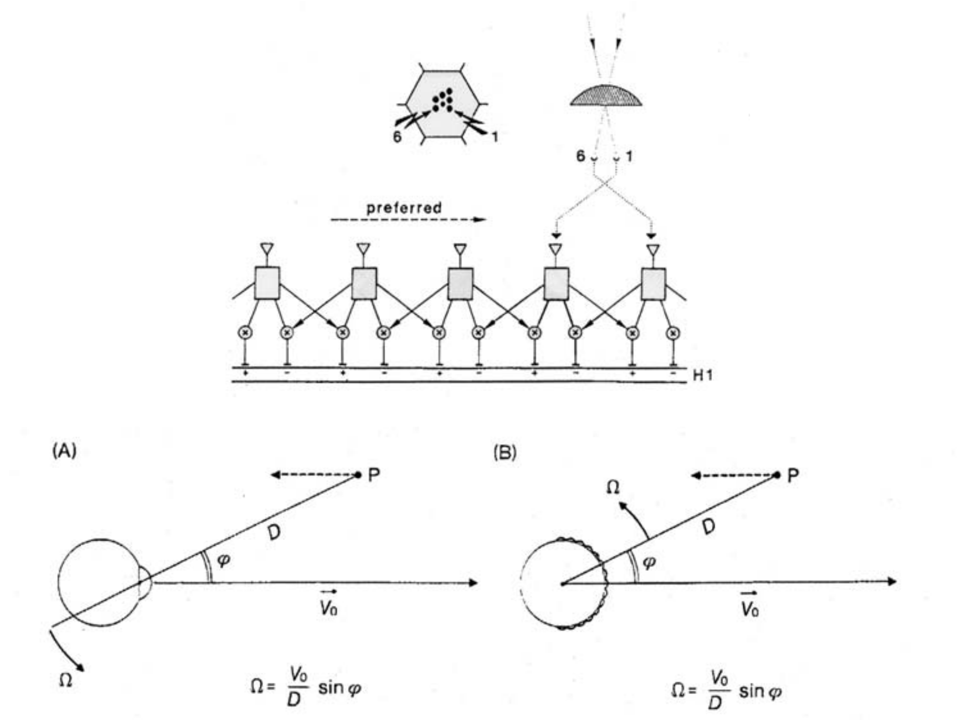

Figure 192 – a) Basic topology of lateral interactions leading to directionally selective motion detection in H1 neuron of the fly – each puppet drives H1 with opposite polarities with two legs whose transmittance is controlled by signals from adjacent puppets on each side; b) principle of motion parallax so a mobile agent can determine its distance to a point P at azimuth ϕ based on P's measured angular speed Ω across the visual field (Fransceschini et al. 1992)

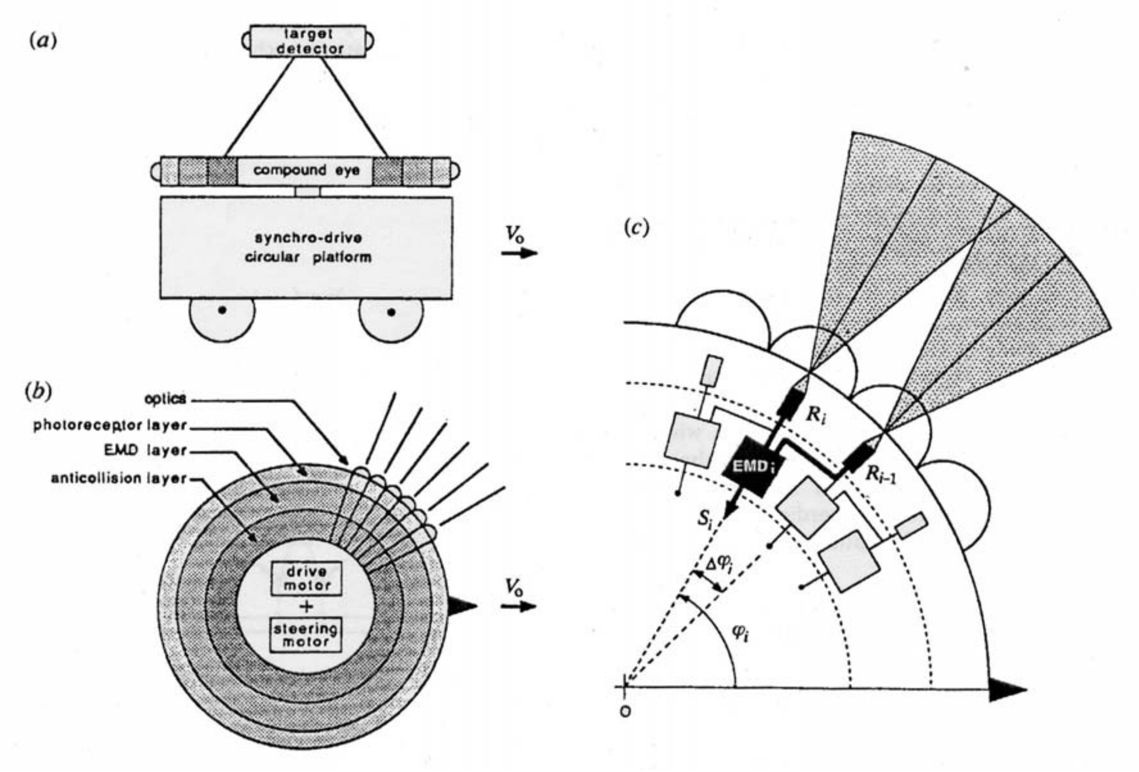

Figure 193 – Mobile robot simulator with panoramic compound eye based on layered columnar organisation similar to insect optic ganglia for collision avoidance and accessory vision system views horizon to provide orientation while compound eye. Collision avoidance is based on motion parallax determined through array of elementary motion detectors (Franceschini et al. 1992).

Kuperstein et al (1987) developed a neural network model of adaptive visual-motor coordination to demonstrate visually-guided behaviours. This involves the process of mapping directly between a sensory input and a motor output through a topological map stored in the connection weights of the network. A neural network was used to adaptively control a visually guided robot manipulator arm using pattern recognition. It essentially learned an adaptive camera-position topographic model of the environment from an egocentric frame of reference based on sensory-motor reactions over a sequence of trials of different movements. The manipulator was controlled using consistency between the signals used to drive it and the signals received on sensing the results of the movement. The robot learned to correlate visual end effector positions from stereo cameras with joint motor activity signals producing those end effector positions, i.e.. joint angle to end effector position mappings. This provided visually guided behaviours.

Humanoid robots must learn within human social environments in order to develop human-level intelligence [Brooks 1996]. Such social interaction will require VOR, OKN and other active orienting and posture control capabilities. Cog is a humanoid-type robot which exploits active vision to generate human-like behaviour [Brooks 1997; Brooks & Stein 1994]. It is a mechanical torso with two peripheral vision cameras and two high resolution cameras which can rotate independently. Its microphone allows it to respond to noise. The torso has three degree of freedom hips; the neck has three degrees of freedom for the head. The eyes have two degrees of freedom (pan and tilt) and a vestibular system of three orthogonal gyros and two inclinometers. Cog integrates active vision, sound processing, dextrous manipulation and aspects of human cognition through a large scale parallel MIMD computer. The key process is that Cog will learn abstract concepts through bodily experiences. Cog comprises a head with two compliant manipulators but no mobility. Cog's control system is based on hierarchies of reactive behaviours to develop the rudiments of social interaction by focussing its attention on tracking objects. It has a gyroscopically-controlled head mounted onto a position-sensing neck to exploit balance. Due to the great advantages in active vision, this has major applications in space exploration in providing automatic slewing for target acquisition and tracking while offering reduced data transmission requirements.

3.12.2. Vibration Sensing

Vibration sensing through the use of hair-like mechanical sensors (which sense pressure through cantilever operation) are ubiquitous in the biological world.

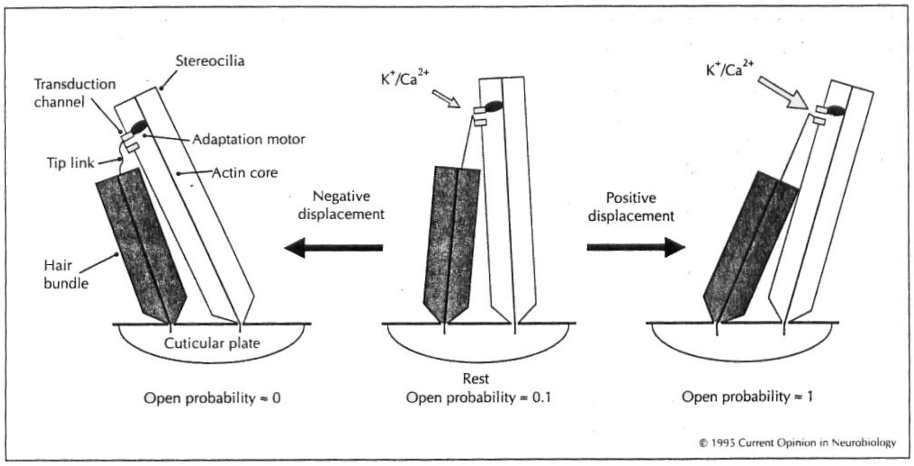

© 1995 Current Opinion in Neurobiology

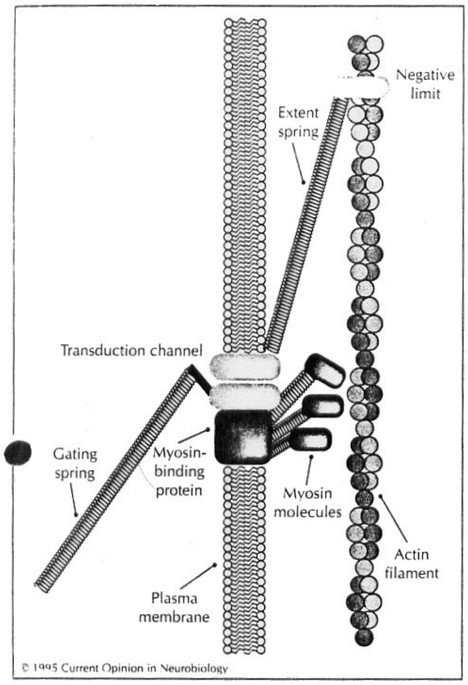

© 1995 Current Opinion in Neurobiology

Figure 194 – a) Gatings-spring model of the human auditory receptors – a mechanical stimulus deflects a hair bundle towards the taller sterocilia generating stretch in the tip links between adjacent sterocilia; b) hair receptor transduction mechanism – myosin molecules act as adaptation motor in parallel with the extended spring which slips down thje actin cytoskeleton (Gillespie 1995)

Vibration sensing, or more commonly, hearing is generally achieved through mechanical transduction receptors in the form of hair cells (Figure 194). Mammalian evolution has added a cochlea for hearing to the vestibular system, so there are two divisions to the 8th cranial nerve – the vestibular and acoustic pathways. The cochlea receives sound vibrations through the perilymph of its spiral chambers. The organ of Corti perform audio-frequency analysis. Fibres synapse with the hair cells of the organ of Corti which pass as the cochlear nerve into the medulla to synapse on the cochlear nuclei at the junction with the pons. Some fibres from the cochlear nuclei pass to the nucleus of the abducent nerve, and onto the nucleus of the oculomotor nerve to provide conjugate horizontal eye movement in response to sound. Other fibres from the cochlear nuclei pass to the medial geniculate body of the thalamus which relay the signals to the primary auditory cortex (area 41). The cochlear of the inner ear is a transducer which converts mechanical vibration into an electrical output. Receptor cells in the cochlea elongate due to changes in the electrical fields which enhance low-amplitude sound. The primary mammalian hearing receptors are hair cells which line the inside surface of a fluid-filled cavity, the labyrinth of the inner ear. Bending of the hair cell generates an action potential. Sound is conveyed into the ear through the pinna to excite the tympanic membrane. The tympanic membrane transmits its vibration to the three bones of the middle ear (hammer, anvil and stirrup) to the inner ear. Much of our knowledge of the function of the cochlear comes from oto-acoustic emissions. In insects, hearing sensors are located on the body, eg. on the legs of the cricket. Some animals utilise hearing in an active form – bats use echolocation by emitting high frequency pulses of sonar ~30 kHz+ through the mouth. The reflection of the pulse from objects is detected and processed to provide spatial information similar to vision. The moustached bat uses pulses which have a constant frequency component (for doppler shift measurements) and a downward frequency modulated sweep (for distance and texture measurements). Although microphones possess a vibrating diaphragm similar to an ear drum, micromachined hair cells have been suggested for mechanotransduction [Li et al 2002]. Echolocation by sonar has long been of interest [Carmena & Hallam 2001]. Sonar cannot be used in a vacuum though radar and lidar are suitable electromagnetic versions of such active techniques.

Bat sonar is a sophisticated echo-locating system that operates by doppler shifts provide information on the relative velocities of objects such as flying insects. The amplitude of the echo combined with the delay provides information on the size of the target. Amplitudes of the component frequencies correspond to the size of a number of features of the target. The differences between the ears in intensity and arrival time of sound give the target azimuth while the interference pattern created by sound waves reflected gives the elevation.

3.12.3. Vestibular Sensing

Attitude sensing is based on the vestibular system which resides in the middle ear. The vestibular system is the most ancient sensory system in providing balance equilibrium and orientation. Indeed, vestibular sensing is evolutionarily more fundamental than hearing deriving from the lateral line in fish. The vestibular system bears similarities with the acoustic system, particularly the transduction mechanism. The labyrinth of the inner ear contains two chambers – the saccule and the utricle. The vestibular sensors comprise of hair cells within the utricle and saccule and hair cells in the ampullae of the three orthogonal semicircular canals. The horizontal canal lies at an angle of 30^o away from the horizontal. Three perpendicular semicircular canals bud off from the utricle which are filled with endolymphatic fluid. Hair cells at the cupula bend in response to the motion of the fluid in the three axes of the semicircular canals – this is the Barany model. This provides the basis for the detection of rotation and translation of the head. Mammalian evolution has added a cochlea for hearing to this system, so there are two divisions to the 8th cranial nerve – the vestibular and acoustic pathways.

Afferent fibres from the 8th nerve provide the basis for coordination of eye movements, and head, neck and body posture. Axons of the vestibular ganglion follow the cochlear nerve to the medulla terminating on the four vestibular nuclei. Some fibres pass directly to the cerebellum which connect to the fastigiobulbar tract which exit the cerebellum and returns to all four vestibular nuclei in the brainstem. Two vestibulospinal tracts emerge from the vestibular nuclei which pass down the spinal cord to the muscles of the limbs for posture maintenance. Fibres from two of the vestibular nuclei form the vestibulomesencephalic tract to the motor nuclei of the oculomotor nerve which control the ocular muscles of the eyes. This enables the eyes to remain fixed on objects while the head and body move. In the human brain, vestibular data is integrated with vision to provide coordination of balance and active visual slewing. The coupling of vestibular sensing to vision is thus a fundamental property of animal sensing whereby sensing and actuation are blurred – this is a component of active vision. This approach to active sensing may be generalised to other sensors and typically must include some form of sensor fusion.

3.12.4. Tactile Sensing

Touch is closely related to proprioception in some respects. Sensory feedback is provided by the proprioceptive sensors in the muscles, joints and skin while exteroceptive sensors provide direct observation feedback (vision, etc). Reafferent connections allow comparison between command and actual performance for control. Proprioceptive awareness of the body is central to controlling the agent in relation to the outside world – it may even be central to consciousness in representing the agent in relation to the world. The body image is used to predict and control the outcomes of behaviour. The self-model provides a reference frame from which to integrate sensorimotor transformations. Touch is associated with mechanical stress or pressure but is a highly complex sense in its most general form as it is often associated with sensitivity to temperature, pain, moisture, etc. Touch is a distributed sensing modality which imposes the requirement for significant local processing to relieve bus bandwidth requirements. Touch sensing is of primary importance possibly more important than vision for fine manipulation and may be regarded as a form of refined force sensing [Paul 1987]. Touch is an active, exploratory process in which sensing and actuation are intimately linked such that the actuation process through movement of the hand and fingers define the sensory data directly. Indeed, blind people use touch extensively and can survive adequately in the world - it is debatable whether a man devoid of touch could do so well. Taction does not suffer from inoperability in darkness, or from problems due to shadows, lighting constraints, colour or obscuration though it does suffer from restricted views of the object. The distinction between touch as a sensing modality and its utility as an integral part of actuation is blurred. Tactile sensibility is provided by two different sensors:

1. simple touch for crude localisation;
2. tactile discriminators to determine pressure, fine localisation, and size/shape perception.

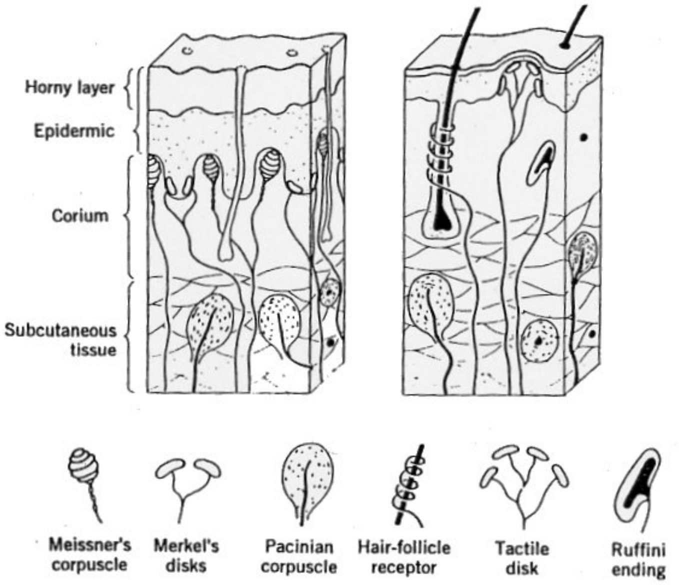

Figure 195 – Structure and location of tactile receptors in the skin

Tactile discrimination data follows the same pathway as proprioception. Behind the motor cortex of the human brain lies the primary somatosensory cortex (area 3) which receives signals from the skin, bones, joints and muscles from all regions of the body via the ventral nucleus of the thalamus. It receives proprioceptive feedback from the muscles and joints. Both of these cortical regions have regions which correspond to the parts of the body forming a cognitive map of the human body (homunculus map).

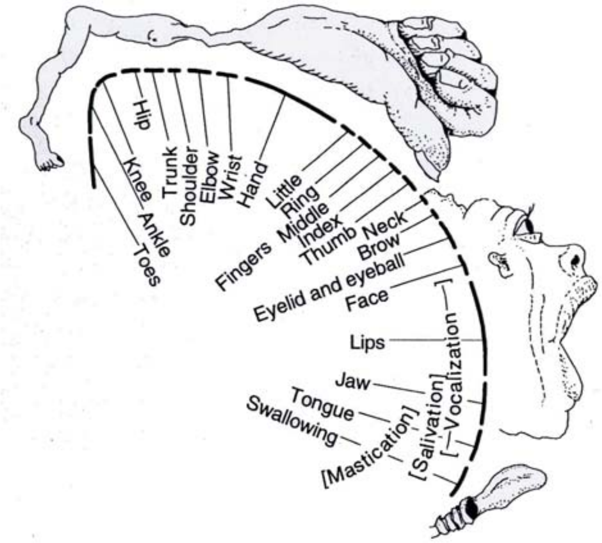

Figure 196 – Homunculus representation of the motor cortex

The distortions in the homuniculi relate to the precision to which the body part is controlled. The motor cortex represents the body topographically from the bottom to the top of the cortex: larynx, tongue, face, thumb, hand, forearm, arm, thorax, abdomen, thigh, calf, foot and perineum muscles – the areas for the hand, tongue and larynx are disproportionately large. The mapping between the body and cortex is distorted since the area of the cortex is reflected of the sensitivity and degree of fine control of that body part. The parietal cortex adjoining the somatosensory cortex is characterised by multiple pathways to other cortical areas. Sensory feedback is provided by the proprioceptive sensors in the muscles, joints and skin while exteroceptive sensors provide direct observation feedback (vision, etc). Reafferent connections allow comparison between command and actual performance for control. Proprioceptive awareness of the body is central to controlling the agent in relation to the outside world – it may even be central toconsciousness in representing the agent in relation to the world. The body image is used to predict and control the outcomes of behaviour. The self-model provides a reference frame from which to integrate sensorimotor transformations.

Tactile sensibility is provided by two different sensors:

1. simple touch for crude localisation;
2. tactile discriminators to determine pressure, fine localisation, and size/shape perception.

Tactile discrimination data follows the same pathway as proprioception. Simple touch passes through the spinal cord through the spinothalamic tract to the thalamus, from which thalamocortical fibres relay the signal to the postcentral gyrus of the parietal cortex. Pain sensors are naked terminals of the branching nerve network – they enter the spinal cord and fibres of the lateral spinothalamic tract convey the signals directly to the thalamus (where pain recognition occurs). Thalamocortical fibres relay these signals to the postcentral gyrus of the parietal cortex (for integration with other sensory stimuli). Cold receptors and warm receptors which are distributed across all body surfaces both follow the same pathway as pain.

Sensory pathways separate into pain and temperature fibres and position and touch fibres after entering the spinal cord through the dorsal root. Both sets of fibres reach the lower medulla, and then onto the thalamus with pain pathways being very extensive. Tactile sensing is particularly important in grasping manipulation through modification of impedance of the hand. Grasping is a fundamental component of manipulation involving interaction with a high variation in impedances from zero (prior to contact) to high impedance if the payload is stiff. Internal models are likely to become complex in certain tasks which involve rapidly changing motion and/or forces. In such situations, low impedance is usually adopted so that the limb can "give" to unexpected disturbances so that large forces are not generated. Once a greater area of contact is obtained, greater forces can be exerted by stiffening impedance of the grip. Preflexes maintain stability while minimising the effects of disturbances. Preflexes are often implemented – preflexes are near zero-order responses due to the intrinsic, non-linear properties of the structure whereby disturbances are rejected within ~20-70ms without the time delays ~200-500ms inherent in feedback control from proprioception. Changing compliance properties of muscles through sliding actin-myosin protein fibres act to stabilise locomotion – contraction is caused by the sliding motion of the head section of myosin along the actin molecule. Myosin acts as a molecular motor whereby ATP-generated energy provides for the cycle of power stroke followed by a recovery stroke. Each myosin molecule thus "walks". The muscles impart spring-like properties to alter the compliance to stabilise the reflex system. Muscle-based actuators offer the required dynamics for biomimetic robotics. They comprise of two fibre bundles – one agonistic contractile fibre, one antagonistic elastic fibre tendon - to provide the required tendon stiffness for tension and the required force generation.

Some types of sensors and actuators can be synthesised within the same device for both sensing and actuation with the advantage of co-located sensor and actuator. Muscle has this property for tactile sensing. Piezoelectric materials in particular have been deployed in this fashion by exploiting self-sensing actuation – a reference capacitor estimates the electrical properties to be subtracted from the piezoelectric signal (determined by both mechanical and electrical effects) to yield mechanical strain. An optimised electrostrictive elastomer and piezoelectric polymer combination offers a piezoelectric sensor with electrostrictive actuation capabilities [Bar-Cohen 2002]. Together with the proprioceptive sensors they provide the means for servo-level control feedback information. In the animal muscular system, sensori-actuator functions are provided by muscle spindles, which operate similarly to a piezoelectric actuator. The human touch receptors which reside at the cutaneous sections of the skin are sensitive to intensity, velocity and acceleration. Subcutaneous receptors are proprioceptive receptors often associated with muscles to detect body posture and movement – these sensors are correlated with data from the vestibular system for providing balance primarily in the cerebellum. Biological touch sensors are often hair cells, stress sensors (campaniform sensors in insects which reside at the base of wings) or vibration sensors (chordotonal sensors in insects analogous to hearing sensors). The primary sensitivity to touch is pressure for which there are several alternatives. Most force sensors are thin film strain gauges, piezoelectric sensors or capacitative array sensors. Carbon-impregnated polymers offer variable resistance depending on applied mechanical pressure and are used as tactile sensing "skins" in robotics. Piezoelectric sensors are the primary sensors used for detecting pressure, but micromechanical cantilever sensors may also be utilised to emulate hair cells. Electroactive polymers may be used to convert mechanical strain to electric current thereby acting as sensors. The most biomimetic sensors are MEMS-based cantilever sensors which detect bending similar to hair cells. The eye-in-hand configuration of laser rangefinders is the current approach to providing "tactile-type" imaging. Biological mechanotransduction is performed by connective tissue, fibrillar collagens in particular [Lambert et al 1998]. Tactile sensors may take the form of an "artificial skin" distributed as arrays of pressure-sensitive pads over the robot links on a thin, flexible compliant substrate for continuous tactile sensing in all directions. The most common approach to tactile sensing is through the use of piezoelectric sensors, though Rossi et al (1988) suggest the use of polymeric gels sandwiched within Ag/AgCl thin film with mylar outer layers and Ag/AgCl mesh with □odeli backing to mimic unaxially arranged keratin-type α-helical fibrils in the dermis and randomly-oriented collagen in the epidermis which provides a complex piezoelectric function. A simpler versuion of this would be a tactile array embedded in a compliant elastic medium with six PVDF piezoelectric elements with different orientations. The transducers must be durable and wear-resistant particularly to slip friction. Harmon (1982) discussed the most desirable properties for touch sensors. Fast transduction and response times are requires ~1-10 ms compared with human reaction times of ~200 ms. Resolution requires 1-2 mm spacing between sensors corresponding to 5 x 10 to 10 x 20 elementary points for a finger tip – this brackets the human fingertip characteristic of 10 x 15 points. The dynamic range required is from 1 g to 1 kg with a sensitivity of 1 g. Finally, hysteresis should be negligible.

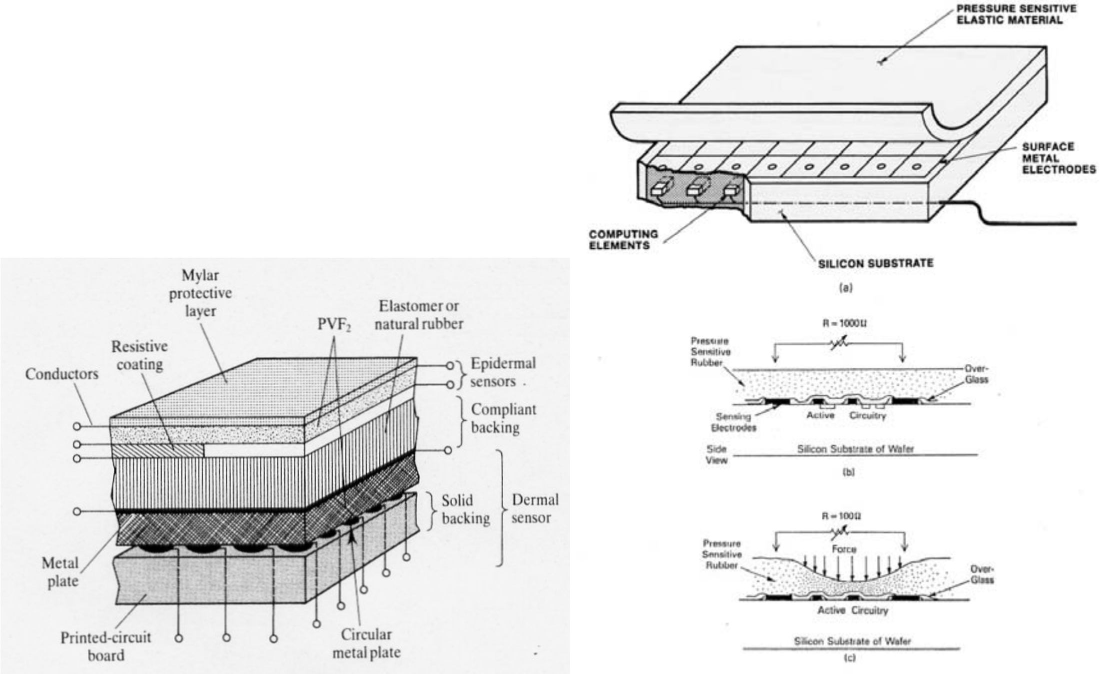

Figure 197 – a) PVF2 tactile sensor; b) Conductive elastomer tactile sensor (Fearing 1990)

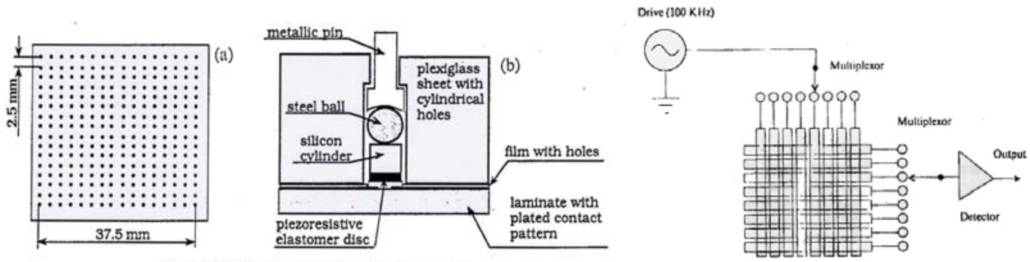

Fig. 5. (a) Tactile sensor matrix. (b) Tactile element.

Figure 198 – a) Tactile sensor matrix; b) Taxel (Odeburg 1995); c) Capacitative Sensor Configuration (Fearing 1990)

3.12.4.1.1. Biomimetic Flow and Contact/Bending MEMS Sensors

Contact and bending sensors are to be mounted on the lobster robot's flexible antennae in order to provide it with a sense of touch. The MEMS flow sensors will allow the lobster to sense water currents and thus to adapt its behaviour accordingly. The MEMS sensors developed here act as simple on-off switches, thereby reducing cost, power consumptions and simplicity in the integration of the electronics.

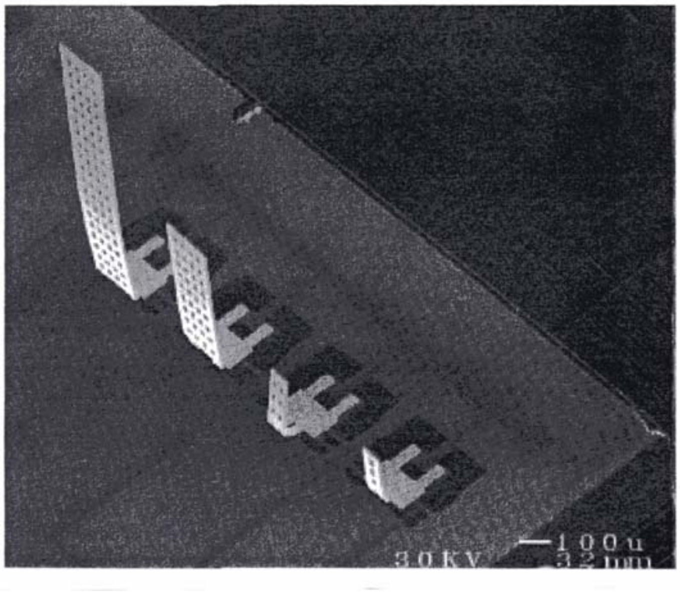

Figure 199 – Paddle flow sensors

The initial rationale for developing an underwater system like this is that there are few devices in existence that can detect sea mines. The ability to detect and destroy sea mines would benefit both military and civilian populations, and an inexpensive ambulatory robot is being designed to reduce the risk to human life.

A lobster has become the biological baseline for this design, as they have survived over hundreds of thousands of years to survive in this aquatic environment. Both the robot's design and behaviour are patterned after the lobster. Similar to a real lobster searching for and recognizing prey, the lobster will search for and classify objects as mines. To do so, the lobster needs to have artificial senses to explore with its antennae, fell the ground that it is walking on, sense if it is tilted or flipped over, and sense extreme currents in the ocean and make appropriate adjustments. A real lobster feels the bending in its antennae as well as detecting flows through the bending of small hairs located in the pits of the claws. Through the concept of microelectromechanical systems (MEMS) sensors, the robot will be able to "feel".

Existing bending sensors can be categorized as thin-film metal foil strain sensors, piezoresistive sensors, or capacitive array sensors. With a conventional Wheatstone bridge, each of the above can be used to measure strain. The accuracy and ability to measure a continuous range of strain are the advantages of these sensors, however their main disadvantage are that they need a continuous source of electrical power and the information collected requires processing due to the continuous range of sensitivities supplied. However, the bending sensors developed for the lobster antennae closely resemble switches rather than traditional gauge sensors. They do not draw current until the antenna is bent to a radius sufficient to activate the switch. A response can be triggered when a particular switch is activated.

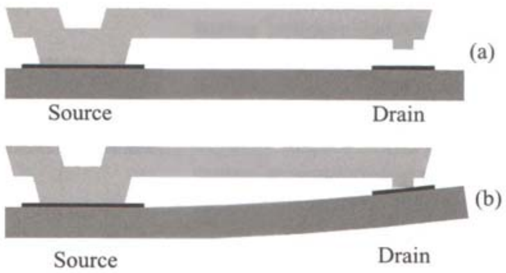

Figure 200 – Standard Sensor with source and drain a) open/free and b) closed/contact.

As shown in Figure 200, the contact/bending sensor is a gold cantilever beam that is surface micromachined on a very flexible and thin silicon substrate, thereby making it possible to design a switch that closes when the silicon substrate deflects to a certain curvature.

However, wouldn't sufficient water flow also cause antenna bending? Analysis has shown that the taper of the antennae that causes its deflected shape due to contact differs significantly from its deflected shape due to flow.

Radius of curvature calculations

The first step in designing both contact sensors is to develop the relationship between the length of the switch and the radius of curvature that it will detect. The relationship between switch length and sensitivity is:

$$R^2 = L^2 + Y^2$$

where:

R = radius of curvature
L = length of switch
Y = vertical location of the switch tip from the centre of the circle

Rearranging in to find Y:

$$Y = R - h$$

where:

h = switch gap = 0.25 μm = 2500 Angstroms

Solving for L gives:

$$L = \sqrt{2Rh}$$ (neglecting h^2)

Theoretical calculations for the cantilever design were based entirely on small deflection Euler-Bernoulli beam theory. The deflection of the end of a cantilever beam with an end load is:

$$h = \frac{PL^3}{3EI}$$

where:

h = *switch gap*
$P = 100\mu N$ *stiction load*
L = *length of switch*
E = *elastic modulus of the beam*
I = *area moment given by:* $I = \frac{1}{12} wt^3$
w & t = *width and thickness of the beam*

The standard cantilever sensor is limited to a length of less than 200 μm, which corresponds to a radius of curvature of less than 80 mm. After testing the constructed stopper sensors, it was deduced that the switch gap was much larger than 2500 Å, resulting in the switches closing at a smaller radius than theoretically calculated.

Flow sensors

There are several existing technologies that are used as flow sensors. One is based on a device that generates heat and two sensors measure the temperatures upstream and downstream of the source, allowing for asymmetry in the readings to be used to determine flow rate. Another design utilizes a flexible membrane that is deflected by the flow, allowing a change in the flow to be detected by a change in the capacitance between two electrodes. A third sensor uses a piezoresistor to convert fluid flow to electrical information. These all require more wiring and power than desirable on the lobster robot, and are also a step away from the goal of making the robot biomimetic. Therefore, small 1-bit on-off switches are used to simulate a hair sensing water flow around it, as shown in Figure 201.

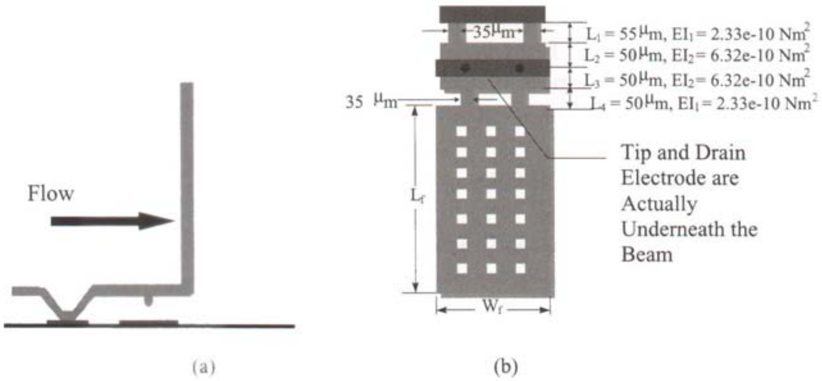

Figure 201 – MEMS 1-bit switch a) profile view after bending and b) head-on view before bending.

Four switches were developed to detect different flow speeds, specifically at flow speed sensitivities of 0.5, 1.0, 3.4, and 5.4 m/sec. The moment necessary to close either switch can be found by:

$$h = \frac{ML_1^2}{2EI_1} + \frac{ML_2^2}{2EI_2} + \frac{ML_1L_2}{EI_1}$$

where:

h = 2500 Å
L_1 = 55 μm
L_2 = 50 μm
EI_1 = 2.33e-10 Nm^2
EI_2 = 6.32e-10 Nm^2
M = Moment required to close the switch; found to be: 1.23e-8 Nm

The flow velocity must then be calculated that causes this moment. It can be found through a derivation of the Bernoulli equation, which is applicable for steady, inviscid, incompressible flow, as shown below:

$$F = \frac{1}{2}\rho A V^2$$

where:

F = force necessary to close the switch (N)
ρ = density of the seawater = 1030 kg/m^3
A = total area of flow paddle (m^2)
V = velocity of current detected (m/sec)

The maximum bending stress in the structure is calculated using strength of materials theory. The equation for the bending stress is:

$$\sigma = \frac{Mc}{I} \quad \text{where: } I = \frac{wt^3}{12}$$

where:

σ = maximum bending stress (Pa)
M = moment caused by design flow (1.23e-8 Nm)
c = half the thickness of the beam (4 μm)
I = cross=section area moment of inertia (2.98e-21 m^4)
w = width of the beam at point of maximum stress (70 μm)
t = thickness of beam (8 μm)

The maximum bending stress was calculated to be 16.5 MPa and since the yield stress of gold is approximately 130 MPa, the stress caused by the design flow speed should not cause yielding. However, another concern is the stiction between the gold tip and the gold contact. The retraction force must be great enough to lift the beam from the drain once the flow has slowed/stopped. This force is:

$$h = F\left[\frac{L_1^3}{3EI_1} + \frac{L_1^2 L_2}{EI_1} + \frac{L_2^2 L_1}{EI_1} + \frac{L_2^3}{3EI_2}\right]$$

where:

F = *the retraction force of about 162* μN

From experimentation, it was found that a retraction force of 100 μN was more than sufficient to lift the beam of a MEMS device upon unloading.

After waterproofing the switches with a silicon shell, tests were conducted on the developed prototypes. Reynolds number for pipe flow is determined by:

$$\mathrm{Re} = \frac{\rho V D}{\mu}$$

where:

Re = Reynolds number
ρ = *fluid density (water = 1000* kg/m^3*)*
V = *fluid velocity (m/sec)*
D = *pipe diameter (9.525e-3 m)*
μ = *fluid viscosity (1.12e-3* $N\text{-}sec/m^2$*)*

Using a maximum Reynolds number of 45,900, the minimum distance for fully developed flow for the fastest flow detectable by the switches can be determined by:

$$l_e = 4.4D(\mathrm{Re})^{1/6}$$

where:

l_e = *length for fully developed flow in metres*

Because the water out of the faucet only reaches about 1.2 gpm, this method of testing only works for the switches that measure smaller flows. The two less-sensitive switches require a pump to increase the flow to the point that these switches can close. Therefore only the two low-flow-rate switches were tested and the results agree well with the calculations. The average switch closed at about 1.1 gpm and lifted off at about 0.9 gpm (McGruer et al. 2002).

3.12.5. Chemical Sensing

Chemical sensing is the most fundamental of biological sensing modalities, common to almost all life forms, and is exhibited by single celled microbes for finding food (attractant) and avoiding poisons (repellent). Odour in particular is such a sensing modality. This is also a ubiquitous scientific sensor onboard planetary spacecraft which attempt to determine the elementary composition of planetary atmospheres and soil, eg. the gas chromatograph-mass spectrometer. Chemosensory cells of high sensitivity are ubiquitous throughout the animal kingdom. This is the basis for the electronic nose. Taste and smell are chemical senses – taste is in fact augmented by smell indicating the partial fusion of these senses in the brain. In fact, pheromone release is used as a broadcast signal in biology and animals with directional capability can track odour plumes or chemical trails. Smell detects volatile components (mediated by nasal fluid) while taste dissolves solid and liquid chemicals. Taste receptors are located in groups and are generally associated with the mouth in most species, though flies also possess taste receptors on the tarsi at the ends of their legs. Molecules for detection bind to the receptor proteins which depolarise the olfactory neuron membrane generating an action potential. This is similar to the behaviour of some SAW (surface acoustic wave) sensors and coated micromachined cantilever-based micro-sensors. Piezoelectric materials provide the basis for many MEMS-based sensors include the highly versatile SAW sensors. Tin oxide thick films have resistance which increase in the presence of oxygen but can be doped with metal catalysts for sensitivity to a range of reducing and combustible gases in an oxygen environment. The quartz microbalance uses a piezoelectric quartz crystal as a sensitive balance. Coated with specific-affinity chemicals cases a drop in the quartz resonant frequency on exposure to the specific target gas. Conductive polymers such as polypyrrole exhibit resistance changes on exposure to volatile chemicals. Many animals possess nasal cavities which are highly convoluted offering large surface areas for a wide range of smell abilities. Insects have their smell receptors located on the antennae which they wave to provide directional information. This is particularly important in eusocial insects such as termites and ants which communicate through pheromones. Given that most odours comprise an array of chemicals, the pattern of activation allows discrimination of a wide variety of smells. This makes such sensors suited to neural network processing and classification. In some mammals, directionality data can be obtained from movement of the nose, but lobsters which deploy their chemical-sensitive antennae can determine directionality data with great accuracy. Chemical analysis by electronic noses is based on the production of ionic or electronic signals due to the action of gas molecules on the active sensor receptor. The human olfactory system comprises of three main components [Pearce 1997]: the olfactory delivery system, the main olfactory bulb, and the olfactory cortex. Odour, which mostly comprises lipid soluble molecules, is delivered to the olfactory epithelium through the nasal passages. Odorants diffuse through a mucus layer to bind with chemically sensitive membranes of the olfactory receptors of the epithelium which cover ~5cm^2 of the nasal cavity. The mammalian nose possesses ~2.5x10^7 neurons capable of differentiating 10,000 different odours, the intensities of which are codified logarithmically. Smell, although of primitive heritage, has been implicated in episodic memory formation suggesting a cognitive component to this sense. One of the most important, though not the only, determinants for an odour molecule is stereochemical shape. Small soluble acidic proteins – odorant-bonding proteins (OBP) - in olfactory mucus which belong to the licocalin carrier protein family (which transport hydrophobic molecules in aqueous media) have been implicated in olfaction. Organic materials operate by key-lock induced fit interactions on specific recognition centres [Gopel 1996]. Electronic noses have yet to replicate the remarkable capacities of biological odour recognition but have been employed for specific compound detection [Gopal et al 1998]. Attempts to replicate the wide chemical capacities of the human nose have been made such as through the use of dyes and polymer mixtures which alter their fluorescence on exposure to certain compounds [White et al. 2002]. We feel that the most promising technology for such array approaches suited to planetary exploration (such as the detection of methane, amino acids, hopanoids, etc) may be arrays of specifically coated micromachined cantilevers or SAW sensors.

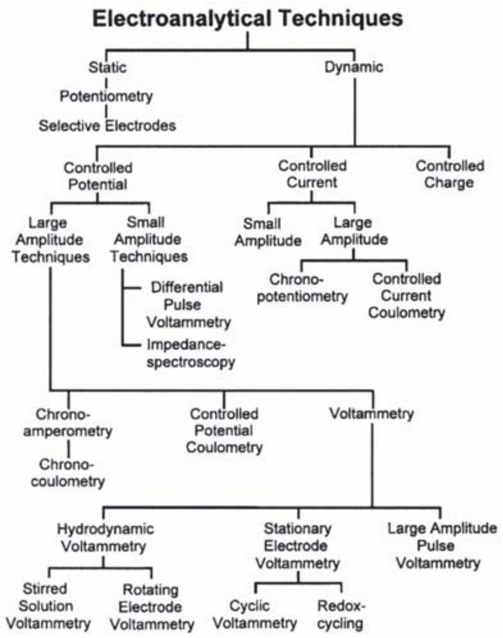

Figure 202 – Different types of electro-analytical techniques for chemical/biochemical sensing (Ziegler et al. 1998)

Microsystems offer the possibility of sophisticated instrumentation such as chemical systems on a single chip. The most traditional form a gas detection is through field effect transistors (FET) such as the hydrogen-sensitive palladium-gate FET [Lundstrom 1996]. There are two types of acoustic wave sensor which are sensitive to mass – the quartz crystal; microbalance (QCM) and the surface acoustic wave (SAW) detector. Both devices are based on the generation of acoustic waves in piezoelectric materials by the application of metal electrodes either side of the crystal. A change in mass causes a shift in the frequency/velocity of the acoustic waves. The application of thin film coatings turn such devices into chemical sensors. The QCM is a resonator which generates shear acoustic waves of opposite polarity at the electrodes on each face of the crystal. Resonance occurs due to constructive interference between waves when the crystal thickness h equals multiples of the acoustic half-wavelength n(λ/2). The acoustic phase velocity is given by:

$$f_n = v / 2h$$

where:

$$v = \left(\frac{\mu}{\rho}\right)^{0.5}$$

h = crystal thickness ~0.033 cm typically
f_n = frequency of nth mode ~5 MHz typically
v = phase velocity of shear wave
μ = shear stiffness ~2.95 $x10^{11}$ dyne/cm^2 typically
ρ = density ~2.65 g/cm^3 typically

The mass at the surface of the QCM alters the resonant frequency of the quartz crystal:

$$\frac{\Delta f}{f} = -\left(\frac{\rho_s}{h\rho}\right)$$

where:

ρ_s = *surface film density*

Bulk acoustic wave (BAW) devices have been proposed but surface acoustic wave (SAW) devices offer much greater flexibility. SAW devices are constructed from a piezoelectric substrate – typically $LiTaO_3$ or $LiNbO_3$ - with metallised electronics including reflector gratings implanted on the surface. The reflector gratings act as mirrors for the SAW and the electrode strip width and spacing are $<\lambda/4$ and determine the limit of sensitivity currently ~0.3μm. Operating the SAW at high frequencies increases its sensitivity. An electrical input to an interdigital transducer (IDT) is transformed by the piezoelectric effect into a micro-acoustic Rayleigh wave propagating on the surface at ~3-5 km/s with an oscillation frequency ~30 MHz – 3 GHz. The design parameters are the phase velocity of the SAW and the coupling coefficient K^2 which are determined by the physical properties of the substrate including elasticity, piezoelectricity, dielectric permitivity and mass density [Springer et al. 1999]. The coupling coefficient quantifies the electrical to mechanical energy conversion efficiency. SAW sensors pass surface acoustic waves along a thin wafer of material, so any change in wafer mass causes a change in wave velocity – they are highly sensitive to surface mass loading which lowers the SAW velocity:

$$\Delta f = f_{load} - f_0 = K^2 f_o^2 \Delta_m$$

where:

$K = 1.26x10\ m^2/kgHz$ *for quartz*
Δ_m = *mass/unit area added due to gas adsorption*

The SAW generates an electric output signal on the IDT. Covering the wafer with imprinted polymer will ensure that the mass will change when target molecules are present. The detection of such mass changes also occurs in biological systems – this provides the basis for the recognition of antibodies by epitomes of antigens. Two types of SAW devices are used as chemical sensors – the delay line and the resonator. The delay line has two sets of electrodes – one set to create an acoustic wave determined by the electrode spacing and a second set to detect the wave after the propagation delay. The resonator is an acoustic cavity created by a series of ridges in the SAW surface either side of the IDTs. One electrode creates the SAW while the second feeds back to the drive circuit.

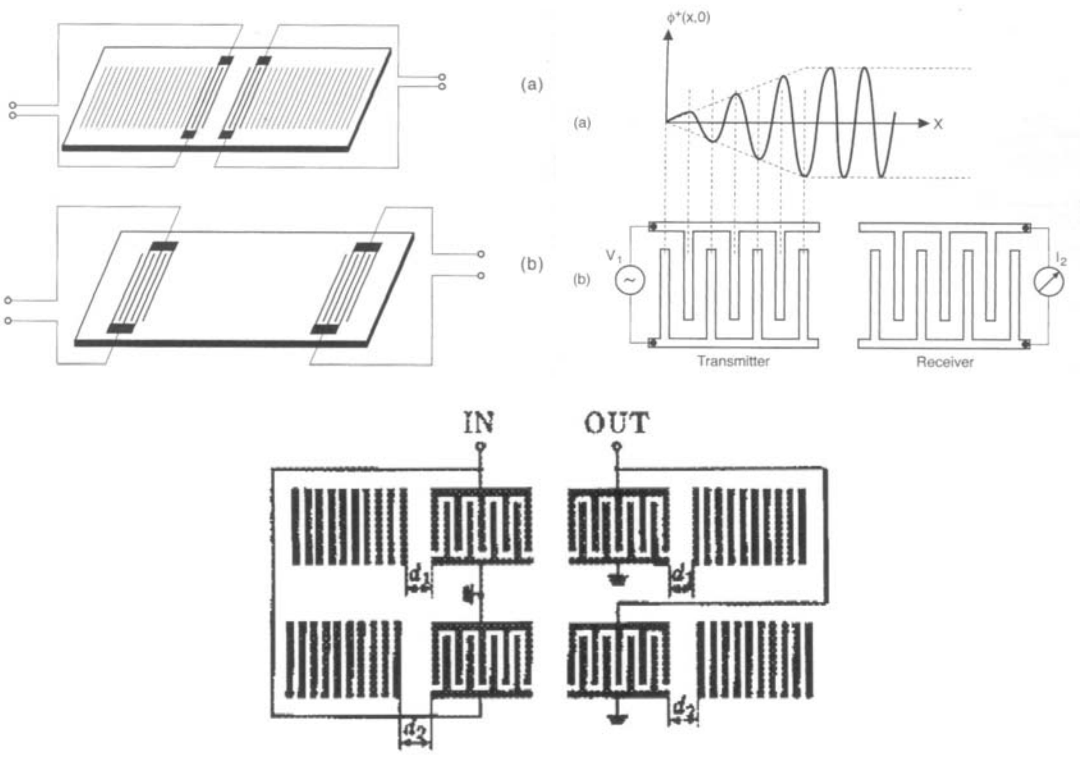

Figure 203 – a) SAW resonator with delay line; b) transmitter and receiver and generated surface acoustic wave; c) interdigital transducer with reflector gratings (Weiller 1999)

The use of surface acoustic wave oscillator fabricated on the surface of a thin diaphragm in contact with the medium offer the possibility of simultaneous measurement of pressure and temperature [Reeder & Cullen 1976]. SAW sensors are particularly suited as gas sensors – a SAW delay line coated with a gas sensitive layer will have an altered time delay if gas is adsorbed by the layer. The time delay is translated to a frequency by using the delay line as a resonator for an oscillator. Polymer coatings are typically used so that the vapour of interest diffuses into the coating, generating a mass loading effect – this alters the attenuation and phase of the SAW. The frequency shift is given by:

$$\Delta f = -kS\Delta m$$

where:

Δm = *change in mass/unit area*
k=*geometric constant*

Generally, responsivity increases with temperature. Coatings include activated charcoal, silica, alumin gels, zeolites and porous polymers. Chemisorptive reactions require significant energy inputs. Metals such as Pd and Ag are used for the detection of H_2 and O/O_3 respectively. ZrO_2 is used as a chemical sensor to monitor O_2 due to the high temperature O^{2-} ion conduction. $ZnCl_2$ is hygroscopic so will absorb water forming $ZnCl_2.(H_2O)_x$ which generates mass amplification to an accuracy of 5 ppm. Polymer coatings such as polybutadiene offer sensitivity to O_3 by the irreversible formation of ozonides – it has a sensitivity of 10 ppb. Polymer films are used primarily to detect organic gases but tend to be of limited specificity, so several sensors with different coatings are used to detect specific substances. However, polyethylene maleate (PEM) offers sensitivity to a range of organic gases such as acetone and methanol. Copper phthalocyanine (CuPc) coatings may be used as an oxidising species detector – the conductivity of the coating increases on exposure to oxidant causing a change of attenuation of the SAW. Horizontally polarised SAW may be used for detecting immunoglobins in liquids [Welsch et al 1997]. The antibody, which does not need to be labelled, is immobilised on the $LiTaO_3$ substrate by protein binding while the antigen to be detected is dissolved in ~5μl of buffer solution. Immunoglobin detection is highly specific.

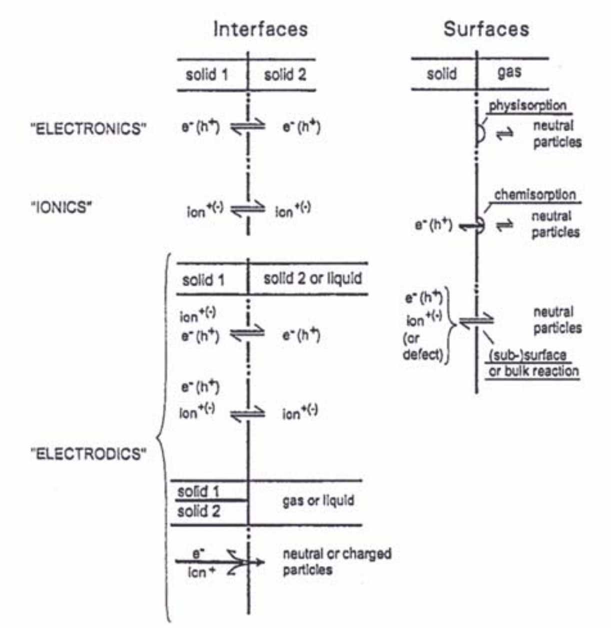

Figure 204 – Types of interfaces of electronic and ionic devices

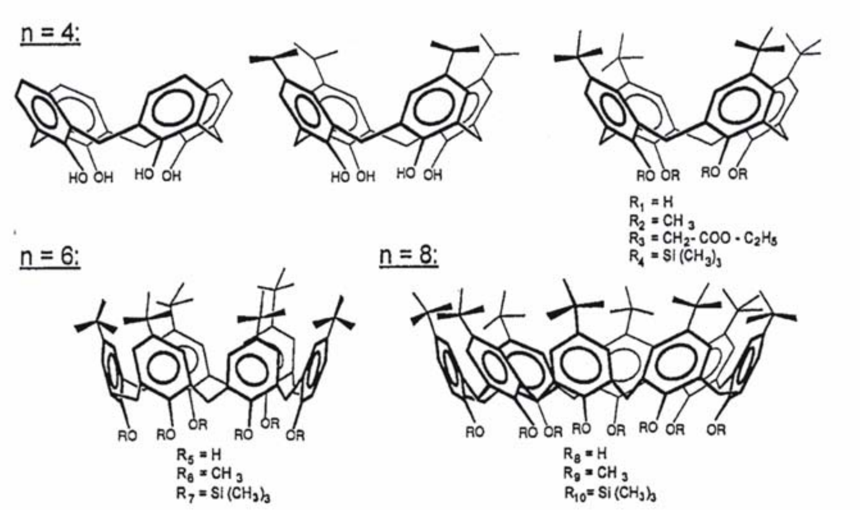

Figure 205 – Calixarenes as coatings to acoustic wave devices provide the basis for volatile organic compounds. Such coatings act as molecular cages

The most basic biological chemical detector is taste. A taste sensor comprising 8 kinds of lipid/polymer 200μ membranes (eg. oleic acid) has been developed to differentiate several taste qualities of different groups amino acid groups [Toko 1998]. These are grouped into the five basic taste qualities: saltiness from salts such as NaCl, sourness due to H+ ions from acids such as HCl, sweetness due to sugars such as sucrose, and bitterness due to quinine etc, and umami (deliciousness) due to a number of sodium compounds such as monosodium glutamate. Taste by itself is not a particularly useful mode of classification for organic materials but smell offers a broader capability. An electronic nose may be defined as an array of electronic-chemical sensors with partial specificity in conjunction with a pattern recognition system capable of recognising odours. Of particular importance is the necessity for the delivery of chemicals to the sensors. Some kind of flow injection system is required unless the chemical environment is gaseous – this will involve filters, non-return valves, two-way purging valves, mass-flow controllers, compressor pumps, flowmeters, etc. The electronic nose can be used as a replacement of the mass spectrometer in conjunction with a gas chromatograph.

Chemical analysis by electronic noses is based on the production of ionic or electronic signals due to the action of gas molecules on the active sensor receptor. The human olfactory system comprises of three main components [Pearce 1997]: the olfactory delivery system, the main olfactory bulb, and the olfactory cortex. Odour, which mostly comprises lipid soluble molecules, is delivered to the olfactory epithelium through the nasal passages. Odorants diffuse through a mucus layer to bind with chemically sensitive membranes of the olfactory receptors of the epthelium which cover ~5cm² of the nasal cavity. Olfactory neurons are unique among neurons in their ability to regenerate. The mammalian nose possesses ~2.5x10^7 neurons capable of differentiating 10,000 different odours, the intensities of which are codified logarithmically – logarithmic compression of the output to stimuli has the form with output O: $O = I(t)e^{-O(t)t}$. One of the most important, though not the only, determinants for an odour molecule is stereochemical shape. Small soluble acidic proteins – odorant-bonding proteins (OBP) - in olfactory mucus which belong to the licocalin carrier protein family (which transport hydrophobic molecules in aqueous media) have been implicated in olfaction. OBPs are weakly specific carriers of odour ligands within the mucus to the olfactory receptor proteins. The odour molecule binds to a receptor which causes a specific G-protein to bind to guanosine triphosphate (GTP) to release GDP. Subsequently, olfactory adenylate cyclase (AC) is activated by the G-protein. ATP is converted into cyclic AMP which reaches its maximum production after 50 ms and opens the olfactory cyclic nucleotide-gated ion channel for Na+ ions for neural transmission. Organic materials operate by key-lock induced fit interactions on specific recognition centres [Gopel 1996].

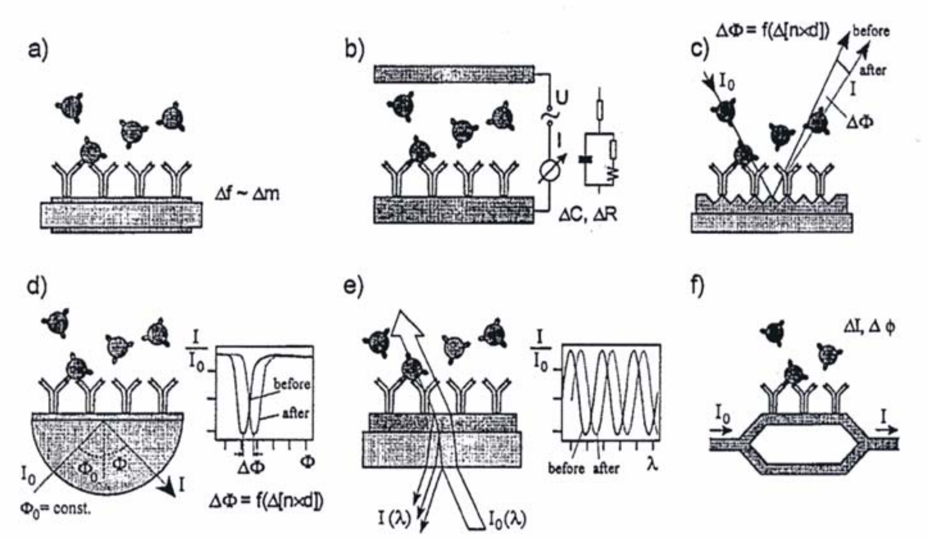

Figure 206 – Different approaches to bioaffinity/immuno-sensors (Gopel 1996)

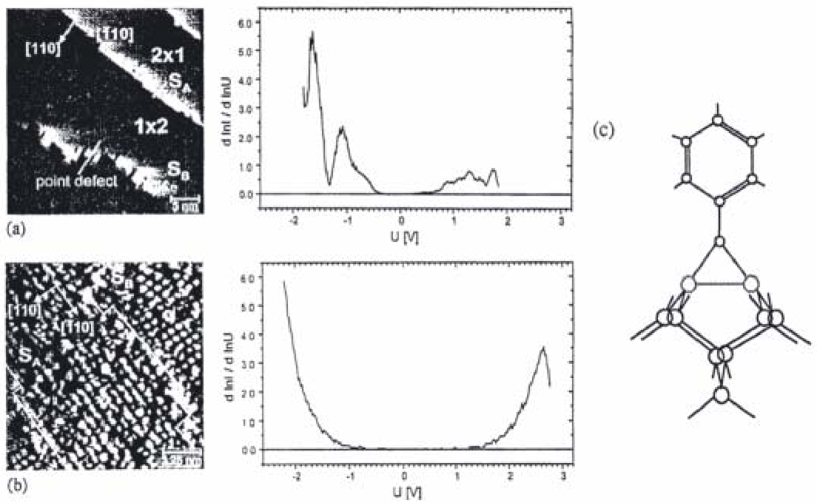

Figure 207 – Organic/inorganic interface due to chemisorption of organic analine overlayer on Si (100) substrate (Gopel 1996)

The olfactory cells are depolarised differentially according to their individual sensitivities and position on the epithelium. A characteristic pattern of electrical activity is transmitted in a coarsely topographic manner to the main olfactory bulb. Receptor cells with similar responses are localised together into bundles of ~25,000 neurons (glomeruli). The main olfactory bulb is highly stratified into distinct layers of 10^4 neurons. The neuronal activity at each layer may be considered to represent a molecular "image" in which each layer successively performs complex computations. However, olfaction does not represent spatial information in that the position of the stimulation on the receptive surface does not represent the locality of the stimulus. The locality on the receptor surface does map the specific properties of the stimulant. The lateral olfactory tract from the olfactory bulb projects and connects non-topographically to the pyramidal cells of the olfactory cortex. Olfactory associative memory provide the basis for the recognition of complex odours. Such olfactory memory resides in the piriform association cortex which bears similarities of structure to the hippocampus. The human nose is a highly sensitive instrument of olfaction which may be emulated by arrays of electrochemical sensors connected to a microprocessor and which may be trained to recognise gaseous chemicals, especially organic chemicals. The response times for biological olfactory neurons are generally much faster than artificial chemical sensors ~20-60s. Olfactory receptor output is based on the neuron modelled as a sigmoid function: $S = \sum \frac{1}{\exp(b_i - a_i x)} + a_0$ while artificial receptors typically follow either linear $\left(S = \sum a_i x_i + a_0\right)$ or power law $\left(S = \sum a_i x_i^n + a_0\right)$. Nasal cavity design may be crucial in determining airflow patterns over the epithelial layer of the nose, eg. dogs have more intricate folding of the nasal cavity than humans. The nasal epithelium cavity is studded with ~10M sensory cells (each with ~10^3 different chemical receptors) for decoding in the olfactory bulb.

Chemical microsensors react selectively with chemical species to be detected by means of a number of physical properties. Hundreds of such microsensors may be implemented on a single Si wafer together with its signal processing electronics. Furthermore, multichip modules allows integration with microprocessors and wireless transceivers. There are six main methods for the detection of (bio)chemicals without labelling [Gopel 1996]:

- mass-sensitive detection, eg. piezoelectric microbalance;
- electrochemical impedance detection, eg. ISFETs
- optical methods:
    - grating couplers for changes in optical thickness, i.e.. refractive index
    - surface plasmon resonance
    - reflectometric interference
    - Mach-Zehnder interferometer.

Chemiresistor is the simplest of chemical sensors in which the resistance of a thin film material may be measured – these include metals (Pd sensitive to H_2 and Ag sensitive to O atoms), organic conductors (phthalocyanines sensitive to NO_2 and O_3), and metal oxides (SnO_2 sensitive to most combustible inorganic and organic gases except N and S based organic compounds). Metal chemiresistors such as Pd undergo increased resistance on exposure to H_2 gas through the formation of hydride – this is irreversible and is strongly effected by the presence of ambient O_2. WO_3 is sensitive to H_2S but requires high temperatures. Metal oxides have limited utility in terms of their specificity, high temperature operation up to ~500^{o}C, and requirement for ambient O_2. Conducting polymers may be used as chemical sensors – intrinsic conducting polymers are conductive without doping and extrinsic conductors are made conductive by impregnating with carbon black. Polyacetylene, polypyrrole and polyanaline are intrinsic conductors due to their pi bond conjugations. They are sensitive to a number of gases, eg. polypyrrole is sensitive to NH_3. Extrinsic conducting polymers are more versatile due to the wide range of polymers that can be used and they are almost entirely reversible. Gas absorption causes swelling of the polymer which lowers the conductivity. They are particularly suited to deployment as polymer coatings for SAW sensors.

Phthalocyanines are macrocyclic organic molecules with a central metal ion surrounded by molecular ring structures. They are effectively organic semiconductors which are stable to high temperatures ~500^{o}C and exhibit a conductivity that increases with temperature: $\sigma = \sigma_0 e^{-E/2kT}$. Their conductivity is low at room temperature so are typically operated at 120-200^{o}C and their conductivity can be increased by using IDTs. This conductivity is sensitive to exposed gases making them useful as chemiresistors. Perovskite as an electrode material offers the possibility of detection of CO, NO and hydrocarbons – they also offer high potential for electro-catalysis in fuel cells.

The most traditional transducer in electronic noses is the gas-sensitive field effect transistor in which sensor arrays are used for pattern recognition. The selectivity and sensitivity are controlled by the composition and thickness of the gate metal, and the temperature of operation. The detected molecule causes a shift in the voltage characteristics of the transistor due to electric polarisation in the metal-semiconductor interface. The hydrogen-sensitive Pd-gate FET may be used to detect low stability hydrogen-containing molecules which release their hydrogen on contact with the catalytically active Pd gate, eg. some alcohols, some alkenes but not amines [Lundstrom 1996]. This device can be made sensitive to NH_3 or H_2S by modifying the gate metal. Ion-selective FETs (ISFET) are commonly used for biosensors in which an insulating semiconductor oxide layer on Si substrate results in an inversion layer in the underlying substrate. Equivalent RC circuits formally describe the frequency behaviour of such devices. Inorganic oxides show electron, ion or mixed conduction changes on chemisorption of gases or upon catalytic reactions – Pt is often used to improve the selectivity of gas detection by improving the catalytic activity of the surface. The biolipid transporter molecule lactose permease transports protons and lactose molecules in a 1:1 relation across cell membranes – it can be used in ISFETs by monitoring the protons (i.e.. pH changes) at the electrode which directly correlates with lactose concentrations. Modulation of the electric field in the insulator layer generates a modulation of the electric field in the silicon which alters the resistance of the source-drain channel. The drain current can thus be modulated as a function of a liquid/oxide interface potential as a function of the electrolyte pH. However, SiO_2 does not interface well with fluids so an additional interface oxide layer is adopted, eg. Al_2O_3 [Bergveld 1996]. Maximum sensitivity to pH is –59.6 mV/pH at 298 K (Nernstian response). Such devices can also measure any substance with a surface charge but are totally unsuitable for immunoglobin detection due to immunoglobin Debye-Huckel charge shielding by counter ions of the solution.

Electronic noses possess arrays of sensors, each sensor producing an electronic output as a function of sensor-specific chemical interactions. The combined electronic signals from the array give a characteristic pattern which can be correlated with a given odour or set of volatile compounds. The strength of the signal reflects the concentrations of each set of compounds. One difficulty is the degree of selectivity which suggests that a range of relatively unselective sensors should be employed with different responses to a range of gases – such an approach would require the use of neural networks trained to deconvolute the data. Gas-sensitive devices may be used for distributed sensor arrays with sensors of different selectivities and sensitivities to obtain olfactory images which represent feature vectors of the gas composition. It is desirable for the gas-receptor to be reversible effectively eliminating covalent bonding active materials such NiO. Strong bindings of gas species with coatings can be highly selective but their irreversibility makes them act as dosimeters to total dose which can suffer saturation. Thus reversible sensors are preferred.

Molecular absorption at a surface may be quantified by the rate of surface absorption [Weiller 1999]:

$$R_a = k_a p(1-\theta) = p(1-\theta) A_0 e^{-E_a/RT}$$

where:

k_a = *rate constant for absorption*
R = *gas constant*
p = *partial pressure of absorbent*
θ = *fraction of occupied surface*
A_a = *frequency factor for absorption*

Rate of desorption from the surface is given by:

$$R_d = k_d\theta = \theta A_d e^{-E_d / RT}$$

At equilibrium, $R_a=R_d$, so: $K_a = \frac{k_a}{k_d} = \frac{\theta}{p(1-\theta)} = \frac{A_a}{A_d} e^{(-E_a - E_d)/RT}$

Hence, $\theta = \left(\frac{pK_a}{1+pK_a}\right)$

As the partial pressure increases, so the surface coverage approaches complete coverage. The equilibrium constant relates the both enthalpy and entropy thus:

$$\Delta G = -RT\ln(K_a) = \Delta H - T\Delta S$$

$$\rightarrow K_a = e^{(-\Delta H / RT)} e^{(\Delta S / R)}$$

where:

$$\Delta H = E_a - E_d$$
$$\frac{A_a}{A_d} = e^{(\Delta S / R)}$$

Gibbs free energy must be negative for a reaction to proceed spontaneously, and exothermy implies that enthalpy and entropy are both also negative. Thus, the rate of surface reaction is given by:

$$R = \frac{k_s K_a p}{(1+K_a p)}$$

where:

k_s = *surface reaction rate constant*

For the sorption of gases in polymers, the concentrations absorbed gas is related to its boiling point:

$$C = C_g e^{(-T_b \Delta S / RT)}$$

where:

C_g = *concentration of gas phase*
T_b = *boiling point*
ΔS = *entropy of condensation for simple non-polar liquids = 20.3 cal/molK*

Hence, boiling point of the species determines the response of a sensor/coating combination.

An electronic nose based on a micromechanical array of Si cantilevers has been developed in which the cantilevers are sensitised with a number of different coatings such as metals, monolayers and polymers [Lang et al 1999]. The chemical reaction is thus transduced to mechanical deflection of the specific cantilever. Each cantilever was 500μ in length, 0.8μ in thickness and 100μ in width with a spring constant of 0.02 N/m. Uncoated cantilevers provide references. Mass change from the accumulation of a small sample at the apex of the cantilever coated with a selective sensor will alter the resonant frequency of the cantilever, typically ~50 kHz. The decrease in oscillation frequency due to mass uptake is given by:

$$\Delta m = \frac{k}{0.27\pi^2}\left(\frac{1}{f_f^2} - \frac{1}{f_i^2}\right)$$

where:

$f_{i,r}$ = *initial and final resonance frequencies*
Δm = *mass change*
k = *cantilever spring constant*

An array of such devices with individually coated cantilevers provides the basis for a wide range of sensitivities. The cantilevers are sequentially read-out optically by recording the position of an incident light beam reflected from the cantilever by a photodetector array. The oscillatory response of the cantilver deflection is processed by a lock-in amplifier. A cantilever coating of zeolite ZSM-5 which absorbs water selectively provides the basis for a water sensor.

The pattern of response across the sensor array provides the basis of odour recognition. Partial specificity of the array elements must overlap to maximise the information from the environment with a high degree of redundancy – indeed, this is the basis for biological hyperacuity through sensitivity enhancement [Pearce 1997]. An electronic nose with enhanced noise rejection and sensitivity comprises an array of sensors D_i of a single class replicated n times. Each sensor with a signal-to-noise ratio S_D will generate a combined averaged signal with a signal-to-noise ratio of $S_D\sqrt{n}$. A set of known chemical prototypes may be learned corresponding to each chemical class – the self-organising map Kohonen network or Boltzmann machine ANNs would be suited to such a task. Such a device can act as a content-addressable memory which maps input and output vectors with reproducible pattern recall, generalisation and robustness. Ziegler et al (1998) suggests however that the key to electronic noses relies on the specific recognition of molecules within an array of different recognition elements – this is essential for maximising the quality of the signal with minimisation of interference from other molecules. Indeed, different transduction mechanisms should be employed for maximum orthogonality for independent chemical determination.

Chemical sensors made from different conductivity polymers embedded with carbon black particles provide an array of large bandwidth. When the polymers absorb molecules from odorous chemicals they swell due to hydrophobicity, eg. benzene is hydrophobic. Plastics comprise of tangled fibres which swell when organic molecules are absorbed between the strands. Carbon black may be scattered among the strands and if the polymer is non-conducting, (i.e.. any plastic can be used), carbon black provides the only pathway for electric current to flow. If the polymer swells, the electrical pathway is disrupted reducing the conductivity of the material. Different plastics swell to different degrees depending on their chemical affinities. Even doping concentration will affect the behaviour of the material, eg. 10-20% carbon black will exhibit different behaviours. The organic matter impregnating the plastic will alter their conductivity and so alter the magnitude of the electric current passing through them. Different polymers respond differently, so each chemical species may yield a unique electrical fingerprint. A greater variety of sensors yields a greater diversity of discrimination among a wide variety of chemicals. A 5-10 sensor head is reckoned to provide good results. Commercially available plastics are suitable, eg. polystyrene (hydrophobic), polyvinyl alcohol (hydrophilic). Artificial olfaction may use artificial neural networks to learn patterns of responses associated with different chemicals. This requires training to recognise any given species. Repeated exposure provides good discrimination ability.

Optical techniques also offer prospects for chemical sensing. A sheath of optical fibres with ends coated with plastics mixed with Nile Red dye can act as sensors to odourous organic chemicals. Light shine down the fibres to the tips make the dye flouoresce characteristically at different wavelengths according to the swelling of the polymer. Part of this light may be reflected up the fibre to be measured by CCD camera into a digital signal for analysis by an artificial neural network. The fluorescence offers additional variables for diversity – intensity wavelength and fluorescence duration. Such devices can distinguish between structures differing by only a single carbon atom. Commercial electronic noses of this type take only 1-2 minutes to fingerprint smell but they require a stable reading at a controlled temperature. The use of optical transduction offers a variety of opportunities:

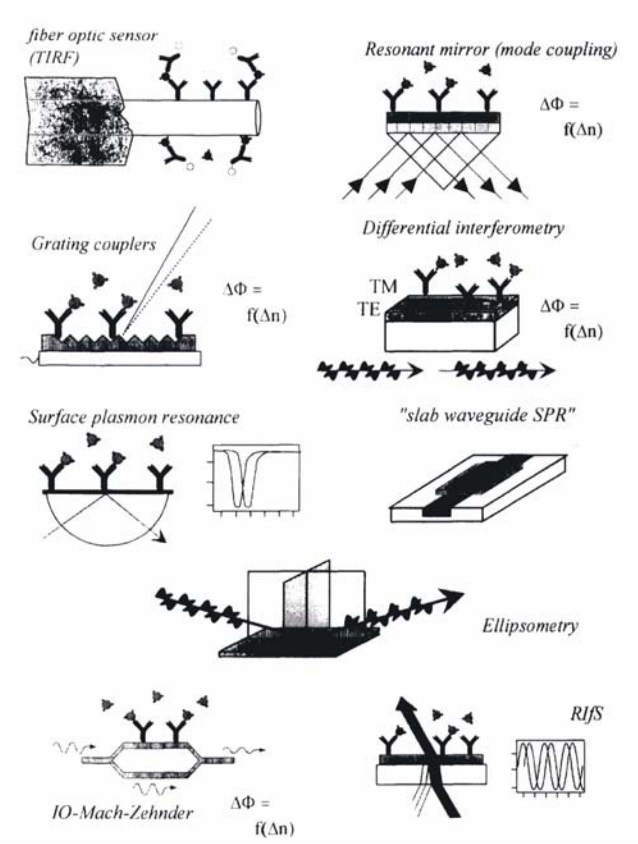

Figure 208 – Types of optical transduction mechanisms (Brecht & Gauglitz 1995)

Light is shone down fibres to the tips to make the dye flouresce characteristically at a different wavelength according to the swelling of the polymer. Part of this light may be reflected up the fibre to be measured by CCD camera into a digital electrical signal for analysis by the ANN. The flourescence offers additional variables for diversity - intensity, wavelength and flourescence duration. These devices can distinguish between structures differing by only a single carbon atom. Commercial noses take around 1-2 minutes to fingerprint smell as they need a steady reading at controlled temperatures. A variety of methods of analysis may potentially cope with this problem, eg. Kalman filtering. Potentially, such architectures could be reduced in size to chip level with sensors and ANN integrated onto a single silicon chip. For example, several 1mm wide wells may be etched onto a 1cm wide silicon wafer connected to a PC. Each well has deposited in it a different type of sensor material. The potential for small, low-mass, low power and relatively inexpensive artificial noses is large, including interplanetary probes for sampling planetary atmospheres for trace amounts of complex organic compounds.

3.12.6. Multi-Sensor Fusion

A distributed and diverse sensor suite enables a system to cope with failures gracefully due to functional redundancy. They can also provide multiple information sources which operate in parallel. Each sensor output is degraded by noise so if multiple sensors are used to determine the same property, uncertainty is reduced. Furthermore one modality may suffer degradation under certain conditions but by using additional modalities such problems can be eliminated. Multisensor fusion is a natural extension of using multiple sensors for verification. For complex environments, it is desirable to simultaneously integrate such sensory data from diverse multiple sources. Multisensor data fusion improves the overall performance of interpretation and classification by providing a more reliable and robust estimation of the environment than each sensor considered alone. Sensor suites offer redundant information with different fidelities or offer complementary information of different modalities. It decreases feature ambiguity, uncertainty due to noise and reduces search spaces during matching at later processing stages. In robotics, the commonest fusion proposals have been the fusion of stereoscopic vision sensors, the fusion of vision information with range finder information, and the fusion of visual and tactile data. Although multisensor fusion is believed to occur in the human brain, ventriloquism seems to indicate that visual information dominates auditory information in the fusion process [Luo & Kay 1989]. Furthermore, space adaptation syndrome (space sickness) also suggests that sensory fusion processes are not complete with vision dominating over vestibular data, the mismatch being the cause of the nausea characteristic of the condition. Such hierarchical dominance between sensors is a natural one in robotics which avoids the many problems of sensor fusion. Vision may have high priority over gross motions with proximity sensors taking over dominance during near-contact guarded motions until tactile sensors become dominant during touch manipulation. Several problems are immediately apparent in sensor fusion [Pau 1987]:

- different sensors provide different physical transduction mechanisms, different locations, with different sensitivities to the environments to which they are exposed, have different bandwidths, different spectral detectivity profiles and differing data processing rates
- registration is required of data points of one sensor corresponding to data points on another sensor;
- merging data from widely differing forms of representation into a unified representation medium

Sensor fusion can occur at different levels from the integration of raw data to the integration of object properties of a world model. Fusion essentially involves the integration of different sensory modalities into a single representational format such as a world model. A common central environment for the integration of sensor data may be provided by a distributed blackboard architecture to provide high level sensor fusion with each sensor module passing messages to the blackboard [Harmon et al 1986]. This architecture provides a means for dealing with uncertainty in visual scenes regarding important objects within the scene. The blackboard provides centralised control of multiple specialist knowledge sources. It provides the mechanism for multisensor fusion at different hierarchical levels [Clemint et al 1993]. Each specialist sensor module is independent but they all share information by communication through the global blackboard database. Objects may be described in terms of features and relations between these features. The blackboard may be organised as a semantic network data structure to hierarchically structure objects and their properties. The sensors update the blackboard with new information to cause blackboard state changes. Shafer et al (1986) called their blackboard model of sensor fusion a whiteboard since it represented a topological map of the scene using potential fields and the modules were separate expert systems dedicated to mobile robotic behaviour for their NAVLAB mobile robot: CAPTAIN, NAVIGATOR, PILOT and PERCEPTION expert subsystems. Sensor fusion was implemented by integrating through the whiteboard map with other processes.

For similar modalities, redundant information may be fused at a lower level of representation. If redundant data is in conflict, a mechanism for assigning confidence to each data value is needed. Confidence may be assigned to provide weighted averages with weightings to each sensor value. Multi-sensor fusion can adopt a number of different methods – weighted average, Kalman filter, Bayesian estimation, fuzzy logic or neural networks. Kalman filters offer optimal estimation of measurements in the presence of Gaussian noise. Although restricted to linear systems, the extended Kalman filter may be used for non-linear systems in which linearisation errors are small. The simplest fusion strategy involves raw sensor measurements of the same property obtained from multiple sensors being combined (direct fusion) [Hackett & Shah 1990]. Before fusion can occur consistency between diverse data sources must be determined to check that they represent the same physical entity. The squared Mahalanabis distance which should be small for the same representative object can determine this:

$$r^2 = (x-\mu)^T \Sigma^{-1} (x-\mu)$$

where:

x = random vector of n random variables with mean μ and variance σ^2.
Σ = covariance matrix=diag(σ_i)

Mahalanobis distance follows a χ^2 distribution. The simplest way to fuse data is to use Bayes maximum likelihood ratio to model the sensor uncertainty. Bayes theorem assigns probability on the basis of evidence. Bayes theorem represents uncertainty as a probability between 0 and 1. Classic logic is limited in that it can apply to certain knowledge only. Bayes method is used to assign confidence factors to expert system production rules. Bayes theorem may quantify the degree of truth via probabilistic induction, eg. in PROSPECTOR, a mineral exploration expert system [Duda, Hart & Nilsson 1978]. Reasoning under uncertain conditions may be modelled probabilistically by Bayes theorem such that belief in a hypothesis increases with the evidence in favour of the hypothesis. The Bayes risk is minimised in selecting a given action. The minimax rule makes the decision that minimises the worst-case risk which can be sub-optimal. The Bayesian approach estimates the minimum expected risk using a priori probability distribution of the state of nature [Hu & Brady 1994].

Posterior probability: $p(\theta \mid z) = \frac{p(z|\theta)\,p(\theta)}{\sum p(z|\theta)\,p(\theta)}$

where:

$p(z \mid \theta)$ *=conditional probability density function of observation z*

This may be weighted by a loss function $L(\theta)\geq 0$ to enable selection of the decision with maximum expected loss according to Bayes theorem. Sometimes, conclusions are in conflict though not inconsistent due to the variable strengths of the certainty factors between the antecedent and the consequent. Conflict resolution is based on the conditional dependence of the implication on assumptions or defaults about the evidence. The primary criticism of Bayes theorem stems from the fact that it cannot represent ignorance – if evidence is only partially in favour of an hypothesis, it must also partially support the negation of the hypothesis such that $p(H|E)+p(\sim H|E)=1$. This is clearly counter-intuitive and contradictory. Any probability not assigned to evidence in favour of the truth of the proposition H is assigned to evidence in favour of ~H.

Dempster-Schafer belief theory overcomes this problem by considering uncertain evidence as representative of subjective ignorance [Pang et al 1987]. It allows a distinction between disbelief and ignorance and distinguishes between degrees of belief (based on evidence for an hypothesis), and plausibility (based on the measure of evidence that fails to refute hypothesis). Evidence does not give a probability distribution over a set of hypotheses. Parts of the belief may be distributed to any subset of $H=\{H_1 \ldots H_n\}$. The distribution of unity belief for a proposition A being true over the set H:

$$H = \Sigma m(A_i) = 1$$

where:

$m(A_i)$ = basic probability assignment $0<m<1$ representing the strength of the evidence with $m(0) = 0$
$= \Sigma p(A_i)/(1-\Sigma p(A_i))$ [Dempster rule]

The rest of the belief does not have to be assigned to the negation of the hypothesis, separating ignorance and disbelief, i.e.. there is no probability distribution over the set of hypotheses. This distribution can be over any subset of H without assigning contradictory evidence to members of ~H. We define the uncertainty interval [Bel(A), Plaus(A)]. Belief represents the lower probability representing reasons to believe proposition B while plausibility represents an upper probability of the hypothesis B representing belief if all unknown facts supported B:

$$Bel(B)=\Sigma m(A) \text{ where } A\subset B$$

$$Plaus(B)= \Sigma m(A) \text{ where } A\cap B\neq\emptyset, \text{ i.e.. } Plaus(B)=1-Bel(\sim B) \text{ with } Bel(B)\leq Plaus(B).$$

The interval between Bel(A) and Plaus(A) denotes the uncertainty about proposition A and true belief will lie in this interval. If this uncertainty were zero, Dempster-Schafer defaults to Bayes theorem. Commitment of a belief to a subset does not require the remaining belief to be committed to its negation, i.e.. $Bel(B)+Bel(\sim B)<1$. The amount of belief committed to a set cannot be subdivided into further subsets. Dempster-Schafer belief theory has high computation time complexity since for n elements, there are 2^n basic probability assignments. It is further not clear if Dempster-Schafer can detect inconsistency.

Human thinking is fuzzy – classes of objects in the real world have the property that the transition between classes is gradual rather than well-defined, i.e.. vague. Fuzzy sets have been advocated to represent that imprecision and vagueness in the form of fuzzy "if…then" rules. The propositions themselves may be of the form "Q A's are B's" where Q is a fuzzy quantifier and A and B are fuzzy sets. Possibility theory utilises fuzzy sets to represent uncertainty as distinct from probability theory [Zadeh 1978, 1983]. In conventional set theory, an element either belongs to or does not belong to a set, i.e..

$x \in A$ or $x \notin A$. Fuzzy sets introduce a membership function π to accommodate the uncertainty of degree of membership of a set where $\pi(p)$ and $\pi(\sim p)$ represent possibilities that p is true and p is false respectively in the real interval [0,1] such that $(\pi(p)+\pi(\sim p))=1$, i.e.. fuzzy logic is inherently analogue. π represents the partial membership of a set to characterise vague concepts [Novak 1986; Munakata & Jani 1994]. If π(x)=1 then $x \in A$, and if π(x)=0 then $x \notin A$ with a gradual range of membership values existing between these extremes. The membership function thus represents the grade of membership of a set. The membership function may be represented graphically as weighting factors to input-output responses. It is usually triangular in shape characterised by height, width, centre point and degree of overlap:

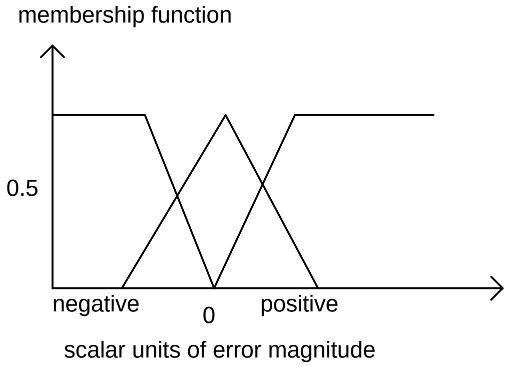

Figure 209 – Fuzzy variables and their overlap

Alternatively, one can use a bell function of the form $\pi(x) = {}^{1}\!\big/_{1+[(x-c/a)^{2}]^{b}}$ where c=membership function centre point, a=half-width, b=gradient at crossover point. The Gaussian bell function has the form $\pi(x) = a.e^{-(x-c/a)^{2}}$. These membership functions determine the firing strength of the "if-then" rule. Hence, fuzzy sets are a generalisation of classical sets. Fuzzy set operations are similar to traditional set-theoretic operations:

Union (disjunction) of fuzzy sets A and B is fuzzy set C:

$$\pi_C(x) = \max[\pi_A(x), \pi_B(x)]$$

$$= \pi(A.B) = (\pi(A) + \pi(B) - \pi(A) * \pi(B)) \to 1$$ such that x∈A or x∈B

Intersection (conjunction) of fuzzy sets A and B is fuzzy set C:

$$\pi_C(x) = \min[\pi_A(x), \pi_B(x)]$$

$$= \pi(A \,\&\, B) = \pi(A) * \pi(B) \to 0$$ such that x∈A and x∈B

Complement (negation) of fuzzy set A is fuzzy set A':

$$\pi_{A'}(x) = [1 - \pi_A(x)]$$

Fuzzy set A is a subset of fuzzy set B if $\pi_A(x) \subset \pi_B(x)$

Distributive properties apply:

$$\pi_A(x) \cap [\pi_B(x) \cup \pi_C(x)] = [\pi_A(x) \cap \pi_B(x)] \cup [\pi_A(x) \cap \pi_C(x)]$$

Associative properties apply:

$$\pi_A(x) \cup [\pi_B(x) \cup \pi_C(x)] = [\pi_A(x) \cup \pi_B(x)] \cup \pi_C(x)$$

Fuzzy sets are not considered mutually exclusive and therefore resemble linguistic measures. This may be illustrated by the concept of the fuzzy relation (composition) between two sets A and B which link them – a relation is a subset of a continuous product:

$\pi R(x,y)=A(x) \times B(y)$ - this is the Cartesian product such that

$\pi_B(y)=\max[\min(\pi R(x,y), A(x))]$ where (x,y) form an ordered pair of elements.

This statement constitutes an inference rule of the form: "if…then" and illustrates the linguistic interpretation of fuzzy sets. Information theory in its statistical aspect is capable of describing information in its syntactic modes, but unable to deal with its semantic modes, and conceivably, fuzzy sets as a theory of possibility may offer a semantic description of language. A language L may be regarded as a correspondence between a set of words x and an universe of discourse U. A word x in a language L may be viewed as a fuzzy subset M(x) of the universe of discourse such that M(x) represents the meaning of x which is not necessarily a well-defined label. The membership function π defines the fuzzy naming relation. Generally, several fuzzy rules can be combined by a union of relations: $R=R_1 + \ldots+R_n$. R represents a fuzzy algorithm from which fuzzy rule-bases can be constructed [Tong 1977]. Fuzzy sets deal with possibilities generated from imprecise past or current events, while probability deals with the random nature of future events [Maiers & Sherif 1985]. More specifically, probability theory deal with precise sets to describe imprecise events, while possibility theory deals with imprecise sets to describe precise events. Hence, the degree of possibility may not necessarily correlate with the same degree of probability. Indeed, possibility theory may be regarded as superior to subjective probability. If a variable x can take values $u_1,\ldots u_n$ with respective possibilities $\pi=(\pi_1\ldots\pi_n)$ and probabilities $p=(p_1\ldots p_n)$, the degree of consistency of probability and possibility distributions may be given by $\gamma = \pi_1 p_1 + \ldots + \pi_n p_n$. Fuzzy sets have been applied to control systems [Tong 1977, Procyk & Mamdani 1979] but are particularly suited for representing inadequate knowledge of expert systems [Zadeh 1983, 1989]. Control involves the use of linguistic rules such as "if error is large and positive, then response is large and negative" and vice versa. These rules act as input-output functions. Fuzzy logic representation is characterised by fuzzy predicates, fuzzy quantifiers and approximate reasoning processes. Fuzzy quantifiers enable information to be communicated in a generalised way and are closely allied to linguistic variables. Fuzzy truth values are multi-valued in the real interval [0,1] but rather than being expressed as a specific value, it may be expressed as a possibility distribution over the unit interval. In this way, it can accommodate possible worlds similar to modal and default logics. The predicates are vague, eg. large/small, much, etc. Such predicate modifiers (hedges) include very/a little, more/less, etc. Quantifiers are also fuzzy: few, several, usually, most, often, etc. Possibilities are also vague: likely/unlikely, very possible/almost impossible. A proposition is regarded as a set of fuzzy constraints of the form "Q A's are B's" where Q is a fuzzy quantifier and A,B are fuzzy sets. Any natural language proposition is a collection of fuzzy constraints. Firstly, variable values which constrain the proposition are evaluated with their implied constraints. Each constraint c_i is characterised by a membership function $\pi_i \in [0,1]$ which represents the degree to which c_i is satisfied. All such partial membership functions are aggregated into an overall membership function $\pi=\Sigma\pi_i$. Certain rules apply for the translation of these constraints C with hedge score π.

Modification

- not-C membership is defined as 1-π
- very C membership is defined as π^2
- more or less C is defined as $\pi^{1/2}$

Composition:

- C_1 and C_2 hedge score is defined as $\pi_1 \wedge \pi_2$ (conjunction/minimum)
- C_1 or C_2 hedge score is defined as $\pi_1 \vee \pi_2$ (disjunction/maximum)
- if C_1 then C_2 hedge score is defined as $1 \wedge (1-\pi_1+\pi_2)$ - implication

Rules of inference provide for deductions similar to standard logical inferencing but uses linguistic variables defined by the membership function. The degree of membership of an element to a set defines its truth value. Fuzzy sets are also often used in neural networks – neuro-fuzzy networks – in which the first layer represents normal input variables, the last layer represents normal output variables, and the hidden layers represent fuzzy rules. The connection weights are fuzzy. The chief problem with possibility theory is that membership functions are not straightforward to define and are often empirical. Defuzzification usually proceeds by computing the area under the membership function and then the fuzzy centroid of the area. The weighted strength of each output membership function are multiplied by their respective output membership function centre points and summed. This area is divided by the sum of weighted membership function strengths to give a crisp output. It is not clear if fuzzy sets have any rigorous utility – truth values of True and False are discrete and fuzziness arises from imprecision in linguistic tasks. However, semantic fuzziness is a characteristic of the real world, its characteristics, relationships and categorisations.

3.13. HIGH-LEVEL COGNITION BY SYMBOL MANIPULATION

Autonomous control systems should be capable of performing well under significant uncertainties in both the system and the environment for extended periods of time without external interaction. The best way to achieve autonomy is to utilise high level decision making techniques, i.e.. intelligence to provide a predictive capability – intelligent control is the means the achieve autonomy in adaptive systems [Antsaklis 1991]. This allows the anticipation of the outcome of behaviour and allows goal-oriented actions to be implemented to achieve some purpose. This involves the process of internalisation of external reality to facilitate its simulation. Planning may be regarded as internalised motion which constructs the potential consequences of actions. The incorporation of intelligence at higher levels of the control system implies the need for flexible decision-making and reasoning. The ultimate goal of artificial intelligence (AI) research is to impart human level intelligence to computing machines. As pointed out by Sutherland (1986), AI has been characterised by different paradigmic approaches. Classical artificial intelligence (AI) methods are based on emulating higher human cognitive functions (physical symbolic systems hypothesis). Clearly, this makes the assumption that the process that we are attempting emulate is that reasoning and intellectual capabilities of the human brain – our intelligence that marks human beings from the rest of the biological complement of Earth. Turing (1950) suggested that a machine may be deemed intelligent if it passed an imitation game (the Turing test). The imitation game is played with a concealed computer pretending to be human whilst being quizzed by a human interrogator via a keyboard. If the computer succeeds in persuading the interrogator that it is human then it may be considered to be intelligent. Although Miller (1981) has questioned whether human intelligence is optimal, it is nonetheless the only model we have available to us. The Turing test is a highly behaviouristic approach to the notion of intelligence. The natural language program ELIZA that simulated a psychiatric interview often convinced test subjects that it was human [Weizenbaum 1966, 1967]. ELIZA simulated a psychiatric interview via a keyboard. ELIZA had no understanding of context and such understanding requires common knowledge of a shared culture. The scripts in ELIZA defined only a limited context and lacked true understanding as it merely used "tricks" that were suited to the Rogerian psychiatry context. A “conversation” was enabled by the program responding to particular types of statements in rigid ways: mention of a family member yielded the response “Tell me more about your family”; statements with “all” or “always” illicited a response requesting for examples; if statements did not fit into the prescribed patterns, the program yielded the response “Does that have anything to do with…” mentioning an earlier pertinent statement which began with “my”. This has been used to illustrate the limitations of the Turing definition, but the author feels that it shows the limitation of human beings and their anthropomorphistic tendencies. It could be contested that it is impossible to pass the Turing test without equivalent human knowledge about the world and human cognitive capabilities given a sufficiently comprehensive version of this test. Fatmi & Young (1970) defined intelligence as a faculty of mind by which order is perceived in a situation previously considered ordered. It is not clear exactly what this means but Barlow’s (1970) definition makes this notion of order more clear. In any case, attributing intelligence to “mind” specifically appears to deny the possibility of intelligent robots. Muller (1970) attempted a definition of intelligence thus: intelligence is "the act, process or result of reasoning from a part to a whole, from particulars to generals, or from the individual to the universal". This definition implies that classification ability is synonymous with intelligence, and indeed, the neural Darwinism paradigm would concur and widen this ability to all higher animals. Barlow (1970) defined intelligence as "the capacity to detect order among different sets of information, essentially reducing the redundancy between the sets of information" (Barlow 1983). Both of these definitions refer to the process of learning through generalisation (inductive inference). Critical to this ability is the form of representation of information in memory. Freitas & Gilbreath (1980) defined intelligence with an emphasis on behaviour - "it is the ability to formulate and revise patterns of order as evidenced by the eventual emergence of successful over unsuccessful actions". Indeed, Albus (1991) defined intelligence similarly as that quality which produces successful behaviour in a hostile environment. These definitions attribute intelligence in terms of its “niche-constructing” nature in manipulating the environment. However, successful behaviour does not necessarily require intelligence. The same applies to Coburn’s (1952) definition: “intelligence is a property of reacting on the basis of a probability as determined by the organism’s incomplete sensory samples of the environment”. Kim (1989) defined an intelligent system with a planning capability as a "purposive system (i.e.. a system that possess goals or objectives) that exhibits a rich variety of behaviours applicable to a variety of environments". The degree of intelligence in this case is defined by the range of behaviours. Adaptation is the modification of those behaviours on the basis of new information about the environment. Learning improves the system's variety of behaviours over time to attain its goals, i.e.. learning enables the system to become more intelligent based on its experiences. Sternberg (1985) defined intelligence as consisting of mental functions purposively employed for the adaptation to and the shaping and selection of real world environments which stresses the cognitive processes. This stressed the goal-oriented nature of intelligence with the degree of intelligence being defined by the range of behaviours with the use of goals in selecting those behaviours. Vincent (1992) defined intelligence explicitly in terms of functional survivability and planning capability: "intelligence generates the correct responses with a minimum of data in short timescales due to the benefit of prediction" Any intelligent system therefore must exhibit sensory integration, sensitivity and repair to damage, ability to make decisions, adaptability to environmental conditions and the ability to plan and predict events. In particular, of interest here, the ability to predict future events implies planning capability. Hewitt (1985) declared that several properties are required of an intelligent machine: it must be able to use partial and incomplete results, use due-process reasoning as well as logic, be capable of self-reflection, and it must perform efficiently. Self-reflection must include a history of its own behaviour, knowledge of its own information processing capabilities, knowledge of the causal relationship between its previous and current behaviour and knowledge of its ability to interact with the external world. Stonier (1988) included the concept of communal intelligence in his definition: intelligence is "a property of advanced information systems which allows such systems to analyse their environment, then engage in processes which enhance the survivability and/or reproducibility of the system". These definitions of intelligence illustrate the nebulosity and diversity of human cognition. There appears to be general agreement that there are two fundamental requirements for intelligence: learning capability and successful adaptive behaviour. However, pre-programming explicit programs for this does not appear to be feasible for generating adaptively successful behaviours in complex, uncertain environments. Hence, learning is an additional primary characteristic of intelligence [Wilkes 1953]. The human brain is primarily a control system which produces goal-seeking biologically advantageous behaviour to find food, avoid danger, compete for territory for survival and attract sexual partners and care for offspring for ultimate gene propagation. Whereas innate behaviour is wholly dependent on the organism's interactions with its environment, purposive behaviour may be only indirectly related to environmental conditions. Organisms are not passive receivers of sensory data but actively explore and manipulate the environment and this requires the imposition of constraints to provide environmental context. The world model must be geared towards objects in the real world that are (biologically for living creatures) important. The internal world model provides the means for "thought experiments" for simulating future events, thereby substituting for energy-expensive physical trials in selecting courses of action. Several properties are required of intelligent machines: they must be able to use partial and incomplete results, use due process empirical reasoning as well as logical deductive reasoning, be capable of self-reflection, and perform tasks efficiently. At a minimum, intelligence requires the ability to sense the environment, make decisions and produce actions to adjust to new situations and function reliably. In fact, this is precisely the definition of a control system. The ability to plan which implies the ability to predict future events forms part of this control system. Here we restrict ourselves to logic, reasoning and planning that form mainstream AI. However, we note that there are three types of knowledge acquisition [Anderson 1989]:

1. rationalism – knowledge that is derived internally through reasoning
2. nativism – knowledge that is innate (e.g. Chomskian language organ)
3. empiricism – knowledge learned through experience

Anderson (1989) suggested that declarative and procedural knowledge is acquired through learning, that the processes of induction and causal inference are innate (similar to the Chomskian language organ of the brain often associated with Broca's and Wernicke's areas of the left hemisphere of the human cerebral cortex), and that deduced knowledge is derived through reasoning algorithms abstracted from our causal inferencing capabilities.

Human behaviour has traditionally been modelled on the behaviourist paradigm rooted in stimulus-response studies whereby the environment is characterised by stimuli and behaviour is characterised by responses. Certain classes of stimuli evoke certain classes of behaviour. This behaviour is represented as an SR (stimulus-response) matrix where the matrix rows represent stimuli and the columns responses with entries specifying the relation between them (probability that a given response is evoked by a given stimulus). Bandura (1974) declared that human behaviour at least is cognitively mediated as well as being responsive to the environment, ie. both self-evaluative as well as externally driven. This gave rise to cognitive science concerned with the internal processing mechanism between stimulus (input) and response (output). Cognitive science has been based on an information processing paradigm known as the physical symbol hypothesis: it states that all aspects of human cognition are the product of information-encoding symbol manipulation within the brain. This paradigm is not open to empirical test in the Popperian sense, but is based on a doctrine of functionalism such that the physical medium be it brain or computer for such information processing is irrelevant to the processing mechanism itself. Cognition and thought therefore is synonymous with computation. The goal of cognitive science is to specify such mechanisms. Newell, Shaw & Simon (1958) first proposed such an information processing paradigm as a model of human problem solving:

SENSORS – PATTERN RECOGNITION – LOGIC PROCESSING (SHORT TERM MEMORY) – LONG TERM MEMORY – MOTOR EFFECTORS.

A central issue to such a paradigm is the role of logic: according to the model, memories store symbolised information connected by ordering relations which are operated on by programs. The crux is that organsmic behaviour is generated by programs. These programs are based on proof generation methods for symbolic logic theorems as a search for solutions to problems. Large problems may be broken down into a heirarchy of subproblems. Newell, Shaw & Simon (1958) suggested that such processes resemble human problem solving. It is debatable whether human thinking is really exclusively based on logical syllogism – much of human thought is a pragmatic device based on correct factual content through associations rather than logical structure, ie. that thought is knowledge-driven. Mental models of situations are generated in context dependent on content rather than through relational structures. Concepts also form a central issue in cognitive science [Smith 1985]. The modern view is that a concept is mentally represented by a prototype such that they are represented by properties that are true of some examples of the concept but not all. Generalisation capability is critical to survival and forms the core of inductive reasoning. Categorisation into concept classes depends on the similarity of the example and that similarity varies according to the particular example. A prototype has strong similarities to frames. Similarity is determined by the number of common features and the number of distinct features.

Ignoring subconscious processes about which little is known, mental processes may be classified as automatic or conscious. Automatic processes such as perceptual processing are characterised by inaccessibility to conscious control whilst consciously controlled processes correspond to higher cognitive functions such as reasoning, planning and problem solving. Conscious processing generates goals and motivations transforming passive information processing into active information processing. The difference is not tight however: some conscious processes can be automated with repeated practice such as motor skills. Automatic processes do not appear to be subject to capacity limitations whereas conscious appear to be inherently limited. These processing limitations are attributable to the limited capacity of 7±2 "chunks" of information within working (STM) memory: this limitation of computational resources is handled by the central executive implying a serial processing nature [Simon 1985]. Automatic processes proceed independently and in parallel. Serial processing is required for logically dependent analyses. Similarly, it appears that automatic processes are data-driven bottom-up processes such that low level raw data is transformed into higher level abstract information. This requires extensive computational operations implying special purpose hardware. Top down concept driven processing relies on high level expectations to interpret lower level information.

Although both connectionism and symbolicism are both information processing paradigms of human cognition, connectionism departs from the Turing machine model. The Turing machine model suggests that any continuous function can be computed by an algorithm to arbitrary precision, i.e.. any process or procedure may be representable as an algorithmic program implemented on a computer. Traditional AI is based on a number of fundamental assumptions [Kirsch 1991]:

- human thought processes are dominated by symbolic knowledge and the manipulation of symbolic knowledge through logical reasoning;
- cognitive processes are independent of the mechanisms of perception and motor actuation;
- cognition processes are linguistic in nature;
- cognition itself is independent of the learning process which installs knowledge;
- a single uniform architecture models all of cognition. According to traditional AI, intelligent behaviour may be represented by a computer program as a model of human reasoning [McCarthy 1961]. The implies that human intelligence is based on the theory of recursive functions (as all computing):
    - there exists a zero function which yields zero given any input;
    - there exists a successor function which yields the successor number given any input;
    - there exists an identity function which returns the identical value to the input;

There are three basic program functions to recursive computing:

4. composition of functions in serial fashion;
5. primitive recursion (for loop);
6. minimisation (while loop);

More complex functions may be constructed from these basic function.

The importance of symbolic processing is based on the Physical Symbol Hypothesis which states that symbol manipulation is a necessary and sufficient condition for general intelligent behaviour [Newell Shaw & Simon 1958; Newell & Simon 1976], i.e.. intelligence is based on rule-based symbol processing. A symbolic representation of a situation in the environment involves a transformation to map the outside world and the internal representation of it. Formal logic holds a central place in AI, the simplest form of which is propositional logic. Logic expresses facts independently of how they may be used in the manipulation of data sets. A formal logical system is a set of abstract symbols together with rules for concatenating these symbols into strings. The finite set of symbols is the alphabet. A grammar is a set of criteria which may be used to form finite sequences of symbols (formula strings) which are admissible in the formal system. There is a set of a priori strings defined for the system (axioms). There is a set of rules of logical inference which determine admissible combinations of strings. Each string following the axioms constitutes a theorem, and the finite sequence of inference rules applied to applied to construct each theorem constitutes a proof. Formal logic is defined syntactically and semantics is defined by the interpretation of the symbols outside the formal system. The formal system is essentially an abstract machine that outputs the list of theorems provable within the system. A formal system is complete if all the statements are permissible as theorems. A formal system is consistent if no contradictory statements can be proven. Truth is thus synonymous with provability. Every formal system can be computed as a sequence of logical instantiations (program) on a universal Turing machine (Church-Turing thesis). Logic essentially constructs formal representational models with the structure of the real world in terms of arbitrary conceptual objects (symbols) which have certain relationships between them. Inference engines are a major component of knowledge processing architectures such as expert systems and they employ logic-based deduction techniques. Knowledge representation almost exclusively uses a logical language. Traditional artificial intelligence methods are based on logical theorem-proving and mental models which encapsulate knowledge. Such top-down approaches which break down problems into hierarchies of sub-problems make assumptions about the task environment. Problem-solving is viewed as a process of a divide-and-conquer in which problems are broken down into simpler sub-problems hierarchically. The rationale behind such approaches are based on the analogy that artificial intelligence is to natural intelligence is as fixed wing aircraft are to birds. This approach is typified by symbol-manipulation architectures such as SOAR and ACT*. Both of these systems are universal computational architectures which attempt to model mental human intelligence, i.e.. the application of computational programming languages such as Church's lambda calculus (encapsulated in Lisp), Post's production systems (encapsulated in OPS5), von Neumann's stored program (encapsulated in C), Turing's universal automaton (encapsulated in augmented transition networks), or variations thereof (such as Prolog).

A logic relation is an attribute of a class of objects with respect to each other. Relationships between terms are represented as predicates and objects as arguments. Predicates define actions, properties and situations about the subject in symbolic form P(x) such that x has the property P. The inference rules of logic are used to create valid theorems. This is the basis of symbol manipulation. Logic makes no provision for exclamation, wishes, supplication, commands or questions. One form of logic is that of Frege whereby an object o with properties may be represented as an axiom of the form: $Has(o,"P") \leftrightarrow P(o)$. The properties are the predicate and the axioms state the relationship between the object and its properties. The atomic formula is the expression of a relationship among terms of the atomic formula. They may be combined by means of the logical operators AND ($\wedge$), OR ($\vee$) and NOT (~) into compound well-formed formulas (if...then implication is composable from these operators, $\sim A \vee B$ is equivalent to $A \rightarrow B$). First order predicate calculus allows the inclusion of the existential and universal quantifiers: $\exists$ "there exists" and $\forall$ "for all" respectively. Both quantifiers are related to each other:

$$\forall x P x = \sim \exists x (\sim P x)$$

$$\exists x P x = \sim \forall x (\sim P x)$$

Problems are solved as conjectures through theorem proving. Each inference rule selects statements which satisfy syntactic rules and generates a conclusion which is a logical consequence of those statements. Proof is a succession of well-formed formulas where each formula follows from the preceding formula according to the rules of inference. The theorem is the last line of the proof. Theorem proving can compose operators that cause changes of state into state transition sequences. Hence it may be implemented as logical inferencing and as rule statements.

3.13.1. Role of Logic in Artificial Intelligence

There is a standard form of first order predicate logic (Horn clause logic) that limits the conclusion of an implication to at most a single atom. Automated reasoning for the discovery of concepts employs theorem proving based on formal first order predicate logic [Wos 1985]. Logical operators include "and", "or", "not", "implies" and universal and existential quantifiers. Well defined inference rules are applied to the clause statements of logic. Each inference rule selects statements that satisfy syntactic rules and generates a conclusion which a logical consequence of those statements. Proof is a succession of such statements. Horn clause logic is computationally equivalent to the universal Turing machine. Atomic formulas which specify hypotheses however may be combined conjunctively: $H_1\&...\&H_n$ where H_i is the condition of the hypothesis. This has the advantage that logical deductions are carried out automatically. PROLOG is a logic programming language based on Horn clause logic and PROLOG rules are similar to context free grammar rules in that they have one element in the consequent, and several elements in the antecedent [Genesareth & Ginsberg 1985, Davis 1985]. PROLOG allows facts about objects and relations to be represented as sentences in its formal logic syntax. PROLOG has generally superseded LISP which was based on Church's λ-calculus. The λ-calculus was developed as an alternative to set theory for the foundation of mathematics. Unlike the Turing machine which is sequential, λ-calculus is an abstract model of computation that retains implicit parallelism. LISP is a functional language that is procedural in nature while PROLOG is a rule based language that is declarative in nature. Inferencing is achieved through Robinson's resolution principle. This allows the use of only one powerful and complete rule of inference for theorem proving for high speed inferencing. It is an inference rule that is widely used for automated theorem proving and logic processing to determine the validity of first order logic statements [Green 1969]: *if* $\exists$ *an axiom of the form* $E_1 \vee E_2$ *and* $\exists$ *another axiom of the form* $\sim E_2 \vee E_3$ *then the resolvent* $E_1 \vee E_3$ *follows logically*. Hence, it is a proof method that proceeds by refutation. The statement to be proved is negated and the negated statement is resolved with a series of known true propositions. If this produces a contradiction, the negation is false and the original proposition is true. Resolution as a rule of inference is sound and complete but it is subject to combinatorial explosion with increasingly complex problems. Unification adds more power to resolution by extending its use to non-identical atoms through substitution. Given a clause with atom p as its conclusion and another clause with p in its condition, resolution creates a new clause with a conclusion that is the union of the conclusions of the original clauses and a condition that is the union of the conditions of the original clauses with the deletion of p from the expression, eg.

Given:

p(A,B):-f(A,B)
g(A,B):-p(A,B),p(B,C)
Concludes:

g(A,C):-f(A,B),p(B,C)

Resolution can also be used when two atoms are not identical as long as they can be made syntactically identical by the appropriate instantiation of the variables by substitution - this unification process find a set of bindings for the variable in the two expressions to make them syntactically identical by substitution:

Given:

p(x,B)
p(A,y)
Substitutes:

x/A, y/B
Concludes:

p(A,B)

Hence unification with resolution allows complex conclusions to be drawn. The unification procedure is essentially a process of generalised syntactic pattern match. It applies the most general substitution that when applied to the atomic formulae in a given set renders them syntactically identical. For example:

Given:

p(x,B):-f(x,B)
g((A,z):-p(A,y),p(y,z)
Concludes:

g(A,z):-f(A,B),p(B,z)

If unification fails then backtracking occurs automatically. One difficulty that PROLOG does have which is its main failing is that it suffers from the negation as failure problem in that it shows ~P by attempting and failing to show P. However, this is a fundamental problem in logic generally.

Approaches other than resolution have been advocated for use in theorem proving particularly domain dependent heuristic knowledge and procedures to limit the search for proofs [Bledsoe 1975]. An example of one such heuristic are demons which are routines that monitor the knowledge base and act only when certain properties become true of a knowledge base. This means that some parts of the knowledge base are static while others are dynamic depending on the context. In essence part of the inference process becomes domain specific rather than just relying on a few powerful general inferencing methods. The use of heuristic knowledge and procedures essentially provides advice for solving specific types of problem. One of the first AI programs was Logic Theorist which used heuristics to prove 38 theorems of the 550 theorems in Russel & Whitehead's *Principia Mathematica* given 5 axioms and just 3 rules of inference [Newell, Shaw & Simon 1956, 1958]. The first axiom of *Principia* was $A \vee A \rightarrow A$. It defined numbers as classes of classes with the same members. The number zero was defined as the class of all classes with no members. There is one class which has no members thus defining the number one. Two can be defined as the class of classes with as many members as the class whose members are zero and one. This process can be continued to yield the set of natural numbers. The Logic Theorist derived these new theorems using a "mechanical mathematics" constrained by only a few heuristics [Minsky 1961].

The culmination of this approach is to use models to aid theorem proving and their application to analogy in knowledge discovery. The automatic mathematician (AM) program developed recently was capable of formulating new theorems in mathematics from 100 elementary axioms of set theory [Lenat 1978]. The AM program utilised an exploratory search of a very large search space constrained by a large body of 216 heuristics in a process of discovery. The heuristics dealt with concepts in set theory. The program emulated the scientific process in unifying large sets of empirical data and could predict new effects. The heuristics were the central component and formed a hierarchy of generalisation/specialisation concepts applicable to the mathematics domain - many of the more specific heuristics were specialisations of the more general ones. These heuristics were used to guide the open-ended search and discovery of new mathematical concepts. Each concept was represented as a frame data structure with slots for binary relations which may assume values. AM was given an initial collection of 115 core concepts of set theory as axioms tacked with the most general relevant heuristic rules that were applicable. Each more specialised concept inherited those rules as being applicable from the more general concept. The program derived the concept of numbers then concept of multiplication in 4 separate ways and then derived the concept of divisors, then perfect squares, and prime numbers and natural numbers including the theorem that every natural number is the product of a unique set of prime numbers (the fundamental theorem of arithmetic). A similar program was BACON which discovered certain scientific laws including Kepler's third law, Ohm's law, Snell's law, Archimedes law, Gay-Lussac's law and Dalton's laws governing chemical reactions whilst at the same time discovering the concepts of inertial mass, specific heat, refractive index, atomic weight, molecular weight and the conservation laws [Walker & Simon 1983]. BACON was data driven being given numerical data inputs from experiments but it employed little search and used a small number of heuristics to search the data for constants, common divisors and differences, etc. It tested for correlations among pairs of variables and when significant correlations were found, BACON tested the product or ratio of the variables for invariance and then introduced these as new variables or constants. The heuristics guided the procedures by trying laws that were symmetric and looked for numbers that were simple multiples of each other. This learning through scientific discovery emulates concept formation by the accretion of data through inductive inference which had the following properties [Lenat 1983]:

- gather empirical data about a phenomenon;

- extract regularities and patterns in the data;
- form hypotheses and experiments to test them;
- abstract new heuristics by compiling general theories as paradigms.

AM and similar programs have since been criticised on the grounds that it is only through the intimate relationship between mathematics and LISP programming language that such programs work thereby limiting automated discovery to mathematical domains only [Lenat 1984]. However the programs have been applied to other domains and it was suggested that they could even be applied to artistic appreciation through an artistic heuristic based on symmetry, unconventionality and utility: IF two apparently disparate parts of the work of art are recognised as being closely related THEN this is an example of harmonious art.

Logic is an inefficient means for modelling intelligent behaviour as logic cannot tolerate contradiction (and much human behaviour is contradictory). Intelligent behaviour is implemented by agents and agents possess physical bodies, goals and motives and perform activities which are selected to satisfy those goals and sate those motives. The body itself is a kind of memory of environmental information (extracted through millions of years of evolution by natural selection – this is the basis of evolutionary psychology as a discipline). Goals represent knowledge about the required behaviour whilst interacting with the physical environment through the body.

3.13.2. Self-Reference in Logic

An important issue in logic concerns self-reference – an agent's representations and reasoning concerning its own structure and behaviour. A robot is in fact part of the environment in which it acts and planning requires an account to be taken of this for the correction of errors [Perlis 1985]. Self-reflection includes a history of the agents' own behaviour (episodic memory), knowledge of its own reasoning processes (introspection), knowledge of causal relationships between its previous and current behaviour (stream of consciousness) and knowledge of its ability to interact with the external world (distinguishing self from non-self). Language accomplishes self-reference through naming distinguishing between names for concepts (second order constructs) and the concept itself (first order objects). The second order construct can be embedded into the first order system as another concept. Concepts carry the meaning but names are referential. Naming is distinguished by using quotation marks: believe I "(a is b)". However, naming devices are prone to inconsistency through the no-truth definition theorem: $True("A") \leftrightarrow A$. This is inconsistent due to the Cretan Liar paradox based on the famous Epimenides the Cretan stating "All Cretans are liars": $True("L") \leftrightarrow L \leftrightarrow \sim True("L")$ with $A = \{L{:}\sim True("L")\}$. This is equivalent to "This sentence is false" - if the sentence is true, the assertion is false, and if the sentence is false, the assertion is true. Such paradox's arise from self-reference whereby the description refers to something that affects the validity of the description. The only way to overcome this is to assume that for some formulas A neither True ("A") nor False ("A") hold, i.e.. the law of the excluded middle does not hold. Hence True ("~A") is not the same as ~A and True ("~A") is stronger than ~A. Although $A \cup \sim A = U$, $True("A") \cup True("\sim A")$ is not required to equate to the universe of discourse. This is accentuates the difference between a concept and its name. This is the difference between belief Bel(A) and truth True(A). This effectively introduces a hierarchy of truth predicates. This necessitates a metalevel/object level distinction between the two levels of truth. Without this distinction, certain concepts or theorems cannot be expressed [Aiello et al 1986].

To overcome the problem of self-contradiction of the Liar paradox, the Tarski Undefinability Theorem suggests a correspondence theory of truth. There is a distinction between an object language and metalanguage [Tarski 1969]. The Liar Paradox is true for the object language hen stated in the metalanguage as an external extension of truth about the language. The object language describes the context of discussion for which a definition of truth is required. From the definition of truth, it is possible to derive the law of contradiction that no two sentences one of which is the negation of the other cannot both be true, and the law of the excluded middle that no two sentences one of which is the negation of the other can both be false. The metalanguage is the language in which the definition of truth is formulated: ("p" is true) - "p" is a sentence in the object language whereas the assertion is a sentence in the metalanguage. The metalanguage includes the object language as a subset such that every sentence in the object language is in the metalanguage. The metalanguage does not have to be semantically universal in that it need not have names for its own components but must have names for the linguistic entities of the object language. In this sense metaknowledge is knowledge about knowledge, metareasoning is reasoning about reasoning, etc. It is possible to communicate between metalevel objects through object level semantic networks of object level linking rules. FOL (First Order Logic) was an expert system whose expertise was reasoning – it was based on the distinction between an object language L and an object simulation structure S. This modelled the distinction between language syntax and its semantic interpretation [Weybrauch 1980]. The language was the object level theory while the simulation was the metalevel model. Metalevel constructs enable self-reflective reasoning through self modification and self-evaluation through the simulation capability to foresee its own behaviour. Self-evaluation involves reasoning at the metalevel about computational complexity and inferencing resources.

Self-reflection generally requires episodic memory. Self-reflection generally requires episodic memory. Metalevel reasoning can provide a mechanism for the formulation of beliefs, default reasoning and changing inferences as a learning process for adapting to changes in the environment. Metacognition as awareness of an agent's thinking and behaviour provides the basis for consciousness – it is critical self-assessment.

3.13.3. Planning – the Importance of Predicting the Future

Plan generation is basically the process of specifying the ordering of actions to achieve desired goals. Plan generation is very much part of the control problem in that it generates the input to the control system as a sequence of commands to achieve the desired goals and monitors the plan execution for feedback. The artificial intelligence paradigm is based on storing knowledge as a set of symbols which are manipulated according to a set of formal algorithmic rules – it is based on the assumption that all knowledge can be so formalised, i.e.. automated logical reasoning [Simon 1978, Post & Sage 1990]. Plan generation has been traditionally viewed as the generalisation of the programming language process in that the plan represents a possible solution to a problem through the execution of an effective well-defined procedure, i.e.. planning is an algorithmic problem solving process. The purpose of planning is to generate a sequence of actions to achieve given objectives and goals. Plan generation as generalised problem solving may be characterised as a state space search for actions which when applied to the outside world change it into a new goal state. Indeed, AI methodology has been characterised *in toto* as search since much of it is concerned with the exploration of a range of possible actions or inferences in the pursuit of well defined goals or solutions [Charniak & McDermott 1985]. Logical assertions may be used as symbolic structures to describe the state of the universe of discourse (the world model) with respect to the effects of all actions expressed as logical implications. First order predicate calculus is usually used as the means of describing the world's state. The problem space comprises a set of symbolic structures or states of the space and a set of operators to change states. The initial world state represents the global representation of the problem to be solved. Each action generates new possible actions creating a branching search tree of possible states. The complete plan is a path through the search space from the initial to the final desired goal state representing a sequence of actions tracing a path from the root of the tree to the desired end node (the solution). The actions act as operators which provide the means for conditional operations to change the state of the world. Hence, operators by representing the effects of actions on the state of the world are expressed as logical implications based on rules of inference such as Robinson’s resolution principle. Such operators comprise of three parts: the precondition list, the effect list and the body of the operator which defines the logical relations between the condition and the effect. In this way, preconditions decide the applicability of operators, effects determine the effects of their execution, and their bodies describe the means to achieve the effects in relation to the world model. An operator is applicable if the operator's preconditions hold in the world model. Planning is a search through the space of inferences to find a temporal sequence of operators such that the first operator in the sequence is applicable to the initial world model and the final operator produces a world model in which the goal is true, i.e.. the world model is transformed through a series of successive states linked by the operators (actions or inferences). Preconditions must be satisfied before the goal can be attained. Each precondition is essentially a subgoal which is further reduced by further branching of the goal tree until the terminal goal state is reached. This process is a process of problem reduction or subgoaling using state space search to find the sequence of preconditions that must be met before the goal can be achieved. It is a divide and conquer strategy. This hierarchical approach to planning is characterised by simplified tasks which decompose into finer grained detailed subplans. The plan is then a representation of actions for achieving the goal or task. State spaces can be searched and several methods exist to find a path from the root of the tree to the desired goal. The popular graph-traversing A* algorithm (and its variant the D* algorithm) can be used to find the minimum cost solution, i.e.. using the minimax principle. It is a branch and bound graph search algorithm which uses an evaluation function to order nodes, combining lower bound remaining distance estimation with dynamic programming principles and is effective at dealing with average case situations though it fares less well under worst case scenarios. The planning process may be defined as a sequence of states to be achieved represented as a tuple [Georgeff & Lansky 1986]:

$$P = \langle N, E, \delta, n_i, N_f, \alpha \rangle$$

where:

N = *set of subgoal states*
E = *set of decomposition rules between subgoals*
$\delta: N \times E$ = *process control function*
$n_i \in N$ = *initial state*
$N_f \in N$ = *set of goal states*
$\alpha: E \rightarrow A$ = *function associating actions with decomposition rules*

The sets of actions may be combined with conjunctive (AND) and disjunctive (OR) operators to form AND/OR graphs. Process associations of the form c<P>g consist of the process description P, the precondition c of world states that must be satisfied for the process to be applicable and the effect g of the process generated behaviours on the world.

The domain-independent GPS (General Problem Solver) system represented the states of the world as logical assertions of first order predicate calculus. GPS used means end analysis as a domain-independent heuristic to compare the current initial world state (the problem) with the desired final world state (the solution) and then search for procedural operators that reduced that difference. Means end analysis was used recursive to incrementally bring the current state closer to the goal state. GPS used depth first search and forward chaining from the current state to the goal state using its "find-and-reduce-difference heuristic". GPS was the basis of STRIPS (STanford Research Institute Problem Solver) planner [Fikes, Hart & Nilsson 1972]. STRIPS has provided the basis for many subsequent planners and it has a central role in planning systems [Simon 1991]. The state of the world was modelled as a set of first order predicate logical clauses. Elementary robot actions were defined with preconditions for their applicability to the situation, and consequences which defined the action on the world. These formulas represented theorems about the state of the world. Each action acted to change the state of the world towards the final goal state by describing the actions in terms of preconditions and effects. The sequence of actions selected thus comprised a search for the correct sequence of actions to achieve the goal.

The operator was characterised by an ADD-list which listed the consequent changes to add to the world model by the application of the operator, a DELETE-list which listed antecedent deletions to the world model by the application of the operator, and a PRECONDITION that provided the context of logical formulae to be fulfilled so that the operator may be applied. The ADD and DELETE list effectively described the assertions that would be true after execution, i.e.. the effects. These properties of each operator localised the effects of each operator. Plans were built from these operator schemata in an ordered sequence through the search space of intermediate states from the initial to the final goal state by transforming the world model from one intermediate state to another. The operators were essentially sets of condition-action propositions like production rules. The conditions defined the circumstances that must exist before the operator may be fired. The actions define the effects on the world and the body of the rule describes the decomposition of actions into subactions. The plan was a tree of nodes representing a hierarchy of actions and their temporal order down to the level of primitives. Propositions represented the initial state and goal state of the world and a library of possible actions. Most planners are based on STRIPS which in conjunction with a plan execution module PLANEX was implemented on the robot Shakey which operated in an artificial world of coloured blocks. Such supervision of the execution of a plan is required – the supervisor can implement reflexive behaviours to react to unexpected events – it is at this level that the control loop is closed. The planning cycle would take 5-10 minutes of computation time limiting its applicability to near-static environments.

The environment is constantly changing and reasoning in the real world must reflect this. STRIPS and similar planning programs suffer from the "frame" problem which arises out of the dynamic nature of the world [Hayes 1973]. In the process of planning, each action changes certain specific features of the situation. The frame problem arises from the need to represent those aspects of the world which remain invariant during state changes. The difficulty is in keeping track of the consequences of performing actions and the required alteration to the world model representation in real time. The frame problem highlights that it is difficult to differentiate between the outside world (auβenwelt) and the part of the outside world in which the organism or robot interacts (unwelt) in von Uexkull’s terminology. In addition, as the world model becomes more complex so the computation time required to update it explodes exponentially. It is necessary to avoid re-inferring assertions for each new instant of time. Since successive states of the world are logically independent, the difficulty is in deciding which statements continue to be true and persist after the performance of a given action. Indeed, generally almost all the assertions will continue to be true after the action has been applied as most of the world is unaffected by events as they tend to make only localised changes to the world. It is necessary to state that actions change only features of the situation to which it directly refers. STRIPS localises the effect of actions so that almost all the statements continue to be true to the next state unchanged unless explicitly eliminated. One possibility is to include "always true" frame axioms to explicitly represent the invariants but this limits flexibility and nullifies the inherent logical correctness of the planning process. Furthermore, the number of such frame axioms becomes very large if the world model is complex. The performance of actions must qualify what does not change when actions are performed through frame axioms. This also introduces the possibility that frame axioms may become invalid as they will vary from the type of state change. This approach was used in STRIPS and many subsequent planners. If the components of the world interact, an action can have a global effect on the world, eg. if a block is withdrawn from a stack, the whole stack falls. Hence, it is important to note interactions and relationships between the world's components. The chief problem of such techniques being applied to planning in real world environments is the frame problem which has plagued such techniques since their inception (STRIPS). The frame problem suggests that there is a fundamental mis-mapping between the representation scheme which is based on logic (in which propositions are statically true or false) and the changeability of the real world environment (in which propositions alter their truth value). Assessing which aspects of the environment change on the basis of actuations on the world are non-trivial.

Ultimately, plan generation may be regarded as a form of self-programming and this is a major goal in robotics to develop autonomous capabilities to machines. Church's λ-calculus was developed as an alternative to set theory for the foundation of mathematics. Unlike the Turing machine which is sequential, λ-calculus is an abstract model of computation that retains implicit parallelism. Self-programming through automated program synthesis may be possible by regarding programs as algorithmic mathematical expressions. Such synthesis requires both programming capability and knowledge about the subject domain. A program generator when given the specification of a problem will produce an encoded program (i.e.. planning). Self programming is based on formal software validation theory. This is a major difficulty in software engineering. Formally a program Q essentially provides a functional relationship between the precondition P and the results of its execution R of output data sets f(x): P{Q}R. The program Q maps P onto R. The automatic generation of programs using propositional logic theorem proving techniques involves deriving that $P(x_1 \ldots x_n) \to R[f(x_1) \ldots f(x_n)]$. The specification Q for a program P describing the intended behaviour of the program is given as input-output assertions [Manna & Waldringer 1975, 1978]. Structured programming is developed by the successive application of transformation rules to the specification by replacing non-primitive constructs of the specification by primitive constructs of the programming language. Structured programming languages employ three constructs: sequences, if-then rules, and do-while rules (go-to is forbidden). Conditional expressions such as (if p then q else r) deal with uncertainty by splitting the world into two parts to reduce the uncertainty in each – one in which p is true, and in which p is false. Recursion or looping involves generating and solving subsidiary problems. The purpose of assessing the validity is to ascertain whether it carries out its intended function correctly and this can be done using mathematical logic [Hoare 1983]. There are three types of program command:

1. axiom of assignment such that variables can be assigned expressions: $\perp P_0\{x := f\}P$;
2. inference rule of consequence such that assertions can be formed as theorems: if $\perp P\{Q\}R$ and $\perp R \supset S \& \perp T \supset P$ then $\perp T\{Q\}S$;
3. rule of iteration to allow recursion: if $\perp P \wedge B\{P\}P$ then $\perp P$ {while B do S}$\sim B \wedge P$.

Goals may be generalised to match lower level goals: let P{f(a)} be a subgoal of f(x) where f(a) is a recursive call. By proving that $P\{f(a)\}\forall x$, satisfaction is guaranteed. No proof can exist to show that a program will successfully terminate (due to the possibility of infinite loops), i.e.. Turing's halting problem. However by recasting the program as a mathematical logic system a program can be proven to yield the required function provided the inputs are suitable. Lehman (1980) classified programs into two classes. S-programs are those whose function are formally defined by a specification from which the program is derivable through valid logically connected sequences of statements. They are essentially static and large S-programs can be easily decomposed into hierarchical structures. A-programs are specified by formal rules and under-deterministic decision rules which are dependent on feedback from the real world environment. These programs alter according to the dynamic state of the world. Any uncertainty is either implicitly included or explicitly represented in the specification. It essentially conforms to Hewitt's (1985) ideas that search needs to be exploratory to cope with ill-defined and changing goals. Automated program synthesis is still in its infancy, but a variant – genetic programming – shows promise although the logical structure of such programs are not easily extracted. If algorithmic planning is viewed as an automatic programming system where subroutines comprise the subgoal hierarchy then three possible outcomes can occur: success in generating a plan, failure when no plan is found after exhaustive search, or non-termination as the plan expands ad infinitum never converging to a solution (the equivalent of Turing's halting problem). If a plan is produced it correctly solves the problem but if failure or the non-termination state is generated no solution exists. Since any algorithm with its inputs can be encoded as a Turing machine, the planning problem is undecidable and no upper bound exists on the time required to solve a problem due to the existence of the possibility of infinite loops [Chapman 1987].

Newell's SOAR (Smalltalk-On-A-RISC) is an extension and generalisation of the GPS architecture system for general problem solving. It is implemented on RISC processor hardware and uses Smalltalk, an object oriented programming environment. It is an approach to derive a unified computational architecture for general intelligence particularly for goal oriented problem solving [Laird, Newell & Rosenbloom 1987]. The RISC computer employs only the most frequently used instructions in hardware and less used instructions in software. SOAR is firmly rooted in the Physical Symbol Hypothesis. It is capable of working on a range of tasks using a range of problem solving techniques. It is essentially a production rule based expert system which uses a heuristically guided A* search to find operators in the problem space that can transform the current state to a desired goal state through the generation of hierarchies of subgoals to resolve impasses that occur when knowledge is incomplete. In effect, the use of production rules as a universal symbol processing mechanism re-introduces stimulus-response behaviourism to artificial intelligence. The A* algorithm is a graph search which uses the evaluation function for ordering nodes. All long term knowledge is stored as production rules based on OPS5 in a uniform declarative long term memory. The condition-action production rules include default knowledge and when they are fired they are held in working memory and describe the processing state of the problem solving process. The production rules are fired in parallel without conflict resolution and decision making occurs in working memory. This knowledge in working memory is used to aid the search and subgoal decomposition through means end analysis. Impasses occur when contradictory production rules are activated in parallel to create a new problem to be solved. The generation of impasses effectively implements subgoaling until subgoals are generated that can be solved trivially and deeper impasses can be resolved. This general mechanism is viewed as the characteristic of problem solving intelligence. This impasse resolution implements a single generalised mechanism for a limited form of learning from experience on the basis of its performance. Previously solved impasses are stored by "chunking". Impasses are generated when procedures fail at which point new goals are generated through the chunking process. Chunking allows it to learn operator generalisation through the variablisation of constants into macro-operators for complete goal or subgoal specification to reduce future search requirements so implementing more reflexive behaviours. It uses preferences to represent knowledge of how to continue from its current state. These preferences decide acceptability, unacceptability, preference or indifference to procedures. It is a domain independent problem solver which may be applied to any domain dependent knowledge based expert system.

The chief problem of STRIPS and all its variants is that they are capable of dealing with only simple planning problems. Problems such as the Tower of Hanoi problem which require overcoming obstacles to goal reduction are not soluble by these methods based on reducing searches. The Towers of Hanoi problem comprises of 3 pegs and a set of different sized disks stacked on one of the pegs in decreasing order of size. The problem is to transfer all the disks from one peg to another with the constraints that:

- only one disk is moved at a time;
- the larger disk cannot be placed onto a smaller disk.

These planners were also restricted to simplistic "blocks-world" type environments – SHRDLU was a planning program similar to STRIPS designed to demonstrate natural language communication which operated in a limited "blocks" microworld of simple 3D objects such as cubes and pyramids on a flat surface in an attempt to tackle the problem of context [Winograd 1980]. Another important problem is that of the conjunction of interacting subgoals whereby if any one of the subproblems is solved independently, the solutions to others may be destroyed [Chapman 1987]. Some goals are conjunctive goals which satisfy several constraints simultaneously. The traditional strategy for achieving goals P and Q is to first derive a solution to achieve P first then modify it to achieve Q also. But this will not necessarily work as actions can have side effects (variant on the "frame problem"). Subgoals may also have different priorities which may also in turn change according to the situation. Motor actions may have long-term consequences which may affect future motor actions. It is often necessary to achieve several goals simultaneously particularly in complex tasks. Such goal interactions are ubiquitous in planning. One strategy for dealing with this is to reach subgoals independently and then treat interactions as they arise. Re-ordering of subobjectives will eliminate the interaction problem in many cases, eg. Sussman's HACKER which was based on SHRDLU. HACKER achieved this through a form of trial and error learning. A least commitment strategy of partial plans is the best approach for dealing with conjunctive goals. Decisions of ordering are taken only when interference with other plans will definitely not occur. General goals are addressed first with details being filled in later. Hence, planning may be linear whereby the sequence is totally ordered or non-linear whereby the sequence is partially ordered.

Given the dynamic nature of the physical world, these approaches necessitate the repair of plans which may lead to infinite looping cycles. As the world changes over time, it is necessary for the planning system to represent time. Time is a major problem for planning systems that regard the world's evolution as a sequence of states with actions causing state changes in the world. A plan is essentially a representation of some aspects of the future. Most actions take time and several can occur at once. Time specifies when conditions should be achieved and for how long they should be maintained. However, many temporal constraints are imprecise due to their relative nature, a lack of a clear relationship between temporal markers and different temporal scales. Planning of activities occurs in a temporal universe with deadlines and external events. Each activity is defined by a temporal window with a start time and a duration. Since much temporal knowledge is relative rather than involving absolute dates and times represented by temporal coordinates, Allen & Litman (1986) introduced a temporal logic based on the temporal interval to accommodate inexact relationships between intervals through a disjunction of primitive ordering relations such as $X<Y$ (X before Y), $Y>X$ (Y after X), $X=Y$ (X concurrent with Y), and XoY (X overlaps Y). This allows the formation of a constraint network comprising directed acyclic graphs of nodes representing interval connected by arcs which relate the intervals.

3.13.4. <u>Knowledge-Based Expert Systems</u>

Planning viewed as a search problem suffers from certain problems. Many cognitive functions such as search (such as the A* and D* algorithms) invoke NP hard type problems which cannot be readily solved in real-time. This is similar to the "travelling salesman problem" (TSP) who must find the optimal route between n cities visiting each only once – indeed, TSP is algorithmically fundamental to the whole class of NP problems. As there are n! possible combinations of routes to search, this combinatorial explosion becomes prohibitive even for small n. The search space can become large and grow exponentially and searching for a path can involve exponential computation. Furthermore, finding an optimal plan in even a simple blocks world is an NP hard problem and so fundamentally intractable due to the conditional dependency of effects of actions on the input situations and the frame problem of discovering which propositions are unchanged by an action and derived side effects. When search trees become large, the computational cost of solution may become impractically large. This suggests the use of knowledge to prune search trees (knowledge principle) [Newell 1982]. Such expert systems encode knowledge as situation-action production rules (or similar data structures) in memory about the task environment. However, they are mostly concerned with semantic domain knowledge captured from experts (expert systems) which comprises a computational model for useful functions such as medical diagnosis or geological prospecting. AI has thus undergone a paradigm shift from power-based inferential search techniques to knowledge-rich strategies which offer a means to reduce the search requirements. This shift represents the Knowledge Principle whereby search and reasoning alone is regarded as insufficient for intelligent behaviour - a great deal of knowledge of the world in which the agent operates is required [Lenat & Feigenbaum 1991]. Knowledge is one of the prerequisites for intelligent behaviour and in effect represents the agent's learning history. Much of this knowledge should be sufficiently general to enable its use for multiple but specific domains (the Breadth Principle). Such knowledge imposes descriptions of regularities in the world including declarative maps, planning constraints, etc. Expert systems are programs that solve problems normally requiring the knowledge and skills of human experts [Hayes-Roth 1984, 1985a, b, Hayes-Roth & Jacobstein 1994]. Essentially, they model human decision making processes through the use of knowledge bases. This is model-based reasoning which draws heavily on world knowledge rather than logical inferencing and assumes that the mental models correspond to the real world. McCarthy & Hayes (1969) considered that intelligence absolutely requires an adequate model of the world to enable the agent to perform tasks in the external world according to its goals and abilities and if necessary add information from the external world observed through the senses to the internal world model. A world model M has the structure:

$$< \theta, P, E, A, T, B, G, \varphi >$$

where:

θ = *a set of actions*
P = *set of other agents*
E = *set of primitive events*
$A \in E \to P$ = *agent of events*
$T \subseteq E$ = *set of possible events/worlds*
$B \subseteq T \times P$ = *set of beliefs*
$G \subseteq T \times P$ = *set of goals*
φ=*predicate interpreter.*

Problems are described using objects, attributes and relations. The objects are data structures which represent knowledge about conceptual entities. Although incomplete, the world model uses past information to predict the future and constrains subsequent sensory processing based on expectations. Intelligence integrates knowledge and sensory feedback information in a goal directed control system that plans purposeful actions. Furthermore, purposiveness requires anticipation or expectancy concerning the choices in achieving goals. Albus (1999)'s NASREM derivative model is another example of such an intelligent robotic architecture based on these notions – he suggested that full human level intelligence requires a number of different modular components encapsulated in a number of axioms:

- the function elements of an intelligent system are: behaviour generation (planning and behaviour generation), sensory perception (filtering of perceptual data according to expectations), world modelling (to generate expectations from internal world models), and value judgement (to assess costs and risks associated with generated behaviours, i.e.. emotional colouring);
- an intelligent system must be supported by a knowledge base that stores *a priori* and dynamic information about the world – the symbol-based world reference model;
- an intelligent system may be implemented by a set of interconnected computational modules forming a control architecture;
- the complexity within an intelligent system may be accommodated through hierarchical decomposition layering whereby upper levels of the hierarchy are characterised by long-range strategic planning and lower levels being characterised by short-term tactical activity to combine goal-driven and reflexive behaviours;
- the complexity of the real world environment can be managed through focussing attention to control sensing and the diligent deployment of resources.

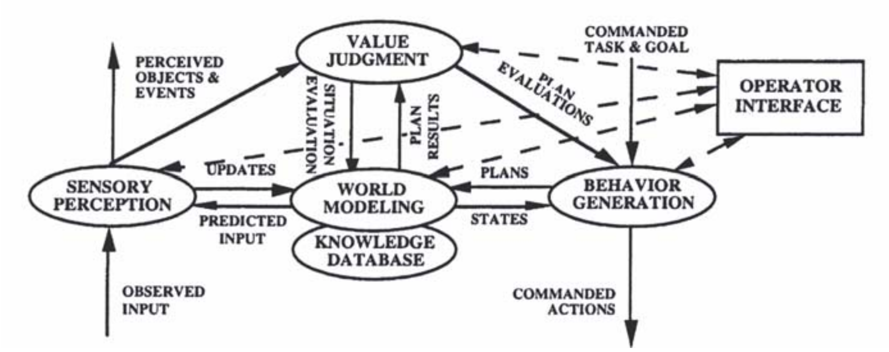

Figure 210 – Functional modules of NASREM include sensory perception (SP), behavior generation (BG), world modelling (WM) and value judgement (VJ)

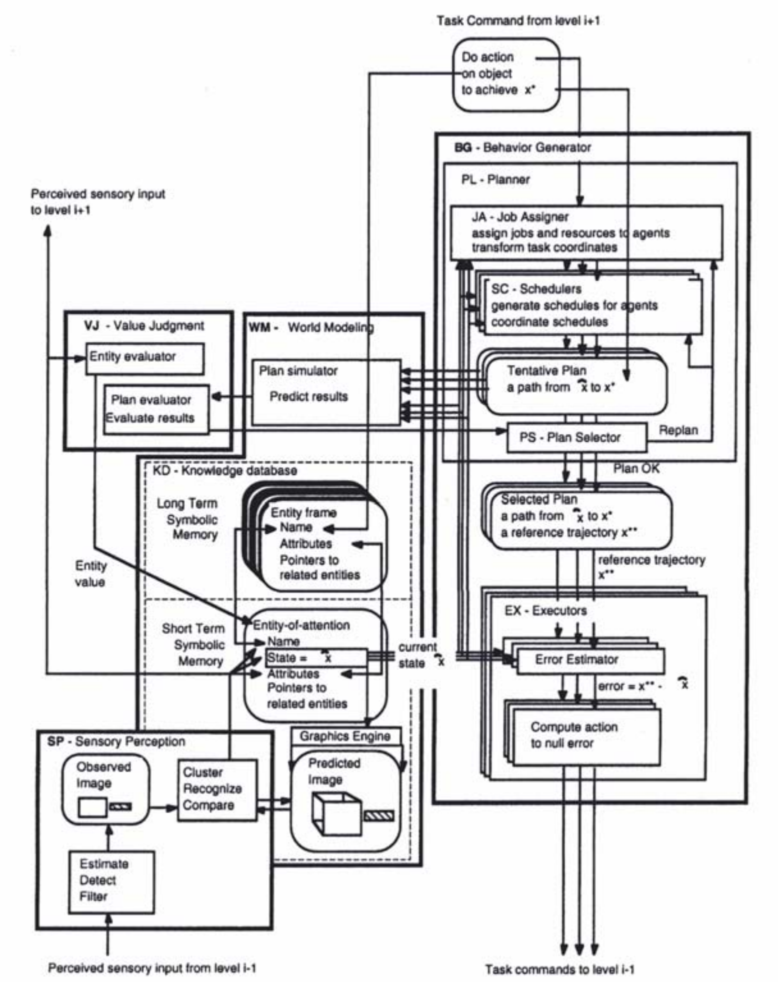

Figure 211 – NASREM architecture: relationship between components within each module

The knowledge base loosely represents stable long term memory (LTM) and they are limited to specific problem domains. Real world data in a dynamic global working memory database (loosely correlating to short term memory (STM)) must be matched against knowledge stored in the knowledge base. Events comprise the inputs to the working memory which is acted upon by an inference engine. This requires the expert system to search the knowledge base for knowledge that matches the data using heuristic control procedures encoded as rules of "good judgement" to avoid the combinatorial explosion by narrowing the search. The inference engine is independent of the knowledge base and emulates deductive behaviour through logical procedures for searching and deriving conclusions from the sets of facts in working memory. The inference engine interprets the current state of working memory and applies the relevant logic procedures which specify the changes to be made. It uses predicate logic to construct deductive proofs to make inferences. Hence, working memory functions as the world model containing all the task relevant data representing the state of the world containing operating knowledge (from the knowledge base) and environmental knowledge (from the sensors). A production system consists of a knowledge base containing all the task relevant information in the form of a set of production rules. Production rules represent relatively independent chunks of knowledge and are used by the inference engine to infer solutions to problems. Production rules define relationships between the rule's terms through implication. These rules consist of one or more antecedent precondition statements linked to one or more consequent action statements (condition-action rules). They essentially model plausible hypothesis-conclusion inferencing logic for problem solving through deduction. Production rules are a subset of predicate logic with the addition of prescriptive components to indicate how the rules are used during reasoning. They can be used to model routines with particular input-output characteristics compactly as stimulus-response behaviours. They incorporate practical conditional if-then knowledge and can solve reasonably complex problems by combining the rules into sequences and indicate explicit lines of reasoning. The precondition statements determine the applicability of the rules in terms of the conditions that must be true for implementation of the rule. The action statements are the operators which determine the effects of the action if the rule is applied to the world model.

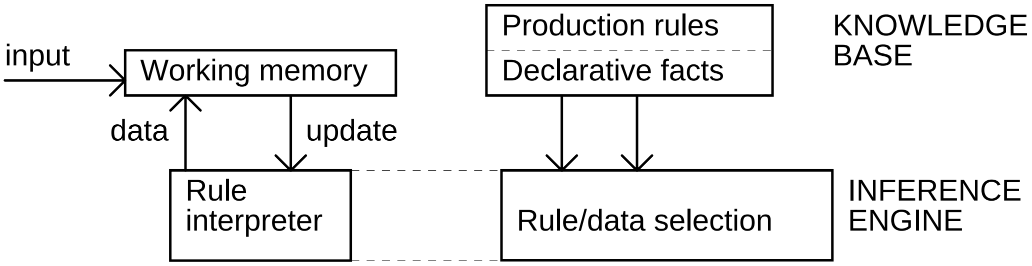

Figure 212 – Knowledge-based expert system

If more than one rule is satisfied and contradiction occurs, conflict resolution priority strategies are employed: specificity - use the most specific rule before the most general, i.e.. the rule with the most constraining conditions; strength – use the rule with the greatest certainty factor; recency - use the most recently fired rule; goal-orientedness – use the rule that satisfies the highest priority goal. Production rules often have numerical weights assigned to them to quantify the level of belief in the consequent given that the antecedent is true. These weights (certainty factors) comprise numerical confidence ratings which quantify plausibility and are the basis of heuristic control to guide the program search for plausible inferences. The certainty factor is a measure of the association between premises and actions to model uncertainty. The basis of determining uncertainty is probability theory either through Bayesian statistics (usually) or Dempster-Shafer theory of plausible inference (rarely as it is complex to compute) to give a truth value as a probability on the real unit interval [0,1]. Alternatively, fuzzy logic may be used based on possibility theory [Tong 1977; Maiers & Sherif 1985]. Most often, degree of certainty is quantified in terms of probability via Bayes theorem which determines the likelihood of an occurrence, eg PROSPECTOR, a mineral exploration expert system [Duda, Hart & Nilsson 1978].

Decision theory uses probabilities to generate utility values which encode risk preferences to accommodate conflicting goals such that the maximum utility is selected. Preferences are represented as real-valued linear maximised function over a probability distribution:

$$\text{Expected utility of i, } \langle U\rangle(Pi) = \sum P(Oj|Pi)*U(Oj)$$

where:

P(Oj|Pi) = probability of jth result after executing ith plan.
U(Oj) = utility of jth outcome
$\sum P(Oj) = 1$

P(E)>0 for any event E, and P(E)=1 if E is certain to occur. P(E or F)=P(E)+P(F) if E and F are mutually exclusive.

Pattern matching that is used by the rule interpreter is a basic operation in symbolic processing and around 90% of execution time in production-based expert systems is spent in this phase. It proceeds through unification which is a form of bidirectional pattern matching used in logic programming and it determines whether two terms can be made identical by finding a set of substitutions for variables in terms. The chief difficulty of production rules is that they are inflexible to iteration and recursion. They are incapable learning as they lack the means to perform induction to abstract properties from examples. They are difficult to modify or extend while maintaining consistency (truth maintenance). As with all domain dependent expert systems, they are brittle and limited to their specialised universe of discourse and tend to fail dramatically in problems outside their area of expertise and in unforeseen circumstances.

Knowledge-based systems use hierarchical network organisations of conceptual knowledge according to resolution whereby general concepts are higher in the hierarchy and more specific concepts are lower in the hierarchy. However, these hierarchical networks are not static as each concept may change both its internal structure and its relationship with other concepts similar to a distributed message-passing architecture of modular experts. Modularity between units (both connectionist and symbolic) is fundamental in combating complexity and hierarchical structure provides the means for relational determination. Subgoaling is a form of hierarchisation of problems where subgoals act as building blocks for solving more complex problems.

Often, there are a number of possible courses of action which can achieve goals, particularly if the goal solution is open-ended. There is then, a choice of actions [Sage 1981]. The expected utility model is based on the rational actor model of economic theory. It is assumed that a set of well-defined goals and objectives have been defined with preferences indicating the degree of satisfaction of these goals. Alternative courses of action are identified with their respective consequences and the choice between alternatives is based on maximising the degree of satisfaction of those goals quantified by the maximum utility value. Alternatives are weighted according to objective utility values which encode risk preferences to accommodate conflicting goals. There are certain elements in the rational actor model:

- a set of alternative actions, $a \in A$;
- a set of possible consequences of those alternatives, $s \in S$;
- a utility function U(s) for each member of S;
- a probability of outcome $s \in S$ if alternative $a \in A$ is chosen, $p_a(s)$.

The expected utility of alternative a:

$$< U(a_i) >= \sum_{j=1}^{n} p[s_j(a_i)] \cup [s_j(a_i)]$$

where:

$s_j(a_i)$ = *states resulting from* a_i.
$P[s_j(a_i)]$ = *probabilities of occurrence of those states given the situation* a_i.

Utility is comprised of a number of components (attributes) and their relative weightings. A linear utility model is given by:

$$U(s_j) = \sum_{j=1}^{m} w_j u_j(s_i)$$

where:

$$\sum_{j=1}^{m} w_j = 1$$

w_j = *weight*
$u_j(s_i)$ = *value*

The assumption is that there are objective probabilities $p[s_j(a_i)] = p_j(a_i) = p_j > 0$ for any event a_i such that $\sum_{j=1}^{n} p_j = 1$. The subjective expected utility model replaces p_j for $f(p_j)$ such that $\sum_{j=1}^{n} f(p_j) = 1$.

Savage's independence axiom states that when dividing between two alternative gambles with a common outcome, the value of this outcome does not offset the choice because the choice is independent of the outcome's utility value. Consider the choice between two courses of action whereby the outcome of each action is influenced by one of two possible intervening events. The action chosen will depend on which intervening event occurs yielding two different outcomes. The action chosen may result in a less desirable outcome if one intervening event occurs rather than the other. Von Neumann-Morgenstern relative utility is used to rate the degrees of desired outcome (preferences) which depend on the relative probabilities between the two intervening events. When ranked in order of preference, each action is assigned an expected utility value (sum of probability weighted utilities of possible outcome). The rational choice dictates that the highest expected utility action is selected. Gain is preferred to status quo which in turen is preferred to loss. Such preferences are transitive, i.e. A>B and B>C implies that A>C [North 1968]. Transitivity ensures the existence of a maximum. However, humans violate the independence axiom such that linearity in probabilities no longer applies.

3.13.5. Frames & Semantic Networks

There are two main forms of knowledge representation commonly used in expert systems, frames and production rules, of which the latter are the most popular [Nau 1983]. Frames are schematic, modular data structures for representing generic concepts concerning objects, classes of objects and situations. Frames provide a natural representation for declarative knowledge (knowledge about something) while production rules are usually used for procedural knowledge (knowledge about how to do something) [Fikes & Kehlar 1985]. Frames and production rules may be integrated into a hybrid representation scheme so that frames can represent procedural action oriented knowledge. Production rules can be attached to frame slots as values to be invoked as demons when the frame is invoked. This ability to attach rules or rule classes as demons to frame slots provides great flexibility in controlling reasoning particularly for reactive functions.

Conceptual dependency structures can be built with groups of production rules appearing in different concepts.

Part of the problem is to ensure that the symbolic representation scheme and their manipulation accurately reflects and corresponds to the real world object and event counterparts – the symbol grounding problem [Harnad 1990]. Expert systems are of limited domain applicability though there have been attempts to broaden their utility into multiple "general knowledge" domains through the use of distributed techniques such as "blackboards", eg. CYC. Humans however possess episodic memory which are often modelled by scripts which encapsulate typical scenarios, eg. "restaurant" script (though it appears that the hippocampus stores both types of memory). Some work has been done on attempting to model emotion and motivation by implementing value systems.

Many workers have proposed that human thinking might depend on knowledge structures comprising a core of knowledge (schemas) which provide context in structural form. This provides a model for episodic memory of remembered events. Most situations we encounter are classified into a small number of classes, each of which may be represented as schemas. Such schemas are often referred to as frames (not to be confused with the frame problem) or scripts in that they represent modular chunks of knowledge, particularly declarative knowledge through the classification and categorisation of events and objects. Frames (or schemas or scripts) are modular template data structures for representing declarative generic concepts concerning objects, classes of objects, situations, concepts, event sequences and action sequences. All relevant information is collected together in modular form. The frame essentially represents a stereotypical model of features characterising a concept or situation which model situations that occur in the external world. Each schema has a number of conditions, a sequence of scenes and a set of results. They are packets of information comprising of a fixed part that is always true of the instances of the concept and a variable part (eg. people's roles or props). They provide background knowledge to typical situations. The variable parts if not specified explicitly are assigned default values, i.e.. they permit the representation of default knowledge in the absence of complete information. The learning process involves the construction of such frames by continual modification and refinement of its structure as they are recalled and adapted to cope with new variant conditions. Each frame has sets of attribute descriptions held in predetermined fields (slots) that describe the situation including function, form and criteria for its applicability. Each slot of the frame corresponds to a variable in the frame, eg:

```
Name of frame:……………
Type of frame:…………….
Source:…………………….
Destination:………………
Agent:……………………..
Object:…………………….
```

The list of attributes and rules describe the typical features of the frame's concept. They are unique and constitute the necessary and sufficient properties to characterise the concept. Attributes may be procedural attachments such as production rules or frames can be nested to form collections of frames organised into a network by the values of the attributes of the frames referring to other frames and linking them together. The frame can include both declarative and procedural information in specified relations. Specific instances of the situation are created by copies of the frame and assigning specific values to the attributes. A frame representing a class of objects can contain prototypical descriptions of the members of the class as well as of the class as a whole. Two kinds of slots are possible: "own" slots which can occur in any frame to describe attributes of the frame, and "member" slots which can occur in frames representing the class to describe the attributes of each specific member of the class. Hence slots can be used to describe object attributes. Slots may also have multiple values. Each slot is restricted to a particular data type specification to preserve the semantic integrity of the frame. A member frame inherits the slots of its superclass frames to which it has membership links. Each slot is characterised by a set of facets. The value facet contains the current value of the slot. Additional facets may restrict the type of value or other characteristics in adding metalevel information to the frame. A slot may contain a pointer or symbolic tag to another primitive object enabling frames to refer to each other as linked lists. Hence, frames may be factored into taxonomies of concepts whereby each class may be regarded as specialisations of more generic classes providing definitions through specialisation. Frames may be used to index memory to evoke expectations. Scripts, a variation on frames, encode knowledge about stereotypical everyday situations. Scripts are schemas for frequently occurring sequences of events such as social events. The oft-quoted example is the “restaurant” script. Each script can be subdivided into prototypical scenes. They store background information which is implicit in the frame, eg. a restaurant scenario and such frames form the basis of socially shared knowledge which reduces the requirement for information exchange through language. Scripts can be used to guide hypotheses as they focus on the context of behaviour. Hence these sorts of constructs may provide the basis and the structure required for understanding unstated goals in communication between individuals. Not all situations comprise of high frequency events and more general abstract schemas may be required. These comprise plans which are formulated to satisfy specific motivations and goals. This implies problem solving capability such that larger plans are constructed from smaller more fundamental data elements at a lower level of abstraction. Hence frames may be organised hierarchically and dynamically. They may be embedded to represent many hierarchical levels of abstraction as schemas, subschemas and superschemas. Frames do suffer from certain disadvantages in that typicality and plausibility are vague concepts and deviations from stereotype are difficult to handle. However they provide a means for object oriented programming where each frame represents an individual uniform object encapsulating all data and procedures. Each object frame may interact and communicate with other object frames by sending and receiving messages. The objects may be manipulated as data or describe procedures or both. Furthermore they may be structured as semantic networks with each node representing a single frame. Hence the semantic network/frame approach may be regarded as a software distributed architecture with the communicating links acting as messages passed between the object frames. The script is a formal construct of background information, eg. the restaurant script:

```
PLAYERS:  customer, waiter, cashier
PROPS:    restaurant, table, menu. food, bill, payment, tip
ACTIONS: (i)    customer goes to restaurant
         (ii)   customer goes to table
         (iii)  waiter brings menu
         (iv)   customer orders food
         (v)    waiter brings food
         (vi)   customer eats food
         (vii)  waiter brings bill
         (viii) customer leaves tip for waiter
         (ix)   customer gives payment to cashier
         (x)    customer leaves restaurant
```

The script provides a mechanism for using and assuming knowledge about the world which may not have been explicitly stated. They are particularly essential in written communication such as narrative where there is no other available information to resolve ambiguities. The difficulty lies in producing the scripts which are sufficiently complex and complete. Scripts are inherently linear and sequential. The everyday world may be reduced to ~100 generic situations (schema) with ~30 variations of each which may be used as the world model. Understanding (semantics) provides the basis within which to generate interpretations in a recursive process, ie. partial understanding is used to further understanding. This pre-understanding may be provided by a set of stored scripts or schemas with new schemas being generated through the interaction of other schemas. Such schemas provide a formal representation of patterns of regularity in nature. Semantics which effectively places sensory inputs from the external environment into internal context is part of the sensorimotor process with the purpose of generating bodily behaviour interactions with the external world. Perception of referents is a cognitive process and may be modelled using artificial intelligence techniques. Meaning is then a result of perceiving referents in a context of the environment and meaning represents a psychological context of experience.

Semantic networks provide the means for organising such knowledge, concepts and/or words. Semantic networks are one form of knowledge representation which comprise directed graphs whose nodes represent semantic entities and whose links represent relations between those entities. Semantic networks are directed labelled graphs of nodes are connected by binary relations with certain strength values measuring the relevance relations between the nodes. Each object is associated with other objects interconnected into a tree structure. They are similar in architecture to relational databases and form a dependency network of concepts. They define a relational logic through network links. Semantics is viewed as a contextual property of the concepts in the semantic network. Each node represents a concept and the concepts are represented as semantic attributes and stand in terms of their specific relationship to other concepts as represented by links. Relevance of the links is an equivalence relation less transitivity. The link structure provides a means for organising knowledge. The structure of the nodes and links provides a means for organising knowledge. The properties of the concepts may be represented as binormal associative or relational links between the nodes forming a highly interconnected structure. The pattern of the relationships between concepts represents the meaning of those concepts. The relations act as pointers to the node data structures (frames). The semantic memory so formed may be organised primarily in terms of noun categories (i.e.. a dictionary) with the relations being verbal or adjectival. Hence concept names are stored as a lexical network with the form of a tangled hierarchy. The semantic network models semantic relations rather than syntactic information. The more properties two concepts have in common the more links that exist between the concept nodes. This determines the semantic relatedness as the aggregate of all common paths. The links between concepts collectively form the definition of the concept in that it is defined holistically by its association with other concepts. Concepts are thus defined as the aggregate of its relational links. The total number of links from all other concepts connected to a concept defines its fan. The reaction time for memory retrieval of a concept increases in proportion to the fan of the concept. Links may also have values indicating how essential each one is to the meaning of the concept. There are in general 5 different types of links:

1. set membership links;
2. modifier links;
3. disjunctive links;
4. conjunctive links;
5. verb concept relationship link.

Information may be clustered into relevance classes equivalent to pictorial segmentation. Semantic nets are often used for defining taxonomic hierarchies of classes. This effectively performs the repeated application of modus ponens inferencing through the network. Hence semantic networks are formally equivalent to a subset of first order predicate logic. The semantic network provides a basis for classification into categories which provide the basis for semantic meaning. Each class may be regarded as a specialisation or subclass of other more generic classes. There are two main types of link between concept nodes adopted in such networks but both are "inheritance" links (superc) to provide a categorisation and classification capability: "is-a" and "is-a-kind-of" link. These two links provide the semantic inheritance network with the representation of both generic-generic relations of subtypes and generic-individual relations. The generic-generic relation represents the subset relation between generalisation/specialisation (universally quantified conditional) while the generic-individual relation represents the set membership between type/instantiation (predication), i.e.. it can accommodate predication in that an individual is of a certain type and universally quantified conditionals in that some types can be subtypes of another type. Higher nodes act as parents to lower nodes and each child (subset or instance of the parent concept) automatically inherits all the attributes of the parent class but with the provision to override standard attributes with unique attributes, i.e.. the frames can be organised into generalisation/specialisation hierarchies. Properties common to more than one concept are expressed at the most general level and inherited by all the more specific nodes. A taxonomic hierarchy of differing resolutions of types gives a partial ordering and strongly resembles the biological classification taxonomy, eg. Animal-Vertebrate-Mammal-Primate-Homo-Homo sapiens. Hence the semantic net makes apparent the interactions between concepts. Semantic network links can be defined in terms of frame slots within each frame to support inheritance. This definition by specialisation with inheritance provides an efficient and concise memory storage scheme in that properties are shared in a distributed manner through the superc link across multiple nodes maintaining semantic integrity. This enables updating of the network and the maintenance of consistency [Fikes & Kehlar 1985; Woods 1986]. One desirable capability is a "programmable" link such that attributes can be turned off and disinherited by more specific nodes, i.e.. to represent exceptions over defaults. This capability would allow the expression of the difference between logic and default logic (universal or default truth) and modal logic (necessary or contingent truth) [Brachman 1983]. Semantic networks are formally equivalent to a subset of first order predicate logic - nodes may represent single argument predicates and links may represent binary argument predicates, eg. in KL-1. N-ary predicates may be generated using conjunctives. Such semantic nets can provide the means for organising scripts. Knowledge is modelled as a network in which concepts are activated automatically and each activation triggers its neighbours causing the activation to spread. Concepts that are appropriate to the input are activated more vigorously while irrelevant concepts fade. Semantic nets are a suitable representation system as lookup tables of production rules in graphical form. Both procedural knowledge and declarative knowledge should be included. Declarative knowledge comprises the thesaurus of world information. The descriptive world is a global semantic network comprising the universe of discourse. It represents all contextual knowledge as lists of relations between labels. This list should allow new relationships to be added through learning. Generalisations organise all labels into classes with relations comprising semantic interpretations, i.e.. associative memory. Such a memory may exhibit varying degrees of resolution through the resolution of membership classes.

Knowledge retrieval can involve searching the whole network but great speedup can be obtained by simultaneously initiating multiple searches from many nodes. Memory search is a parallel activation state spreading as a breadth first search along the links from two or more concepts stored as nodes in the semantic network. This activation state constantly expands as spheres of activation until an intersection occurs between the two nodes. The activation represents the amount of information processing in the semantic network structure. The intersection indicates a potential relationship between the two concepts by effectively constructing a path between the two nodes. This is the spreading activation inferencing mechanism which spreads through marker passing and was originally proposed as a theory of human semantic processing through human long-term memory [Collins & Loftus 1976]. The sources of activiation are the outputs from the perceptual system and goal structures. Consciousness may be characterised by the temporary network of nodes activated. Activation spreads out along the network in a decreasing gradient that is inversely proportional to the strength of the link, i.e.. the signal is attenuated. Furthermore the activation decreases over time. The spread of activation determines the latency of recognition. The fan effect is another property of the network. This means that the probability that the activation will spread through a link is a function of the ratio of the strength of the particular link to the sum of strengths of all the links emanating from the node, i.e.. activation is inversely proportional to the number of links that fan out from the node. The more paths from a concept the less the activation spread [Anderson 1983]. More frequently accessed facts about a concept increase the strength of the link enabling its activation to be more rapid. Hence recognition time $t = 0.36 + 0.77(P - 0.5)^{-0.36}$ evolves through learning to $t = 0.36 + 1.15(P - 0.5)^{-0.36}$ where P=learning parameter. There is a power law decay process that weakens the strength of connections over time if i is not accessed: $s = \sum_{i=1}^{n} t_i^{-d}$. This builds a strong recency component into the network in conjunction with the frequency effect. Priming involves using the same trace as for the memory search so that the activation is prepared and spread to an unspecified depth. Hence when another concept is subsequently presented the memory search expands to intersect with the activation of the primed concept. This priming process reduces the reaction time for the association of concepts thereby increasing the speed of memory retrieval.

Frames have no direct facility for declaratively describing how knowledge stored in frames are to be used procedurally. Such procedural knowledge is usually cast in the form of production rules which by themselves are inadequate for describing domain objects and their relationships. Prototypes in frames containing necessary but not sufficient conditions can be used by a classifier to conclude that an item is not a member of a class but cannot be used to conclude that the item is a member. Production rules provide the means for augmenting class descriptions to include conditions for determining sufficiency of membership in a class. In fact default descriptions do not even specify necessary conditions. Production rules can be used for successive refinement classification. The rule condition is that an item be a member of the class provided it satisfies the local conditions for membership. Hence production rules provide a guiding structure for organising rules to specify sufficient conditions for class membership. Simon (1991) suggested that the human mind is dominated by semantic networks of frames which provide both inferencing capability and declarative knowledge and that humans learn only declaratively initially and procedural knowledge is obtained from such declarative knowledge.

General-purpose expert systems available include KEE (Knowledge Engineering Environment) and ART (Automated Reasoning Tool) [Fikes & Kehlar 1985]. These provide a means for experimenting with facts and situations to determine the viability of alternative interpretations or effects (alternative worlds) without disturbing the current system (present world). This allows the extension of rule sets by merging of these worlds. A declarative representation comprises a set of domain specific facts in the form of first order predicate logic. Facts are declarative and express assertions about properties, relations and propositions which are usually static since they have absolute truth values and may be stored in a knowledge base separately from the procedural rules of inference. The frames provide the foundational structure for describing objects and their relationships and taxonomies. They also provide a means for organising into chunks and controlling large collections of production rules. Frames provide several ways to attach production rules to frames to allow the retention of the object oriented approach [Woods 1986]. Production rules can be attached to frame slots as values to be invoked as demons as pieces of "advice" when the frame is invoked. This provides the hybrid representation with two types of inferencing: frame-based inheritance inferencing and production rule inferencing which is effectively "wired in" to the overall frame representation. Such semantic networks provide an alternative encoding to predicate calculus and go beyond it. An alternative is to use frames to represent the production rules themselves. Rules can easily be grouped into classes. Additional descriptive information may be incorporated as frame slots. Rules may then be invoked as a group. This provides a means for overcoming some of the shortcomings of large rule bases of ~1000 rules so that they may be organised into small modules to prevent bugs and uncontrolled interactions. This is essentially a rule management procedure. Large rule bases are best structured hierarchically by aggregating rules into rule sets or frames according to their function. Both KEE and ART support frame representations. KL-1 also offers a hybrid production rule/frame representation [Woods 1983]. The situation parts of all production rules are organised into a structured taxonomy of all the situations and objects in the knowledge base. A collection of concepts are linked together by "is a" links allowing the inheritance of properties between concepts. The action parts of the rules may be attached to the nodes in the structure as pieces of "advice" that apply to the situations described by the rules. Each concept is a configuration of structural attributes which satisfy certain restrictions and have specific relationships to each other. More general concepts subsume more specific subconcepts such that some of a concept's attributes are directly attached to it while others are inherited from more general nodes. For instance, a frame of production rules may represent the concept of ACTIVITY with slots for TIME, PLACE and PARTICIPANTS. A more specialised concept is the PURPOSIVE ACTIVITY which inherits the properties of the more general ACTIVITY concept but differentiates PARTICIPANTS into AGENT and OTHER PARTICIPANTS and introduces a new slot GOAL to represent the purpose of the activity. If conflicts arise the advice associated with the more specific rule can be executed supplementing the more general rule. Semantic nets provide a means for analogical reasoning through graph matching as a means of defining semantic closeness. New cognitive structures can be generated by extending old ones. Many new concepts are learned through analogy to past experience, eg. the principles of electrical conduction are usually learned by analogy to fluid pipes. The semantic net representation facilitates the learning process by analogy through graph matching.

Fahlman & Hinton (1987) described NETL, a coarse coded distributed neural network/semantic network whereby each processing element is tuned to each concept or feature. Silverman et al (1989) suggested that a neural network based associative memory may be a useful representation mode for a semantic network for archival storage and rapid search and retrieval. The nodes of such a semantic net may represent propositions and links may be quantified by connection weights (representing degree of association). The evidence for a feature was represented as the unit's activation level. Such coarse coding reduces the number of units required to represent a range of values by generating fine precision from the simultaneous activation of multiple units which are coarsely valued. The motivation was to attempt to combine the advantages of connectionist neural networks and logical inferencing without their respective disadvantages. Artificial neural networks are suited to multiple constraint satisfaction reminiscent of subcognitive unconscious processing while logical inferencing offers transparent semantics. NETL was a hardware implementation of the semantic network where the nodes represented symbolic noun-like concepts and links represented the relationships. The specific inferencing link adopted was the inheritance "is a" link handled by simultaneous marker propagation through the network and its parallel set intersection capability enabled recognition tasks to be implemented. NETL was chiefly designed for implementation on the 64,000 processor Connection machine. The connection machine is a semantic net processor network which uses memory based reasoning. Semantic networks can be built up with pointers that couple rule elements into a network such that activation of a rule causes the activation of rules to which it points, i.e.. C2 is coupled to C1 if a condition of C2 is satisfied by the result of the generation of the action part of C1. Any logical combination of conditions can be realised by coupling primitives by AND, OR, NOT combinations to form association weights of a coarse (rule) coded neural network. Tags can create hierarchical or heterarchical associations between rules. Concepts may thus be represented as groups of simultaneously activated rules as a spreading activation. Alternatively, NETTalk translates English text into speech using MLP backpropagation algorithm implemented on a 3 layer network. Each node represents a letter of the alphabet. Inputs are binary ON/OFF nodes. After the presentation of 10,000 patterns, the connection weights converge. Letters are symbols encoded at the input layer and phonemic outputs are also symbolic categories. The basic symbolic categories are fixed defined by the designer. The NETtalk internal states are identified by cluster analysis to extract the phonemes. The chief criticism of connectionist architectures is that the interconnection architecture is fixed (although the weights are modifiable). In semantic networks however, the interconnections are dynamic and the interconnection architecture lends concept nodes their semantic content and such semantic entities may change. If ANN nodes are similarly assigned symbolic concepts, their semantic content cannot be altered.

Sophisticated speech recognition and natural language understanding requires a great deal of contextual knowledge. To that end, Lenat et al (1990) have been developing CYC for a decade, a massive knowledge base of ~10^8 axioms spanning all human common sense knowledge equivalent to a one volume, 30,000 article encyclopaedia (or 1/10 of the Encyclopaedia Britannica) - around 300 programmer-years. Most expert system knowledge bases which cover narrow domain specific applications have only ~10^2-10^3 rules. They are extremely brittle such that when confronted with unanticipated situations they deliver inappropriate responses. ASR requires a large knowledge base with general facts and heuristics of human consensus knowledge including cultural knowledge, beliefs and knowledge of other agents. There are both general rules and specific facts about objects and situations. CYC comprises a knowledge base in a frame-based representation embedded in a predicate calculus framework acted upon by a number of inference engines. The frames (schemas) represent schemas which classify and categorise events and objects through inheritance properties. The inference engine implements forward and backward chaining. Over 90% of the assertions are non-monotonic representing default beliefs. Most of the monotonic assertions are definitions or mathematical facts. The non-monotonicity allows three truth states: default true/default false/unknown. It also implements causal relations in its logic representation. Truth maintenance ensures consistency while a machine learning subsystem searches for unexpected symmetries while "sleeping" (bugs usually). The designers view CYC as the formulation of the first true artificial intelligent agent by virtue its general knowledge coverage.

3.13.6. Physical Symbol Processing

The top-down approach based on symbolic processing is based on two hypotheses:

1. physical symbol hypothesis (principle of rationality) whereby representations are physical structures (symbols) which have correlations with aspects of the environment maintained by sensory data – these symbols are defined and processed according to a formal logical system.
2. knowledge level hypothesis that these symbolic representations encapsulate knowledge of that environment – their manipulation provides predictive power by simulating relevant processes of that environment assuming that there is maintenance of strict correspondence with the environment as it changes over time.

The symbolic method is particularly appropriate for implementation on computers and has been very effective in modelling many aspects of higher level reasoning [Wah et al 1989].

The corollary of the Physical Symbol Hypothesis is that both human beings and computers are physical symbolic systems that generate an evolving collection of symbolic structures over time. This is a statement of the strong AI school which suggests that the cognitive states of a machine are functionally equivalent to the mental states of a human mind (the weak school of AI suggests that computers can simulate human cognitive processes but that no correlation exists between the cognitive states of a machine and the human). A formal symbol system comprises of semantically interpretable symbols and rules for manipulating those symbols. These symbols are arbitrary in relation to the objects and events that they represent. It makes the explicit assumption that cognition is independent of the physical substrate on which it is implemented, i.e.. “subcognitive” neuronal dynamics with its characteristic timescales of less than 100 ms has no bearing on cognitive processes. Formal symbolic states are independent of their physical implementation. The physical symbol hypothesis suggests that cognition is the result of symbol manipulation according to rules. Symbols represent the states of physical objects. Their formal transformation and manipulation may therefore be carried out by a Turing machine. The Church-Turing thesis states that computation can represent any formal or physical process. Mental states are therefore implementations of symbolic states. Symbols are labels for representing abstract forms which may be acted on by algorithms and so is appropriate for computer implementation. The hypothesis assumes that internal representations of the world mediate between perception and behaviour, i.e.. a computational model for information processing via well-defined algorithms. The symbol's role is determined by formal logic with well-defined grammars with semantic features arising as a result of an active constructive interpretation process with the symbolic representation. Thus, knowledge arises as a result of structuring environmental stimuli. Such semantic interpretation requires an external interpreter. The physical symbol hypothesis suggests that cognition can be divorced from perception and action processes (brain-in-a-box approach). Piaget's theory of intellectual development supports the Physical Symbol Hypothesis by viewing the development of intelligence as the acquisition of schemas (mental structures for the accomplishment of tasks) through interaction with the environment over a series of consecutive ontogenetic stages, language being one manifestation of a generalised symbolisation capability. After a few months, babies begin babbling, initially with just vowel-like noises and later with consonant-type sounds in addition. Language acquisition begins to appear in the second year of life. The sensorimotor period from birth to 2 years of age involves the learning of sensorimotor coordination and organised behavioural patterns through trial and error exploratory behaviours. The pre-operational period from age 2 to 7 years involves the development of internal representative symbolisation (including language use). The concrete operational stage from 7 to 12 years involves learning the ability to apply logical mental operations to concrete objects in the context of structured relations (including social skills). Concrete operational thinking involves classification and generalisation of empirical data, manipulation of symbols in direct cause-and-effect relationships and the organisation of facts to solve well-structured, familiar problems by analogy to personal experiences. The final stage of intellectual development is the formal operational stage which occurs after the age of 12 when the child gains the ability to apply mental operations to hypothetical abstract objects (symbols) to generate hypothesis-deduction concepts. Formal operational thinking adds the ability to reason formally and hypothetically through deduction about speculative problems beyond personal experience. Much of learning involves the transfer of formal operational knowledge into concrete operational knowledge [Sage 1981]. Experienced and expert human problem solvers solve new problems by exploiting similarities to previously encountered situations. This provides more efficient performance over formal approaches which utilise abstractions. There are 5 stages to performance:

1. novice involves explicit rule-following to determine actions;
2. competence involves examples used as guidelines to determine actions;
3. proficiency involves memorised principles used as maxims to determine action;
4. expertise involves specific situations triggering intuitive actions;
5. mastery involves automated reaction without the need for attention.

The Leibnitz-Whorf hypothesis also lends credence to the Physical Symbol Hypothesis as it regards human thought and reasoning as being entirely determined by the formal and logical structure of language. The rules of such symbol manipulation are determined by the rules of logic. The interpretation of symbols is not inherent in the symbol itself. The symbols and the logical relations between them must carry their own interpretation or such interpretation must be an emergent property of those rule sets, eg. semantic networks. Human thinking is based on universal symbol representations of concepts learned from particular concrete examples, i.e.. an inductive process of going from the particular to the universal which is central to hypothesis generation. Fundamental to this process is pattern classification. This is fundamentally an experiential process. The only way to learn the meaning of symbols in language is through experience of interacting with the object(s) in the environment. Semantics is thus a pragmatic process. Furthermore, language for communication requires the overlap of private meaning of symbols with the public meaning of symbols. This consensus of domain knowledge is enabled through interacting in a common environment, i.e.. it is overlapping of experiences that determine the degree of understanding. Machine intelligence is at present limited to abstract logical representations of symbols which are not grounded in the real world [Watanabe 1974]. Kirsch (1991a,b) declared that knowledge only has meaning in relation to the agent's environment, i.e.. the environment provides the semantic grounding of symbol systems. Although it has been called into question whether human intelligence is optimal, it is the best model we have. The interdependencies of learning, reasoning and knowledge essentially underlie the difficulties that AI has had concerning mode of representation. It is highly doubtful if any single representational framework will be sufficient for intelligent activity [Smith 1991]. After all, humans employ a variety of representations ranging from symbolisation to imagery. Human information processing is not only logical, abstract and analytical – it is also holistic, concrete and based on past experience. Furthermore, the world is non-deterministic and uncertain and organisms must respond flexibly to this world. The problem of control in the real world is ill-posed in that the conditions of stability, uniqueness and solubility are generally not satisfied.

The chief problem is that the maintenance of strict correspondence between the environment and the representation cannot be achieved by sensors – sensors are not full environment detectors, merely determine weak correlations between external and internal states, eg. object locations [Steels 1995]. A host of supplementary techniques have been developed to augment this logical approach has been developed – circumscription, non-monotonic logics, default logics, fuzzy logics, temporal logics, etc - none of which are entirely satisfactory, and indeed, deprive logic of its provability. Furthermore, these methods are not amenable to learning (except through a limited form of "chunking"), eg. SOAR. A variation on the SOAR concept is case-based reasoning which attempts to emulate aspects of analogical reasoning. There have been models which attempt to model the architecture of human cognition, eg. ACT*. Production systems provide a mechanism for problem solving and all forms of cognitive processing. Production rules consist of independent condition-action pairs embedded into a large knowledge base. One general purpose computational architecture is ACT*. Working memory interfaces via the perception and motor systems to the external environment. Long term memory is comprised of declarative and production memory. Declarative memory comprises interconnected hierarchies of factual data stored as temporal knowledge, spatial images and semantic propositions. Production memory comprises of production rules of associated strengths encoding procedural knowledge pertaining to abilities, skills and techniques. Declarative memory functions by spreading activation: the most active nodes comprise the working memory according to their activation strengths. This is the spreading activation inferencing mechanism which spreads through marker passing and was originally proposed as a theory of human semantic processing through human long-term memory [Collins & Loftus 1976]. The sources of activation are the outputs from the perceptual system and goal structures. Declarative memory may be updated easily, but procedural memory may be updated only after much practice. The transition from the declarative knowledge to the procedural knowledge bases constitutes automisation of skills. Working memory comprises an active scratchpad from which information is directly accessed. It also compares and matches its contents to the condition parts of production memory rules. It performs this in a parallel fashion. When a production rule is activated its action part is executed. This pattern matching process is the major computational bottleneck. Goal structures are hierarchically organised into different levels as productions of plans in the production memory. Attention is provided by the degree of activation of goals – activation is switched off by appropriate production rules on completion of goals. This goal-directed processing is serial and aids in conflict resolution through the strength of activation, degree of match and specificity.

There is little doubt that natural language as a representation scheme is a superset of logic offering greater flexibility than syllogistic logic. However, the chief advantage of expert system type techniques is that they are transparent to analysis for algorithmic correctness.

3.13.7. Qualitative Physics as Models of the World

Much common sense knowledge is based on a naive qualitative physics which comprises knowledge of the scientific principles about the behaviour of the everyday real world beyond procedural and declarative limits. The laws of physics provide constraints on the behaviour of the real world, i.e. they constrain the nature of continuous change that the environment undergoes. Qualitative physics attempts to model the structure of the world by identifying the laws that describe the behaviour of physical processes without recourse to mathematical solutions of differential equations. It provides a means to model and simulate as an inference procedure to overcome the shallowness of expert systems which attempt to derive conclusions from observations (usually through a probabilistic certainty factor). However, conclusions and observations are independent, empirical entities.

A set of simultaneous differential equations encapsulating physical laws may be used to describe the behaviour of a physical system in terms of functional relations. Such equations are causal representations of the system behaviour as time must be represented as a totally ordered set of distinguishable time points. The duration of an event is the difference between the terms for its start and end points. All measurement sequences (observables) are totally ordered by the times of measurement. Events are adjacent if no interval exists between them. Any state comprises a finite number of components, some of which are fixed by the measurement sequence. Qualitative process theory defines two major physical processes – billiard dynamics and liquid flow, i.e. motion.

Laws are applied as constraints in order to extrapolate behaviour such as the laws of conservation of energy and quantity. Stability and equilibrium are also important concepts in qualitative physics. Qualitative physics based on processes and devices that cause change is an aspect of causality analysis. Causal models are deeper than expert systems as they model unobservable mechanisms, structure and behaviour thereby enabling explanations for sets of observations. Causality is an asymmetric property characterised by directionality and non-commutability. This requires the existence of asymmetric relations between the interdependent variables of the simultaneous equations (causal ordering). The equations form the basis of a model of distinct conceptual components (mechanisms) in terms of which the function of the complete system can be defined in terms of their interactions and connections. The constraint equations of the components constrain the values of the variables. The system behaviour arises from the interaction between components. These mechanisms are described by laws describing the physical processes. Each law is an equation with physical parameters describing the mechanisms, the variables of which influence other variables of other equations describing other mechanisms. Simulation produces sets of possible behaviours by generating possible directed state transitions given the qualitative constraints and initial conditions. Qualitative mathematical relationships include DERIV, ADD, MULT, MINUS, M+ (monotonically increasing), and M- (monotonically decreasing) [Forbus 1987]. Each parameter characterising the physical process is real-valued and varies continuously over time. It may be realised by a mapping function that is continuous and differentiable. Qualitative constraint equations are applied to the parameters to represent the state of the mechanism. A system of n equations with n unknowns allows analytic/numeric solutions to equations and this defines variable dependency in self-contained structures. If there are fewer equations than variables then some variables are undefined and assumptions may be added to generate additional equations to allow solutions and restore self-containment.

De Kleer & Brown (1986) and Kuipers (1986) regard the device or component as a composition of distinct parts each of which is charaterised by its particular behaviour. Causal ordering involves identifying the mechanisms and their asymmetric relationships in the self-contained dependency structures. Qualitative calculus involves describing how values change causally and in what direction (positive or negative) through sets of qualitative differential equations (called confluences), i.e.. whether an increase in variable x causes y to increase or decrease. Variable values ∂x are thus qualitative values +, 0, - providing an informal description of the physical situations [de Kleer & Brown 1986]. This overcomes the computationally demanding step of solving the differential equations mathematically or algebraically by transforming the equations into their qualitative counterparts and propagating the qualitative disturbance though the structure. Components interact by communicating +, 0 and – locally with their neighbours like a propagating wavefront through the system. Causal ordering is based on the equilibrium equations. Mythical causality involves the dynamic propagation of disturbances of the system from equilibrium values of variables through the network of constraint equations forming a causal dependency structure. Mythical causality describes the trajectory of the system through non-equilibrium states through the constraint network (which acts similarly to Gaussian elimination) until the variables assume new equilibrium values [Iwasaki & Simon 1986]. Mythical causality specifies the sign of the effects of the variables on each other – such signing is a necessary component of asymmetric relations, i.e. sgn(dx)=sgn(dy). Differentials of parameters are replaced by signs and the signs propagate through the system using the signed equation. The signs of the derivatives are determined from conditions for stability of equilibrium. Adjustment is not instantaneous. All regions of the system which are unaffected by the disturbance are invariant and propagation is restricted to the device's topological paths between its components. The no-function-in-structure principle states that component models are independent of their contexts. It assumes that all components are operating normally in that the system is in a valid operating region. Departure of observations from expectations may be utilised as symptoms of faults in a predictive manner. Feedback introduces a complicating factor such that sequences of cause-effect interactions produce effects on the antecedent in the sequence. When input and output are disturbed from equilibrium by changes, they will return to their equilibrium values at rates proportional to their deviation from them.

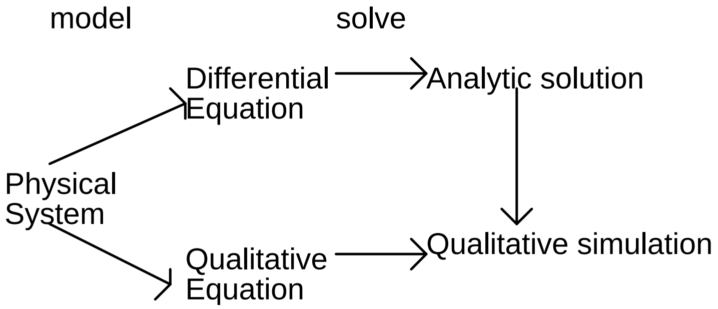

Figure 213 – Qualitative Simulation Process

Qualitative physics can detect instability conditions without the need for a dynamic model. Positive feedback is a condition that usually generates instability so negative feedback is automatically selected. The QSIM algorithm has been shown to be mathematically correct as well as computationally rapid [Kuipers 1986, 1987]. QSIM begins with a mechanism model comprising a qualitative abstraction of an ordinary differential equation to predict qualitative behaviours. It produces a graph of the possible future states of a system with possible behaviours being paths through the graph from the initial state (environment). It cannot be guaranteed not to produce physically unrealisable behaviours but this is a fundamental flaw in localism implied by the state-to-state transitions derived from immediate predecessors. The additive of global constraints such as conservation of energy can generate the correct behaviour and precludes physically unrealisable behaviour. Complex devices can be built from simpler components, the operations of which are described qualitatively. Qualitative simulation may be based on device topology (device structure) and envisioning (device function from structure) via a causal model. The causal model is based on conduits between components which act to give causal mechanisms for the transfer of information between devices. Mechanisms composed of component parts with well-defined behaviours connected by interaction paths forms a natural method of representing both distributed and hierarchical systems found in nature. Many processes were the result of many interacting equilibrium mechanisms. Hence, the temporal asymmetry of causation may be viewed as a symmetric constraint through abstraction to an equilibrium mechanism at a higher level. QSIM regards causation as a result of simulation of event sequences. Forbus (1987) described qualitative physics as being analogous to speech understanding models as one of its main functions is to quantise numerical data into qualitative descriptions. Processes and their effects are central such that all changes in physical systems are caused by processes. The environment requires selection of those behaviours that are consistent with measurement in a dependency-directed search over possible qualitative states (analogous to the imposition of grammatical rules in speech understanding). We can see that qualitative modelling through naïve physics is another consequence of the frame problem to cope with a changing world and the nature of causality.

3.13.8. Non-Monotonic Logics

Logic in AI imposes a variety of problems, almost all of which involve the frame problem in one aspect or another. Logic deals with eternal truths, and all extensions to logic are attempts to cope with causality and change in the real world. Indeed, this is the fundamental requirement in robotic planning. Indeed, it may not be an exaggeration to state that in a fundamental way, artificial intelligence is a sub-discipline of robotics as robotic planning systems use AI techniques which are subjected to the difficulties in coping with a changing world. Logic requires new axioms to be added to a list of provable theorems only when there are inconsistencies (Godel’s theorem). This is insufficient in a dynamic world where incomplete knowledge is the norm. Logic's great strength is in offering rigour to closed world systems (such as mathematics) but this is also its failing. Logic is descriptively universal and requires complete, certain and precise information. First order logic can handle only eternal truths. However, the world is dynamic and constantly undergoing change. Logic lacks the flexibility to cope with a changing world. First order logic appears limited in its ability to model human thought processes. Logic used to be considered the discipline that captured the laws of thought. However deep common sense knowledge is difficult to represent in terms of logic. Indeed only a fraction of human reasoning is based on deductive logic.

The frame problem is essentially one that requires a theory of causal relationships. Kelly (1989) viewed intelligence as the ability to solve the frame problem, i.e.. to select relevant information or features to a given problem or situation over time. A robot requires knowledge about time and causality and it must cope with the external world which changes over time and yields imperfect information to the agent. It is absolutely fundamental to be able to represent situations, goals and the effects of actions for effective behaviour and this involves causality: the possibility and results of actions and the consequences of situations to achieve goals [McCarthy 1960]. Causation however is an implication but it does not give necessary or sufficient conditions. It is based on a naive physics which comprises a set of heuristic rules which govern the world's behaviour and generates potential predictions. Rather than representing frame axioms explicitly, a modal operator UNLESS (W) which may be interpreted as "~W is consistent" may be introduced, i.e.. every state change is assumed to have each statement unaffected unless it is possible to deduce otherwise. This is in effect an example of default logic. Frame axioms do not need to be explicitly represented but may be represented by default: for all relations R which take state variables as arguments and for all state transition functions f: $\frac{R(x,s):MR(x,f(x,s))}{R(x,f(x,s))}$. This states that every action is assumed to leave every relation unaffected unless it is possible to deduce otherwise.

Logics other than the first order type have been developed to attempt to cope with some of the limitations of first order logic: default logic, modal logic and temporal logic [Webber 1983]. Non-monotonic logics attempt to circumvent the rigidity of logic as an approximation to intelligent reasoning while retaining the symbol processing paradigm. Default logic augments the normal deductive axioms with default rules to infer conclusions that cannot be deduced from the axioms. Such conclusions are consistent with the real world but yet not verified and may be rescinded in the advent of further information. They are based on the fact that propositions are founded on presuppositions and inherent assumptions. These presumptions are effectively inferences made in the absence of knowledge which may be refuted if subsequent knowledge is contrary. A default rule comprises three parts: a prerequisite that must be satisfied for the default to be applied, a set of assumptions that are consistent with the consequent, and the consequent itself which is inferred. Modal logic is based on the notions of necessity and possibility through the use of modal operators to qualify truth through such expressions as "it is possible that…" and "it is necessary that…" [McCarthy 1960, 1977; McCarthy & Hayes 1969]. Necessity corresponds to truth in all possible worlds □ (necessary truth) while possibility corresponds to truth in some possible worlds $\Diamond = \sim\Box\sim$ (contingent/possible truth), though if fact may be false. Actions and new knowledge affect the state of the world by changing it, falsifying some facts that were precisely true and making true some facts that were formerly false. By using the possible worlds interpretation (the universe of discourse U=<W,R> where W=set of worlds, R=reachability relation R⊂W), actions and knowledge relate possible worlds through consistency with knowledge or actions in a former world. Statements are not assigned a single truth-value but a spectrum of truth values, one in each possible world. Only when truth is established in all possible worlds is it necessarily true. Temporal logics enable reasoning about changes that occur in sequences of events. It explicitly represents the time dimension to enable reference to past events and rates of change. It reasons from past events to the truth of facts in the future. The past is a linear sequence of time points up to the present (now). The future is then a branching structure of time points from the present reference point to many possible worlds. Operators quantify the temporal interval for which propositions hold valid. Logical axiom schemas apply to temporal assertions in general, i.e. if Pq was true in the past it remains true now (LPq). Other axioms state that propositions hold for all future states (AGq) or hold true to the next future time point (EXq). Nonlogical axioms hold true forever in all possible future states. The future directed operator □ (always) and the operator ◊ (eventually) is similarly defined as in modal logic as $\Diamond f = \sim\Box\sim f$. All these logics are slightly different approaches to cope with a common feature - the frame problem. All these logics are slightly different approaches to cope with a common feature - the frame problem. The frame problem is essentially one that requires a theory of causal relationships. Kelly (1989) viewed intelligence as the ability to solve the frame problem, i.e.. to select relevant information or features to a given problem or situation over time. A robot requires knowledge about time and causality and it must cope with the external world which changes over time and yields imperfect information to the agent. It is absolutely fundamental to be able to represent situations, goals and the effects of actions for effective behaviour and this involves causality: the possibility and results of actions and the consequences of situations to achieve goals [McCarthy 1960]. Causation however is an implication but it does not give necessary or sufficient conditions. It is based on a naive physics which comprises a set of heuristic rules which govern the world's behaviour and generates potential predictions. Non-monotonicity allows the revision of beliefs in the presence of new observations and this implies the need to record the assumptions made in deriving each belief. The consequents of these rules are equivalent to beliefs subject to revision in the light of new information. This overcomes the negation as failure limit of logic where the failure of a proof implies the falsity of the hypothesis.

The situation calculus may alleviate the problem of causality. For dealing with causality, situation variables may be introduced as extra arguments to logic variables. These arguments denote situations as time instants – time(s). The situation represents the complete state of the universe at an instant of time. Relationships between situations (including the consequences of actions) is represented by a result function independently of the problem: s'=R=result(e,s) where s' is situation that results when event e occurs in situation s [McCarthy 1969, 1977, 1987]. The result function represents causation such that event e in situation s causes situation s' (or equivalently the situation s' that follows situation s satisfies event e). There is no a priori logical justification of inferring R(e,s) from s so it represents a belief. Hence R represents a new modal operator for first order predicate logic which provides a means for events changing the world state via causation. Such a modal operator can represent an action e to cause world state changes and such actions may be combined into sequences. By assuming that only certain conditions change from one situation to another all the other conditions are assumed to remain the same. It is possible to adopt a collection of statements (laws of motion) of the form $\varphi[s]$ such that every situation in φ is an occurrence of s to allow assertions about effects: $\varphi[s] \& \psi[e] \rightarrow \theta[R(e,s)]$. These laws of motion effectively denote a qualitative naive physics knowledge of the world. Such knowledge is required to make common sense of the external world and determine all future situations from the current situation. This common sense knowledge includes knowledge of object locations and their changes over time. History is modelled as a sequence of situations representing such knowledge. Indeed, most common general knowledge concerns the effects of actions on the world and the degree of localisation of these effects. Common sense knowledge is fundamentally related to deductive conclusions for the purpose of performing actions on the world on the basis of those deductions. In this sense, AI is fundamentally rooted in robotics. Hayes (1973) regarded these axiomatic laws of motion as frame axioms. They may be implemented at metalanguage level in the form: $\varphi[s] \& \psi[e] \rightarrow \theta[R(e,s)]$ if $\aleph(e,\varphi,\psi)$ where $\aleph$ is a conditional function on e, φ and ψ denoting the conditional relationship between s and R(e,s). This implies that $\aleph$ comprises a series of assumptions. Hence, causal connections can be implemented and their effects delineated in the world to solve the frame problem.

This leads to the concept of non-monotonicity which is closely related to belief, certainty and default reasoning. Human common sense relies heavily on the ability to make decisions in an uncertain world by using default reasoning to compensate for incomplete specifications. Most human thinking is not concerned with universal or even statistical truth but with norms as plausible explanations which may be subsequently invalidated by exceptions to the norm. The fact that the world is dynamic implies the need to revise beliefs to deal with inconsistencies generated by new information and data. First order monotonic logics are incapable of revising old conclusions as they are eternal truths. If new propositions are added to a knowledge base, all the original propositions are still true, and cannot be revised. In non-monotonic logics the acquisition of new knowledge allows revision or retraction of previous statements. Non-monotonicity provides a way of expressing a fact that is true unless there is an abnormal circumstance or exception to the general rule, eg. birds can fly (unless it is a penguin). Nonmonotonicity also solves the frame problem. Axioms are required to state which features of the situation are unchanged by actions and this may be accomplished through circumscription, eg. objects do not normally change their locations [McCarthy 1977, 1980, 1986]. Circumscription is a form of informal non-monotonic deductive reasoning often used by humans as a solution to the frame problem as well as more general problems concerning knowledge and inferences. Indeed, much real world information is based on facts which are almost always true with a few exceptions [Reiter 1980]. In these cases beliefs are held to be true in the absence of information to the contrary. Circumscription yields conjectures (jumping to conclusions) even to the point of allowing inconsistency (to which humans are prone) through over-generalisation of object classes, eg. "all birds fly" (but penguins don't). This is an expression of the standard case as a streamlined expression of probabilistic information to express default which allow the representation of "most" in the form "Most P's have the property Q" in the absence of information to the contrary without appealing to fuzzy logic or probability theory. By default, they represent possible hypotheses As default rules fill in the knowledge bases gaps they effectively act as meta-rules to complete the incomplete first order theory. A number of means have been proposed to suppress extensionality by using the "ordinarily" predicate: ordinarily (c,p) equate to "objects x of the class c ordinarily have the property p". Such axioms are inherited unless an abnormality is present. The only requirement is that the default rules are consistent with respect to all first order logic facts about the world. They may be formed from classical logic by extending it with modality defining consistency: Mp means that p is consistent with everything that is known. There are in fact 3 belief states: true, false or unknown. Unknown states imply the need for default reasoning due to incomplete information. Problem solving is an exploratory incremental process of refining sets of belief which model the world. Logic cannot make deductions on the basis of absence of knowledge but non-monotonic reasoning operates in the absence of contradictory evidence. Circumscription allows conclusions to be reached about relations for which no fact exists in the knowledge base by using assumptions. Circumscription lacks monotonicity and allows reasoning of the form: if $A \to p$ and $A \subset B$ then $B \to p$. Thus non-mononotonicity allows overgeneralisation and it is this overgeneralisation which simulates common sense reasoning. Hence circumscription is a form of heuristic default reasoning in the absence of observations to the contrary to augment first order logic. Default rules serve as rules of inference (conjecture) augmenting those of first order logic in an incompletely specified world. Default rules have the form:

$$\frac{A(x):B_i(x)}{w(x)}$$

where:

A=prerequisite of default
B=justification of default
w=consequent of default

Only positive information is represented in the database with negative information being inferred by default. As most information about the world is negative this represents considerable representational efficiency. A normal default has the form $\frac{A(x):B_i(x)}{B_i(x)}$ so the consequent is equivalent to its justification. They cannot produce inconsistencies and guarantee coherence. No new defaults can refute the justification of an already applied default as they have already asserted their justification. However sometimes defaults can interact generating contradictory consequents. A seminormal default has the form: $\frac{A(x):B_i(x) \wedge C_i(x)}{B_i(x)}$ whereby the justifications entail but are not entailed by their consequents. Once such a default has been applied only those conjuncts of the justification not entailed by its consequent are susceptible to refutation by other defaults as these are what can lead to inconsistency. The role of defaults is to fill in gaps in knowledge and function like meta-rules in that they are instructions on how to extend incomplete first order logic with the property that $W \subseteq E$ where E=extension, W=incompletely specified world. Default logics are a special case of generalised non-monotonic logics but remain in the framework of first order logic. The purpose of default reasoning is to determine one consistent set of beliefs about the world, i.e.. one extension. This extension is used for reasoning until such time as evidence forces a revision of those beliefs. The requirement is for believability rather than logical rigour - the maintenance of consistency is the major constraint. Rationality of beliefs is based on having acceptable reasons for the current set of beliefs.

Non-monotonicity allows the revision of beliefs in the presence of new observations and this implies the need to record the assumptions made in deriving each belief. The consequents of these rules are equivalent to beliefs subject to revision in the light of new information. This overcomes the negation as failure limit of logic where the failure of a proof implies the falsity of the hypothesis.

3.13.9. Truth Maintenance

These techniques are dependent on extracting knowledge from experts, codifying it into programming language, and transferring it into machine. Such programs are currently limited to ~10^6 lines of code. Several additional problems exist with knowledge bases – they are limited to 500-1000 rules for consistency maintenance. Truth maintenance is an important issue in large knowledge bases so that data which has changed should be updated so that the knowledge base remains consistent with the rest of the data in the knowledge base and with environmental data. Conflicting combinations of multiple assumptions can arise in Truth Maintenance Systems which depend on the order in which default rules are applied. This is fundamental to the frame problem. Truth maintenance deals with inconsistent knowledge by recording and maintaining reasons for beliefs, i.e.. truth maintenance which involves maintaining a database of beliefs that are consistent with propositions believed to be true unless there is evidence to the contrary [Doyle 1979]. It records the reasons for beliefs so that beliefs can be revised when subsequent information contradicts the initial (default) assumptions so that consistency is retained. Much reasoning involves presuppositions to infer conclusions that cannot otherwise be deduced because although the conclusions are consistent with the world they have not necessarily been verified. If conditions change, these justifications are examined to update the set of current beliefs incrementally. In TMS (truth maintenance system), each statement records an associated set of justifications (representing inferences), each of which represents a reason for holding the statement as a belief. These justifications function similarly to defaults and act as constraints. These justifications extend the set of beliefs in the light of what is known about the world and are used to determine the set of current beliefs by examining the recorded justifications. This allows a belief to be justified whenever no proof of the negation of the statement is known. For each statement p there are two possible states:

1. p has at least one valid reason and so is a member of the current set of beliefs;
2. p has no valid reason and so is not a member of the current set of beliefs.

Justifications for a belief comprise of other beliefs. The state of the belief depends on the state of the justifying beliefs. Assumptions are current beliefs whose justifications depend on a non-current belief. They essentially provide heuristic guidance. The justifications are actually inferences which are recorded to extend the set of beliefs and the assumptions form hypotheses. New discoveries and belief revision are inevitable in a changing environment. Defaults are based on assumptions which are not irrefutable. Only when situations change over time and affect actions to be taken do these overgeneralizations need to change. Abnormalities are dealt with when they contradict assumptions. Non-monotonic logic allows the introduction of new axioms that can invalidate old theorems to model active processes with incomplete information and subsequent new observations to generate new assumptions.

The Truth Maintenance System (TMS) provides the means for belief revision when discoveries contradict assumptions [McDermott & Doyle 1980]. Assumptions are retracted only by giving justifications for their removal and if the justification becomes invalid then the retraction is no longer valid and the assumption is restored. This provides the means to revise inconsistent beliefs into consistent ones through dependency-directed backtracking. Dependency-directed backtracking involves revising the current set of assumptions by operating only on the reasons for beliefs (justifications) rather than the beliefs themselves. Hence, it traces backwards through the justifications to find the assumptions which yield conflicting beliefs. Each time the search process makes a choice, it creates a world embodying the effect of that choice. These worlds are data structures that model the environment. It then removes at least one of the current assumptions that generates the conflicts. The TMS makes the necessary revisions in the current set of beliefs by updating them. TMS is tasked with maintaining a model database of formulas generated by an independent proof procedure which are globally consistent by detecting inconsistencies and adding axioms to eliminate them. It keeps track of formula p's justifications which entail p. p is labelled 1 only if all the formulas in the justifications of p are labelled 1; Mp is labelled 1 only if ~p is labelled 0. Hence all those formulas labelled 1 comprise the present beliefs and are "in" (for there is a non-circular proof for it) while all other formulas are "out" (there is no such proof). This allows jumping to conclusions. As new justifications are discovered decisions may require revision but with the proviso that ~p and Mp or p are not both "in". Hence each belief has 4 possible arguments: neither P nor ~P is believed (unknown), P is believed (true), ~P is believed (false), or both P and ~P are believed (unknown). The law of the excluded middle does not apply - the truth gap is effectively plugged. When a new justification is discovered for some formula which invalidates a current assumption, TMS must traverse the justifications through dependency-directed backtracking to find the assumptions underlying the conflicting formulas. To rule out an assumption Mp it must derive ~p. This is added as an axiom which invalidates and retracts the assumption Mp and so restores consistency. Hence it is possible to discard old conclusions in favour of new evidence or observations. This is similar to due process reasoning in that arguments are produced for conclusions and then arguments are produced against the opposing conclusions. Conflicts are then resolved either by challenging the premises of the arguments or by challenging those justifications which represent the arguments. Advocates collect evidence in favour of a hypothesis while sceptics collect evidence against the hypothesis. The debate continues until a decision made based on the “minimum mutilation” principle. This due process reasoning overcomes the negation as failure limitation of logic. TMS provides the means for dealing with incomplete but precise information as distinct from vague information as characterised by fuzzy logic. De Kleer (1986) extended TMS to allow manipulation of assumptions as well as justifications (assumption-based TMS) and it used chronological breadth first backtracking. Chronological backtracking revises the last choice of possibilities but dependency directed backtracking is depth first and revises the choice of possibilities that contributed to the failure. These assumptions represent the contexts under which the proposition holds and represent primitive decisions. The world is characterised by a set of assumptions. The primary source of justifications are deduction rules which are declarative truths. If there are n assumptions then there are potentially 2^n contexts. Unfortunately ATMS is computationally expensive to implement. TMS has been implemented in the object oriented knowledge based system KEE which uses a frame type format of object frames [Filman 1988]. Non-monotonic default logics are particularly relevant to semantic nets with inheritance links where higher level properties are assumed to be inherited unless contradicted. Ordinarily inheritance of properties is assumed to be valid unless abnormality conditions arise. Beliefs are essentially psychological states that have direct reference to the environment. They cannot be adequately modelled by logic as they are indeterminate and subject to constant change.

3.13.10. Modelling Uncertainty

There is also a need for numeric representations for probability/confidence/certainty to calculate the degrees of belief in a conclusion [Post & Sage 1990]. This again requires some form of numerical valuation to belief based on probability or possibility theory, i.e. Bayes theorem, Dempster-Shafer theory, or fuzzy set theory. This allows a range of possible outcomes rather than one default outcome. Default choices can thus be made probabilistic rather than restricting defaults to being defined with probability of 1 or 0. Human reasoning is characterised by the problem of uncertainty. Classical logic is limited in that it can apply to certain knowledge only, i.e.. all propositions are either true or false – this stipulation gives the law of the excluded middle its power. Uncertainty arises due to imprecision in knowledge characterised by the possibility of truth or falsity. This implies the need to accommodate propositions with a truth value intermediate between zero and unity [Pang et al 1987]. There are always limits to the accuracy of data available to any computational system concerning the real world. The principle of incompatibility states that as the complexity of a system increases, the ability to make precise statements about its behaviour decreases. Sage (1981) expressed this as an extension to Godel's incompleteness theorem: "Any system is…deficient with respect to external utility measures… These utility measures cannot...be expressed exclusively in terms of system performance measures since they are external to the system".

Shannon's information theorem states that the information carried by a signal may be quantified by:

$$I(x) = \log_2 \frac{P_2(x)}{P_1(x)}$$

where:

P_2 = probability of correct identification of signal x on receipt of the signal
P_1 = probability of correct identification of the signal x prior to the receipt of the signal

The amount of information is determined by the amount of uncertainty the message dispels. Furthermore, the human observer may be regarded as an information channel of finite capacity. This information capacity is ~2.3-3.2 bits between stimulus and response when subjects are required to make absolute judgements when exposed to a single stimulus dimension. Miller's limit suggests that there is a limit of 7 ± 2 items to short term memory capacity of a few seconds duration. Chunking can increase the effective capacity by efficient recoding of the items into shorter lists – this is the process used in SOAR for generalised learning.

Human reasoning and thought is characterised by uncertainty. Uncertainty arises due to random variation or imprecision. There are always limits to the accuracy of data available to any computational system. The principle of incompatibility states that as the complexity of a system increases, the ability to make precise statements about its behaviour decreases. Sage (1981) expressed this as an extension to Godel's incompleteness theorem: "Any system is...deficient with respect to external utility measures. ...These utility measures cannot ... be expressed exclusively in terms of system performance measures since they are external to the system." Heuristics are one example of vague truisms. They are generally used to generate weak forms of implication, leading to logically inconsistent conclusions due to their uncertain validity. Often in real terms, the conclusions are in conflict though not inconsistent. Inferencing techniques used with degrees of belief or subjective probabilities are usually inductive rather than deductive. Traditionally, reasoning under uncertain conditions is modelled probabilistically by Bayes theorem such that belief in a hypothesis increases with the evidence in favour of the hypothesis. However, human beings use judgement heuristics and ignore prior probabilities, eg. the representation heuristic uses typically as the deciding factor in acceptance of beliefs.

3.13.11. Problems with AI-based Cognition

Schmidt et al (1978) suggested that the generate-and-test paradigm for planning was not sufficient in that it can deal only with well formed problem statements. It is best used with fixed search spaces and may be used in a model driven process (generate) or data driven process (test). The model driven process has the advantage that a hypothesis is generated only once and uses the data as evidence for the hypothesis. However, much of human planning is based on ill formed problems and this requires a hypothesise-and-revise paradigm. The hypothesise and revise paradigm involves predicting about actions that have yet to occur and making observations consistent with these expectations. This is necessary due to uncertainties in the situation. The search space may change radically due to inputs. The hypothesise and revise paradigm remains strongly focussed by a single hypothesis which can be revised to accommodate unexpected events. The methodology is based on the fact that humans use observed actions of agents to attribute intentions and goals to the agent. Observers give intentional interpretations to an observed ordered sequence of events. Human cognition is grounded in knowledge and specific knowledge is used by the subject to infer the plans that an agent is pursuing. The subject uses logical connectives between actions that have subgoal-goal relations in the plan that the observer attributes beliefs and goals to the agent. Planning involves obtaining the knowledge necessary to perform actions and this requires mental actions such as observation and reasoning. An active agent such as a robot assumes that knowledge is representative of true beliefs. The robot forms new beliefs by observing the external world and its own internal state. Hence action understanding is a process of plan recognition. Observers generally generate a single hypothesis rather than maintaining a set of alternative hypotheses. If the hypothesised plan is disconfirmed by observations then it is abandoned and an attempt is made to generate a new hypothesis based on the new observations. This is both a model driven and a data driven strategy. The hypothesised plan provides the basis for generating expectations while expectation violations provide the basis for revision. The implementation of this approach was as the BELIEVER system which divided knowledge sources into World, Person & Planning domains. The knowledge sources provided the general domain of concepts and their relations and dependencies. A model was regarded as a collection of instances of the knowledge domain. This provided a distinction between the representation of general semantic knowledge and the representation about situations and events (episodic memory). The World domain comprised concepts and relations concerning objects in the environment. The Person domain comprised concepts and relations concerning agents. The Plan domain comprised concepts and relations about planning. Action sequences represented in plan schemas cause the world model and person model to be changed. The plan model represented the intermediate states. Rules of motivation in the person domains determined the pursuit of goals. Hypothesised plans in the plan model were grounded in observations or revised when observed actions deviate from expectations in the world model and person model are violated. If many alternatives exist from which there is no logical way to select a candidate then a wait-and-see strategy is adopted to postpone the selection of the revised hypothesis until further inputs are obtained.

The world is non-deterministic and uncertain and robots must respond flexibly to this world. The problem of control in the real world is ill-posed in that the conditions of stability, uniqueness and solubility are generally not satisfied. The problem of error correction for human cognition is fundamental in a changeable environment. Dissonance provides the means for model stability. This involves matching of sensory input to stored images in memory. Temporal stability of the model provides continuity of awareness while attention focuses priority on the basis of degree of dissonance. We know that human level performance for computational machines can be achieved - we are the existence proof of that. However the evolutionary record suggests the degree of difficulty in certain capabilities [Moravec 1984]. Animal perception, control and locomotion have evolved over 500 million years since the Cambrian explosion, but logical thought evolved over at most over a few million years with the arrival of the australopethecines. In addition, large segments of the human brain are devoted to perception and motor control while our higher logical functions are learned and propagated culturally (since Aristotle). As Albus (1991) suggests, the prime purpose of animal brains is as a control system for the generation of biologically advantageous behaviour such as searching for food, avoiding danger and procreation of offspring. Lower level skills appear to be more difficult to emulate than higher level skills such as logical reasoning. However, sensory-motor control systems have been evolving for ~10^9y to increase survival rates while human logical reasoning appeared only ~10^6y ago. Furthermore this reasoning capability combined with our learning capability has enabled us to master high level abstract logical thinking. However humans perform this kind of thinking inefficiently and only after a great deal of training yet the computer performs them well. The human brain evolved to provide flexible control in the perception-action cycle. Perception and motor control are performed by a large fraction of the human brain in the form of dedicated circuitry. The cerebral cortex comprises 30 x10^9 neurons devoted to high level reasoning while white matter comprises 70 x10^9 neurons devoted to sensory-motor control in the human brain, i.e.. ~10^{11} neurons in total. The visual system alone occupies ~10% of the brain. It provides the mechanism for early warning information at a distance to enable strategic rather than just tactical behaviour. The retina provides ~10^6 pixel images which are processed in the visual cortex at a frame rate of ~10/s. Neurons selectively respond to local features such as intensity boundaries and lines involving around 100 neurons per pixel, i.e.. ~10^8 neurons in total. Hence around 100 local features may be processed simultaneously by the 10^{10} neuron vision system at a rate of 1000/s. The human brain as a whole has a processing speed of ~10^6 MIPS. Memory resides in the synapses between the neurons: small molecules which dissipate over time control STM while large stable proteins are synthesised for LTM. These molecules affect the firing threshold of the synapse and there are ~10^2-10^3 synapses per neuron suggesting a human memory capacity of ~10^{13} bits. Amateur high level thinking requires ~10^8 IPS while instinctive low level perception action requires ~10^{11} IPS. Expert level performance involves transferring high level thinking into the computationally efficient perceptual-motor parts of the brain and so involves between 100 and 10^5 MIPS processing. These lower level skills-based perceptual and behavioural capabilities are a direct result of mobility which is peculiar to the animal kingdom – mobility requires a wide variety of behaviours and their control to cope with the wide range of environments encountered because of mobility [Moravec 1984]. The corollary of this is that only robots have the prospect of developing intelligence – the "brain-in-a-vat" approach of symbol processing will be insufficient.

Planning is inadequate in situations where events occur unexpectedly. Recovery usually requires replanning and in a dynamic environment this will be required almost continuously. Indeed, it is not clear if rule-following is in fact characteristic of expert behaviour, eg. in learning to drive [Dreyfus & Dreyfus 1987]. In the learning process, several degrees of competence and skill are attained sequentially from novice, through advanced beginner, competent and proficiency, to expert. At the levels of novice to proficiency, rules of increasing generality are used. At the novice level, context-free rules are used with context being ignored. At the advanced beginner level, the learner begins to become context-sensitive by recognising concrete conditions. Competence indicates performance which deviates from the previously learned rules in relation to goals. Proficiency indicates that previously conscious decision-making has become unconscious. At the expert level, rules are no longer used to guide behaviour as the behaviour has become effectively instinctive. The skill is no longer a sequence of problems to be solved or plans to be executed. Symbolic techniques are inflexible by themselves in coping with unforeseen circumstances and are brittle in the face of deviations in the real world [Beer et al. 1990] - much of adaptable behaviour in the real world does not involve logical reasoning. There is little doubt that humans use logic in rational knowledge manipulation in the process of reasoning. Not all human information processing involves logic manipulation however. Episodic memory does not involve the manipulation of rules but attempts to solve problems by direct reference to memory – a pattern recognition process. Indeed, much of human thinking is decidedly fallacious from a logical point of view (even excluding reasoning which violates distribution and negation rules which are common): the fallacy of alluding to facts out of context, argumentum ad misericodiam (fallacy of appealing to sentiment), argument from authority, argumentum ad ignorantium (truth asserted by a lack of evidence against it), argumentum ad hominem (falsity established due to the messenger rather than the message). Autonomy requires a variety of different behaviours, hence behaviour-based methods.

Durfee & Montgomery (1991) broadened the notion of goal directed planning in the form of 6 dimensional behaviours. Each dimension is characterised by a question - goals represent WHAT is to be achieved, plans represent HOW the goal is to be achieved. The other dimensions have the form WHO? WHEN? WHERE? and WHY? The WHY? dimension in particular comprises the motivational component of behaviour. Agre & Chapman (1990) outlined their plan-as-communication view in which plans are merely one resource amongst many used for improvisation. This views the agent as performing goal directed activity through participating in the world rather than controlling it through planning. They outlined their theory of activity that routine activity is independent of plans. Their system PENGO used contingencies to improvise moment to moment in constant interaction with its environment. These contingencies were patterns of directed activity routines without building world models. It did however use visualisation of a series of scenarios. It regularly invoked and interleaved its repertoire of activity routines. The complexity of its actions reflected the complexity of the environment which essentially combinatorially increased the action possibilities. As a dynamic system, the mobile agent indulges in motor interaction with the environment and its sensory feedback is a product of both the environment and its previous actions. This is very similar to reactive behaviour based methods. Learning took place through the agent's interactions with the environment through the incremental evolution of a dependency network of rules which converted perception to action similar to memory based reasoning.

In humans, sense organs receive emissions from the environment which are converted to impulses to the brain. The brain sends impulses to the muscles to perform actions which affect the world. Semantic meaning results from interaction with the environment via the brain and body. The true test of AI techniques can only occur when it is integrated within a perception-action cycle. Planning requires some form of awareness on the part of the intelligent agent which may be characterised by varying degrees of sophistication:

(i) general proprioceptive bodily awareness for body positioning;

(ii) the situational awareness of body positioning with respect to the environment;

(iii) the agent’s awareness of self as an agent;

(iv) the social agent’s awareness between the self and other agents;

(v) linguistic self-awareness involving the symbolisation of the self as an agent.

The sensory and nervous system imposes structure to the environment through experience of interacting with it. The physical coupling between the organism and the environment via sensors and effectors provides a pragmatic grounding semantics in the real world. Language is also dependent on the physical interaction as it itself is merely a complex form of behaviour. Cognition is a form of embodied sensory-motor coupling by which the agent is a dynamic system perturbed by and perturbing its environment [Mallot 1997]. Cognition is fundamentally based on contextual experiences obtained through the body through its senses and actions. Cognition is thus sybaritic which justifies the situated robotics approach with its emphasis on environmental interaction. The problem of semantic meaning of symbols disappears – symbols are syntactic and open to semantic interpretation, and meaning does not reside in the symbol itself. The attribution of meaning is not achieved through symbolic rule-following but acts to provide awareness to the agent of its situation in the environment and provide context to choices of behaviour. The mapping between the symbols and the outside world that they represent is achieved through grounding the symbols via interaction with the real the real world. The manipulation, categorisation and naming of real world objects and events occurs through sensorimotor interactions with the real world. It is through concrete categorisations that symbolic abstraction arises in human cognition from the “bottom-up”. Categorisation marks the beginnings of symbolisation whereby perceptual categorisation provides the basis for filtering sensory data. The emergent internal symbols are still strongly constrained by their grounding in the real world, and so are not arbitrary like formal symbol systems. Patterns of behaviour representing generalisations may be regarded as proto-symbols. Our language is grounded in physical sensorimotor patterns that are generated when we interact with the external world. Much of our language comprises of physical metaphors based on our bodily interactions with the world. Biologically, this is obvious – behaviour evolved through natural selection imposed by the physical environment. Object representation is thus a sensorimotor phenomenon rather than a product of symbol processing. Similarly, the morphology of the robot provides the physical embodiment of important characteristics of the robot-environment interaction. Behaviour as agent-environment interaction is an ongoing process – there is no final output to a computation as data is constantly updated and recomputed. Motor actions alter the state of the environment and determine the sensory data received continuously. The physical environment forms an integral part of the sensorimotor coordination process. The agent generates correlations through interaction with its environment, i.e.. grounded in the environment. Behaviour is an expression of the regularity of the interaction dynamics between the agent and the environment. The computation required in robotics is thus not a linear program but the result of the interaction of many processes.

3.13.12. The Role of Affect in Artificial Intelligence

Few AI systems to date employ affect (emotion and motivation) in their design, though recent activities have suggested the inclusion of affect for value assignation (which is also important for learning). We review the nature of motivation and emotion in cognitive psychology to ascertain their utility in artificial cognitive systems. Very little is known about the role of motives and emotions in human cognitive processing. The James-Lange theory of emotion suggests that emotional experience is the product of cognitive processes indicated by body (visual) responses concerned with the maintenance of homeostasis. Emotion is the product of the limbic system which resides below the cerebral cortex and it reaches from the hippocampus to the hypothalamus to the anterior thalamus. Emotional state then becomes a physiological and cognitive process where cognition is directed by the degree of arousal. There is little doubt that the motivational portion of human cognition is an important complement to the rational component [Kent 1978]. It provides an important regulatory function and all cognitive processes are subservient to the mechanism of goal generation and intention which essentially turn passive information processing into an active system. It is the motivational component which produces goals and desires to be achieved and generates emotions as an alerting device. The intensities of a motive determines the intentions of the agent. The brain evolved as a control system to seek biological needs (pleasures, eg. eat) and avoid biological harm (pain, eg. being eaten). Albus (1991) suggested that emotions provide estimates of biologically valuable state variables and play a central role in behaviour to prolong pleasure (reward) and avoid pain (punishment). Motives provide estimates of biologically valuable state variables with regard to importance and needs. Pleasure is a positive value state variable while pain is a negative value state variable. White (1959) suggested that motivations are innate behaviours that are different from basic drives in that they cannot be sated. They are selective, directed and persistent and represent a general state which gives way to primary drives when they are aroused. The tendency for exploratory behaviour is an innate motivation on which learning is based and is manifest as curiosity. It is well developed in mammals. Its persistence cannot be halted, only decreased as habituation produces familiarity. Basic to this concept is the notion of competence defined as the organism's capacity to interact effectively with its environment, and this has a motivational component. The tendency to explore the environment is on independent primary drive which cannot be learned. It is moderate in intensity but is persistent where other drives (hunger, thirst, etc) can be viscerally sated. There is a strong motive to explore the environment for novel stimuli. Curiosity only declines with familiarity until new novelty becomes available. Its origin appears to be in the hypothalamus. There is a fundamental need for stimulation. Indeed, it is opposed to fear which involves retreat in that it involves approach. It is a behaviour that operates on the environment to provide an effect which is fed back through the senses, but the effects are not geared towards satiation of organic needs. It is fundamental as a motivation that satisfies a basic need to cope with the environment. The exploratory behaviour is not random but directed in that it focuses attention on stimuli with a purpose. Play is a form of exploration through experimentation and manipulation to enable children to learn about their environment. Such goal directed motivational behaviour through emotion is generated in the limbic system at the core of the human brain, the hypothalamus and mesencephalon. The hypothalamus detects the need states while the cortex provides perceptual information of stimuli both of which connect to the limbic system which is the centre of the reward system and maintains certain behavioural outputs. Avoidance behaviours may be implemented through inhibitory connections. Fear originates from the posterior hypothalamus, anger from the amygdala, pleasure from the septum, pain from the hypothalamic nucleus, and arousal level from the pituitary. The cerebral cortex also has a role particularly the right cerebral cortex with regard to appropriate tones of emotional recognition and response. Drives are a set of states that dictate the needs of the organism such as hunger. They are homeostatic in that they operate to maintain a particular level of activation in the organism through feedback mechanisms. Nonhomeostatic reflexes are those states which assign the highest priority to behaviours to remove a particular environmental stimulus and persist until the stimulus has been removed, eg. response to pain. This is analogous to a priority interrupt system such as modelled through behaviour-based architectures. Emotions serve to provide anticipitatory mechanisms to invoke behaviours before such stimuli are encountered, eg. fear associated with potentially painful stimuli to invoke avoidance behaviours. Hence emotions are anticipitatory motivations related to reflex and drive objectives. Pleasure and pleasure seeking behaviour may be similarly characterised.

Cohen & Levesque (1990) maintained that intention maintains a rational balance between beliefs, goals, plans and commitments. The motivational state of an animal or human depends on its internal state, its perception of the external world, and its current behaviour and expected effects on the world. An intelligent agent acts on its intentions, adopts intentions that are feasible, discards intentions which are infeasible and alters intentions as situations change. Intentions differ from plans in that they involve commitment. Communication involves the ability to reason about the intentions of others. Intentions guide the planning process and are determined by the present beliefs. Intentions are a means of settling on one state of affairs from many conflicting requirements to overcome their limited resources. Planning is the means to achieve intentions selected and also provide a consistency standard from auxiliary intentions. The agent believes that it is possible to fulfil the actions and the agent believes that he will in fact do the action. Beliefs are the justification for holding intentions. Negotiation is a form of social commitment and intention as a basis for communication and social interaction. Intentions provide the means for distinguishing between actual and possible events by acting as persistent goals (effectively a chosen set of possible worlds) that are retained as long as certain conditions hold. Their persistence is limited until the goal has been achieved or becomes unreachable. Each set of possible worlds comprises of a sequence of events extended temporally into the past and the future. Each world specifies what could happen in the future. Beliefs and goals characterise what the world model would be like if the agent's beliefs and goals were true and these are the worlds which are selected. As beliefs change so too will the agent's goals. Indeed agents reason about complex plans without having complete knowledge of how the world will change and this information is acquired during plan execution.

Motivational emotions allow specific activities to be selected in accordance with relevant goal stimuli. This selection process appears to be heuristic rather than algorithmic. Emotions may arise due to multiple goals, needs and motives with limited resources in the face of an environment that is not static. Viewed in this way emotions may be regarded as a fundamental requirement for intelligent information processing systems [Sloman & Croucher 1981]. They provide the means for goal revision in a dynamic environment and with conflicting or inconsistent goals. A central administrative processor behaves as an overall operating system and may provide motives as the prioritising decision-making capability to select goals and plans according to the overall motives. Some of these motives may be pre-programmed like drives. Some will be generated according to other factors. Hence a priority system will be required to determine the importance and urgency ordering. Satisfaction or violation of motives produces emotion as pain to invoke corrective action or pleasure to maintain the executed action. However the pleasure/pain criteria for the selected of goals and actions views emotion as an evaluative function. This is rather a narrow view as it is more likely that pleasure/pain evolved in animals to ensure the satisfaction of basic needs [Jacob 1977]. The CogAff architecture attempts to model emotions as the result of a control hierarchy – the first layer is reactive, the second deliberative and the third meta-management [Scheutz et al. 2001]. The meta-management layer provides self-observation and self-reflectiveness by orienting attention. The primary emotions result from the interaction of primary processes, secondary emotions from deliberative processes and tertiary emotions from meta-management processes. It is not currently clear how well the CogAff architecture indicates these emotions through behaviour. Zrehen & Gaussier (1997) used a motivational system for locating unobservable targets with respect to visual landmarks based on their PerAc navigation system. When the PerAc system encounters a target initially, it triggers an exploration reflex which generates looping behaviour around the target whilst storing its sensory stimuli. This provides the basis for learning landmarks while associating landmarks with azimuthal directions to the targets. Landmarks are stored in neurons which act as place cells similar to hippocampal pyramidal neurons. Hebbian learning is used to associate a landmark neurons in the landmark neuronal group with the place cell in the hippocampal neuronal group. The motivation system provides a "drive" to be reduced by appropriately selecting its behaviour. The learning is controlled by the global parameter called vigilance. The drive for hunger invokes exploration behaviours while the drive for sleep invokes homing behaviours.

Motivational systems have a central role to play in providing a mechanism for reinforcement learning in the cerebellum such as operant conditioning. Behaviours which activate the reward system tend to occur more frequently over time. Those which fail to activate the reward tend to occur less frequently over time. Operant conditioning appears to occur between the reward system and the inputs to the basal ganglia. Inputs from the cerebral cortex and the thalamus produce patterns of activation in the basal ganglia to produce the desired behaviour and if the resulting behaviour produces favourable stimuli then sustaining input from the reward system maintains this pattern activity. Hence operant conditioning is merely classical conditioning between the basal ganglia and the reward system of the limbic system with the reward system acting as reinforcement.

Emotions are an essential component of higher level cognition to enable rapid decision-making. It provides greater flexibility than purely reflexive behaviours. Emotions are a major determinant of motivational behaviour: much behaviour may be characterised as pain aversion or pleasure-seeking. The concept of gratification/deprivation provides the basis for stimulus-response behaviour through utility reinforcement. Individuals avoid situations of deprivation which result from unsatisfied needs and attempt to achieve situations of gratification resulting from sated needs. Thus, they are commonly invoked in providing contextual evaluation in reinforcement learning. Such contexts are acquired through past experience and are used to select positive outcomes over negative ones. Emotions are thus the basis of value judgement. Emotions may be regarded as a function of both arousal due to sensory inputs and the current cognitive state. Expectation (related to habituation) provides a stable background against which sensory inputs are analysed. An important aspect of emotion is its social role in communicating an agent's internal state to other agents. One way in which emotions may be modelled is in affecting the learning gain with the gain being proportional to the emotional intensity. However, this role does not provide good performance in robotics experiments in that an agent may still receive negative/positive reinforcement for good/bad action selections respectively, particularly in the case of delayed feedback – this is the action selection and reinforcement apportionment problem [Gadanho & Hallam 1998]. Lazarus (1982) suggested that emotional reaction requires prior cognition before emotional evaluation can occur. Such a view regards emotion as a mechanism of cognitive attention, motivation and expectation to select sensory information within a limited information channel. Furthermore, emotion is the result of organism-environment relationship including social interaction. Zajonc (1980) suggested that emotional reactions to stimuli are the result of separate control systems from cognition, though the interaction between these can occur. For lower organisms, emotions are the only form of behavioural control and provide the basis for involuntary approach-avoidance behaviour without the need for prior cognitive processing. Such emotional reactions are primary in that they provide the basis for motivation and expectancy and they are basic to the evolution of animal behaviour. They are not subject to attentive control and comprise reflexive behaviours. They are activated directly by sensory input and basic perceptual processing providing a limited form of pre-categorisation of stimuli. Indeed, emotion without doubt is a function of physiological arousal generating visceral and motor activity. Furthermore, emotion has been localised to regions of the mid-brain such as the amygdala, hypothalamus and hippocampus under the control of the limbic system with direct pathways from the retina (retinohypothalamic pathway). The limbic system predates the evolution of the neocortex in humans which is responsible for high level cognitive processes. Emotions appear to be characteristic of the right hemisphere of the brain while cognition appears to be characteristic of the left hemisphere. This form of automated behaviour is to be differentiated from learned automated behaviour. The fact that rational cognitive processes are capable of suppressing and controlling emotions does not negate this hypothesis – the evolution of two separate control systems, one limited in repertoire (affect), the other of great flexibility (cognition) would require some form of hierarchical mediation to evolve to avoid conflict. However, it is questionable whether the term emotion is applicable to rigid, built-in avoidance-approach responses or whether emotion does require cognitive appraisal. Indeed, recognisable emotions do appear to be exhibited only by higher animals. It is conceivable that emotion is in fact the result of activation across the cerebral commissures between genetically controlled precursors and higher level cognitive processes in each cerebral hemisphere.

Kemper (1987) suggests that there are four primary biological emotions – fear, anger, depression and satisfaction. However, there are six basic facial expressions of emotion (happiness, sadness, fear, surprise, anger and disgust) with the possible addition of two more (interest and shame). Additional secondary emotions are socially constructed. Interpersonal attitudes are friendly, hostile, superior, inferior and loving which form the basis of social skill. It is significant that there are more negative emotions to invoke avoidance behaviours than positive emotions. The intensity of each emotion depends on the agent’s internal feelings of hunger, pain, temperature, restlessness and eating. Emotions provide feedback on the visceral status of the body. The intensity determines the activation of the endocranial system. From an evolutionary point of view, biological emotions must be adaptive with survival value and indeed, their biological basis lies in the endocrine system of the brain. Their activation indices causal affects in the autonomic peripheral nervous system of the bodily organs, i.e.. heart, lungs, blood vessels, skin and digestive system. These emotions also exhibit universality of facial expression in diverse human cultures. Primary emotions also have their precursors in phylogenetically less advanced species. There is good correspondence between human and facial expression for anger, fear, happiness and possibly depression. Fear and anger energise urgent activity when exposed to threat or danger. Fear invokes the action of epinephrine (adrenaline) and anger invokes the action of norepinephrine (noradrenaline) in the sympathetic nervous system. Satisfaction provides the mechanism for positive reinforcement in learning to pursue important survival-related goals. Depression has survival value in that it enhances social cohesion when evoked by the loss or change in social condition or integration. Satisfaction and depression involve the variable action of acetylcholine in the parasympathetic nervous system. The two branches of the autonomic nervous system regulate the activity of bodily organs in response to emotional states in complementary fashion – the sympathetic nervous system is responsible for activation while the parasympathetic nervous system is responsible for deactivation. These primary emotions have become adapted to social situations involving power and status interactions in social groups: fear results from subjugation to the power of other agents; anger results from the denial of expected status by others; depression results from the denial of status due to own fault; and satisfaction results from the achievement of expected status. Other secondary emotions, of which there is a large class, requires some degree of cognitive awareness and differentiation of self from others and are thereby imparted by social interaction and socialisation in the context of primary biological emotions, eg. pride, shame, guilt, hate and love. They tend to appear late in human post-natal development. Guilt results from anxiety due to the use of excess power against others. It exhibits association with fear of punishment, especially the withdrawal of love. It is the outcome of social learning through punishment. Pride results from satisfaction of status through achievement and self-worth. It is the outcome of social learning through reward. Shame results from anger at self for discredited status and unworthiness. Shame may be the mechanism whereby self-anger is seen as depression. Emotions may not just be linked to individual processing but also higher order combinations may yield emotions, eg. fear-depression generating anxiety. As long as society defines new social situations, so the number of possible emotions are effectively limitless.

3.14. HYBRID APPROACHES TO INTELLIGENT CONTROL

Behaviour-based robots lack explicit goals and goal manipulation. A situation-driven set of behaviours is only effective if enough environmental cues are available to determine what actions are applicable. If a task requires knowledge about the world and requires reasoning and memory rather than just perception and survival, then reactive behaviour in insufficient. Planning involving conditional effects is undecidable such that a plan cannot be guaranteed to succeed without execution time additions to the plan. Hence, the planner must work in close cooperation with the execution supervisor which monitors the results of the application of the plan in real time. Rather than replanning from scratch which is time consuming and wasteful, a more efficient approach is to empower the executor with the necessary decision capability to repair plans dynamically about recovery procedures and error correction. This must be online and event driven. This involves tighter coupling between sensing and action through reactive behaviours to improve responsiveness in dynamic environments. An autonomous system requires both goal-oriented planning and robust reaction.

3.14.1. Integrative Approaches

Payton et al (1990) have considered the problem of integrating high level planning with lower level reactive behaviours. Reactive behaviours require tight coupling between sensing and action so they viewed the lowest level of the control hierarchy as a subsumption-type architecture. Plans were viewed as resources for advice rather than as constraints on predetermined courses of action. This allowed opportunistic advantage to be taken in unexpected situations. At each instant of time, the best choice of goals was selected without discarding alternate plans. The internalised plans were represented as gradient fields computed from graph searches to minimise the abstraction requirement. The low level behaviours implemented were obstacle avoidance and navigation (wandering) but without subsumption-blocking of behaviours allowing each behaviour to communicate with each other.

Maes (1990) suggested that explicit goal handling with reactive capabilities eliminate the need for re-planning or re-programming with altered goals. Goals effectively bias choices of action (action selection). One approach to such action selection is through emergence via the local interactions between action-oriented modules. Each module represents an operator of the classical planner and may be described by the triple $(c_i, a_i, d_i, \alpha_i)$ where c_i=pre-condition list, a_i=action effect list, d_i=action effect delete list, α_i =activation level. The module is executed when all the pre-conditions are true and the activation level exceeds the threshold. Modules are linked into a network with three types of link between modules: successor links, predecessor links and conflictor links. Successor links exist when every proposition p that is a member of x's add list is also a member of the pre-condition list of y. Predecessor links exist for the inverse relation. Conflictor links exist when every proposition p that a member of the delete list of y is also a member of the pre-condition list of x. Modules use these links to activate and inhibit each other with inhibition providing the means to deal with interacting sub-goals. Activation energy is determined by the input from the current situation according to its match to modules and by goals which determine the activation threshold of goals. Protected goals remove activation from modules that would undo them. Spreading activation across the network of modules equilibriates to the best actions determined by the current situation and the current goals of the agent. Successor links propagate action forward to successor modules and predecessor links propagate activation backwards to predecessor modules by a fraction of their respective originator module activation energies. Conflictor links decrease the activation level by a fraction of the originator module's activation energy when a module undoes a true pre-condition of another module. All inputs of activation from a module is weighted by 1/n where n=number of propositions in the pre-condition/add/delete list as appropriate. As the state of the environment changes, so too do all the activation patterns. Global parameters may be used to alter the activation dynamics through threshold adjustment, input activation adjustment, goal activation adjustment, and conflict activation adjustment. The pre-conditions act as subgoals. A sequence of modules are highly activated which transform the current situation into the goal situation through forward spreading from the current state and backward spreading from the goal state. As the links are only local among modules, different paths may be evaluated in parallel. The system is a marker-passing system without variable use in that objects in the environment are represented by only those features relevant to the agent. Hence, new operator modules are not required when encountering new objects. Goals are specified in terms of functional constraints on the objects in the environment. The control structure is emergent when particular actions are activated. The interaction dynamics between action-oriented modules establish the sequence of actions in a distributed manner in response to environmental conditions and global goals. The local interactions between components forms an emergent global structure.

Mataric (1992) suggested that rather than utilising hybrid approaches with separate reactive and planning modules, it is possible to maintain a map representation based on discovered landmarks distributed over the subsumption architecture of behaviours. This represents a fully integrated approach. The robot's control system comprised of behaviours specific to basic navigation, landmark detection and map computation. Basic navigation and obstacle avoidance involved strolling forward unless an object was in front of the robot, turning away from objects, and aligning with boundaries to follow them. Landmark detection involved matching sensory data to stored landmark signatures. The mapping algorithm mapped the structure of the environment based on the spatial relationships of the landmarks. The map was coarse and encoded as a topological graph with each node representing a landmark and the links representing adjacency. Each landmark was a behavioural set of rules and the links operated as message-passing connections. Each landmark behaviour received inputs from the landmatk detector, while sending outputs to other landmark behaviours. Loccalisation involved comparing stored landmark signatures with the current sensory data. Expectation was generated by spreading activation through the graph, so adjacent landmarks were primed forming a path to the goal landmark, ie. path planning.

3.14.2. Symbolic Connectionist Approaches

Handelman et al (1990) introduced a robot control methodology modelled on human motor skills learning. There are three phases of skills learning of sensorimotor reflexes:

(i) The early cognitive phase where the task is approached consciously through inferencing by the beginner;

(ii) The intermediate associative phase where response patterns begin to emerge;

(iii) The final autonomous phase when the execution becomes automatic and reflexive for the expert.

The skills learning process includes a shift from conscious declarative processing to automated procedural processing, ie. the declarative representation becomes "chunked" into a more condensed procedural representation through a compilation mechanism. An alternative representation of motor skills acquisition is through the repeated associative learning of input/output pairs which gradually determines the connection weights of an artificial neural network which becomes fixed on completion of the learning process. Such a method of implementing motor skills learning to robots should offer high levels of dexterity and adaptability. A three-level hierarchy was defined for fproviding reflexive behaviour. A rule-based execution monitor teaches a CMAC (cerebellar model architecure controller) network how to accomplish a task by observing the rule-based task execution. The knowledge-base represents declarative knowledge with inferencing abilities while the neural network represents associative procedural knowledge with pattern matching abilities. The neural network effectively captures the causal relationships between the robot and the environment. The knowledge base determines how to accomplish the control objectives using the rules. These rules include the learning strategies, task execution processes, network training data and performance monitoring. It teaches the neural network to accomplish the same tasks by allowing them to observe the task execution. They learn through generaslisation at which point they assume control responsibility. The nural net then fine-tunes its performance during task execution through reinforcement learning. The knowledge base continuously evaluates the networks performance. If errors occur due to the changing dynamic environment, the knowledge-base is re-engaged and the network re-trained. The basic control system implements low level servo-reflex on the innermost loop, reflex modulators in the middle loop to provide gain adjustment, and execution monitoring to supervise reflex modulator learning.

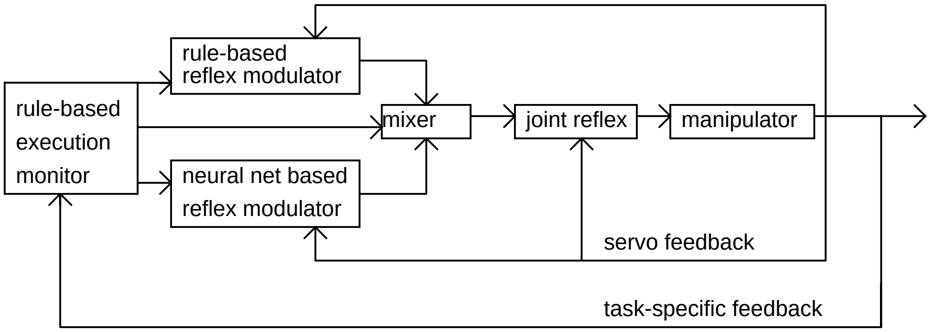

Figure 214 – Motor skills learning architecture (Handelman et al 1990)

The servo-level implements a standard computed torque control law. The rules in the knowledge-base are independent of the dynamic model of the manipulator. A goal-directed backward search is implemented by an inference engine through the knowledge-base of parameters, production rules and procedures. Parameters are values relevant to the control objective while the rules express relationships and dependencies between parameters. Within the conditions and actions are procedures that invoke time-critical control tasks such as low level joint reflexes and control laws. The rules provide the means to build higher level control actions out of the procedure building blocks to generate coordinated motion. All maneouvres are implemented as a series of steps. The neural network implementation is a CMAC which performs a perceptron-like table lookup for generating nonlinear mapping functions between multiple input-output variables. The knowledge-base control executor provides the means for it to learn by example. The necessary generalisations of sensorimotor behaviour are accomplished through connection weight adjustments via a steepest descent updating algorithm. The CMAC learns to duplicate the rule-based reflex commands. When the CMAC's outputs match those of the rule-based outputs, learning ceases and the CMAC outputs directly to the joint reflexes during manipulation. Once trained, the input values cause the CMAC to output control gains and commands to the servocontrol system of the manipulator. Thus the CMAC reflexively executes the required muovements. If the CMAC output deteriorates subsequently, control switches back to the knowledge-base while CMAC undergoes further learning. Hence, knowledge-bases can be used to supervise and train artificial neural networks and indeed, the neural network offers a more robust performance by fine-tuning its performance through reinforcement learning.

3.14.3. Neuro-Fuzzy Approaches

Fuzzy inference system (FIS), in the form of symbolic logic rules, provides human reasoning capabilities to capture uncertainties, which cannot be described by precise mathematical models (Figure 215). It is a powerful tool for control of such mathematically complex, uncertain or noisy systems [Wang 1994]. Fuzzy logic control offers an attractive alternative to the existing highly mathematical conventional controllers proposed for such complex systems. Herein, we refer to a complex system as one for which model-based classical control techniques cannot easily render a reasonable solution due to ambiguities, uncertainties and/or nonlinearities in the complex system's mathematical model.

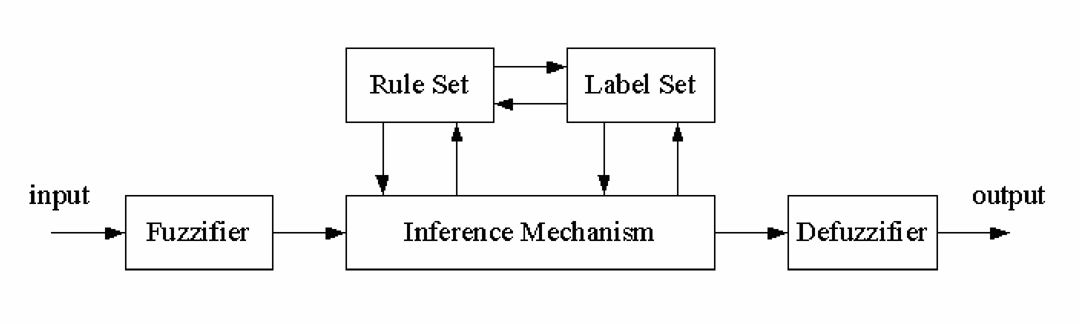

Figure 215 – Fuzzy inference system

The procedure of fuzzy controller design reflects its pragmatic nature by the reliance upon intuitive knowledge and an experiment-oriented approach. Firstly, the objectives of the design process are determined. Secondly, an intuitive understanding of the nature of the process to be controlled must be obtained. The accuracy of the intuitive understanding is vital since the fuzzy controller will be built based upon this. It may be obtained by any means available such as personal experience, advice from experts, mathematical analysis, and experimentation. Once the designer is confident that he or she has sufficient understanding of the plant and the method by which it may be controlled, a suitable type of FIS, such as Tsukamoto type, Mamdani type, or Takagi-Sugeno-Kan type, must be selected. Once this has been done, an iterative procedure aimed at achieving the original design objectives begins. This iterative procedure begins with the selection of sensor signals which the designer feels contain sufficient information about the plant to allow the fuzzy controller to infer the correct control signal. The ranges of these inputs are then determined and fuzzy sets defining concepts such as high, medium, and low, are created. A set of "if-then" rules and a set of appropriate labels are then defined on the basis of the fuzzy sets and a defuzzification method is selected, such as the centroid of area (COA), mean of maximum (MOM), bisector of area (BOA), maximum criterion, etc [Lee 1990; Mamdani & Assilian 1975; Jang et al 1990]. The COA method is the most widely adopted defuzzification strategy, which is reminiscent of the calculation of expected values of probability distributions. This completes the definition of a fuzzy controller; this controller must then be tested and improved further until the design objectives are met.

Once the design objectives for the controller have been met, the system may be commissioned. The application of fuzzy logic relaxes the need for an accurate mathematical model of the system by replacing the mathematical knowledge with human expert knowledge and intuition. As a result, performance of a fuzzy logic controller is a function of the quality of its embedded expert knowledge rather than a highly accurate mathematical model.

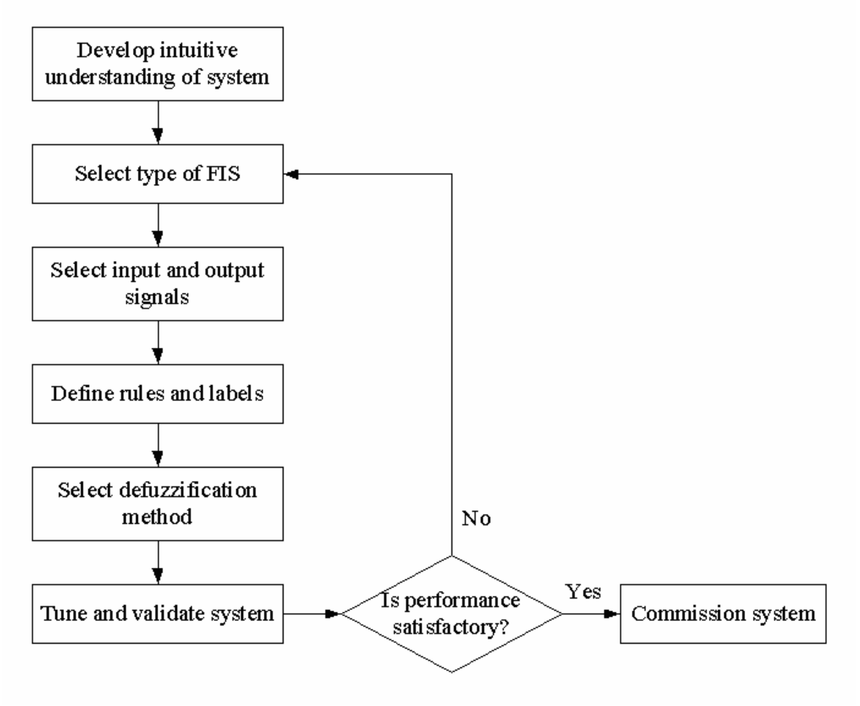

Figure 216 – Design procedure of a fuzzy logic controller

We can easily interpret how fuzzy controller works. They base their decisions on inputs in the form of linguistic variables derived from membership functions which are formulae used to determine the fuzzy set to which a value belongs and the degree of membership in that set. The variables are then matched with the preconditions of linguistic if-then rules, and the response of each rule is obtained through *fuzzy implication*. To perform a compositional rule of inference, the response of each rule is weighted according to the confidence or degree of membership of its inputs, and the centroid of responses is calculated to generate an appropriate control signal. FIS design is a subjective procedure which is adopted to express domain experts' knowledge. Developing quality expert knowledge, however, can be a time-consuming and costly endeavor. Furthermore, a human expert may find it difficult to express his/her control actions, which are often partly decided at a subconscious level, in terms of a set of constrained rules and membership functions. Also, in most instances, the available knowledge-based controller may adequately control a given process, but may not necessarily be the optimal controller. The process of knowledge acquisition is indeed a challenging problem in fuzzy logic; and up until now, there is not yet a systematic method for knowledge acquisition in conventional applications of fuzzy logic control. Consequently, in many practical instances, fuzzy control alone is not sufficient for addressing the complex intelligent control problems of robotic systems. Additional tools are necessary to achieve adaptation and learning capabilities. Some hybrid techniques have been proposed to augment fuzzy logic with other soft computing paradigms [Jang et al 1997] such as artificial neural networks (ANN), genetic algorithms (GA), and genetic programming (GP) to achieve the level of intelligence required of complex robotic systems. The fuzzy-ANN controller takes advantage of ANN for handling complex data patterns and learning from experience, the fuzzy-GA controller utilizes the ability of GA to optimize parameters of membership functions for improved system response, and the fuzzy-GP controllers utilizes the symbolic manipulation capability of GP to evolve fuzzy control rules [Akbarzadeh et al 2000].

There is recognition of the shared characteristics between fuzzy and neural systems and several studies have investigated the functional equivalence between various neural architectures and fuzzy systems to prove the capacity of both types as universal, model-free function approximators. Initially, the rationale between FIS and the interpolative approximation in ANN seems to be the same: to produce an output under similar but not exactly the same conditions. Several authors have shown that neural structures can be constructed that are equivalent to FIS [Buckley & Hayashi 1994, Kosko 1992]. Park and Sandberg (1991) proved that under certain mild conditions on single- hidden-layer RBF networks and fuzzy systems are equivalent. Nie and Linkens (1993) showed that the two systems have almost identical computational procedures whereby the matching degrees of a fuzzy control algorithm corresponds to the response function values of the RBF network. This equivalence is, in general, established under the following conditions:

- The number of centers or receptive field units of the RBF is equal to the number of fizzy if-then rules.
- The output of each fuzzy if -then rule is a scalar constant.
- The membership function within each rule is chosen as a Gaussian function with the same variance.
- The T-norm generator used to compute each rule's firing strength is multiplication.
- Both the RBF network and the FIS use the same defuzzification method (i.e. weighted averaging) to derive their overall outputs.

As a result, RBF network can be viewed as a mechanism for representing rule-based fuzzy knowledge by using its localized network structure, and performing associated fuzzy reasoning using feedforward computational algorithms. From the viewpoint of knowledge representation, RBF network is essentially a network representation of if-then rules, with each hidden unit representing a rule and the basis functions being equivalent to the membership functions in an FIS. We can apply what is known about one model to the other, and vice versa.

The majority of reported studies on fuzzy-ANN control applications address one or more of the following functions [Linkens & Nyongesa 1996]:

- Using ANN to tune fuzzy control systems;
- Extracting furzy rules from given numeric data examples using ANN learning algorithms;
- Developing hybrid systems combining ANN and FIS in various forms to implement.

There are two distinctive approaches to integrate ANN into fuzzy control system [Linkens & Nyongesa 1996]. One would be using ANN as an approximator or classifier for identifying a fuzzy control algorithm. The other one is to realize the process of fuzzy reasoning and inference through the structure of a connectionist network. The latter one is also known as fuzzy neural networks (FNN) approach, which are in general ANN whose nodes have 'localized fields' which can be compared with fuzzy rules and whose connection weights can similarly be equaled to input or output membership functions.

The simplest and most straightforward attempt in the merging of fuzzy and neural controller techniques is to make the ANN learn input--output characteristics of a fuzzy controller [Bouslapda & Ichikawa 1993]. Typically, a conventional fuzzy controller is used to generate the data for training the ANN. This approach, strictly, may not be classifiable as a fuzzy-ANN since the teaching controller is a pure fuzzy system and the resultant trained system is a pure neural system. However, it is recognized that an ANN controller provides smoothing of the actions of the fuzzy controller and has been shown to provide a controller performance that is superior to the original fuzzy controller in terms of smoothness and robustness, which justifies this approach. Extensions of this method combine the fuzzy logic controller with multiple ANN, one of which may be designed to function as a plant emulator and the other as an online compensator for the performance of the fuzzy controller [Khalid et al 1994].

Evidently, the most common trend has been to apply to tune the membership functions of the fuzzy rules. A typical approach is to assume a particular shape of membership functions and define its characteristics which can be learned by ANN. Horikawa *et al.* (1991), for instance, start with a fixed number of rules whose membership functions are subsequently perturbed through backpropagation until they fit given data. One of the problems with merely adapting membership functions is that it may tend to obscure the interpretability of the furzy neural system behavior, especially when the defining membership functions are changed so that they have different meaning from what the designer intended, or the same linguistic values have different meaning in different rules. Akbarzadeh-T *et al.* (2000) provide rule generation mechanism other than adjusting the membership functions alone using ANN. The rule generation mechanism monitors the system response over a period of time to evaluate new fuzzy rules. Non-redundant rules are appended to the existing rule base during tuning cycles. One ANN classifies the temporal response of the system into various patterns such as oscillatory behavior, overshoot in response, or steady state error, etc. This is used to determine how to tune the membership function. Another ANN identifies the controlled system, hence acting as a reference model. This model can be used to determine the stability of the new rules generated before applying to the real system.

A new approach that is rapidly gaining interest is to create special architectures out of standard feedforward ANN, also known as FNN, which can be interpreted as fuzzy controllers. The membership functions and sets of rules are constructed from example data using multi-step procedures that involve learning the membership functions and weights, forming rule representations and constructing computational networks. This readily solves two problems of conventional fuzzy controller design:

1. Lack of systematic design for rules and membership functions because it is basically a heuristic approach;
2. Lack of adaptability for possible changes in the reasoning environment.

Based on the brief overview of the development of FNN, most learning of the membership function and the weights, is done using the back-propagation method [Jang et al 1997; Nie & Linkens 1993; Lin 1994; Juang & Lin 1998; Kim & Kasabov 1999] and least squared methods [Gao & Er 2003]. Number of the fuzzy rules of a FNN controller is normally selected a priori according to expert knowledge. However, such a predefinition in FNN structure is found quite insufficient in many cases, because this reduces the flexibility and numerical processing capability of the controller and may result in redundant or inefficient computation. Many methods have been proposed to generate and prune the fuzzy rules systematically and automatically. Research work [Bossley et al 1995] on parsimonious construction algorithm for linear-parameter ANN such as RBF-NN, B-spline and Bayesian network, overcomes the *curse of dimensionality* associated with FNN structure learning. Nie and Linkens suggested a self-organizing learning scheme for structure learning. Some hierarchically self-organizing paradigm was developed thereafter. The premise structure in [Juang & Lin 1998] is determined by clustering the input space (i.e. so-called aligned clustering-based paradigm) via on-line unsupervised learning. As to the consequent part, only a singleton value selected by a clustering method is assigned to each rule initially. Afterwards, some additional significant terms (input variables) selected via a projection-based correlation measure for each rule will be added to the consequent part incrementally as learning progresses. Simple and effective paradigms whereby the structure is identified by input-output pairs rather than by only input data were presented in [Kim & Kasabov 1999]. These approaches are essentially optimal clustering-based partitioning which are only applicable for offline learning. A tree structure, which is efficient in coping with high-dimension partitioning, was presented to configure FNN in [Sun 1994]. An iterative construction algorithm called adaptive spline modeling of observation data (ASMOD) was developed in [Kavli 1993], in which structure is adapted to the data by means of an incremental refinement procedure. Each step of the incremental model refinement consists of one of three alternative actions, namely univariate addition, tensor multiplication and knot insertion. This paradigm is fundamentally a tree structure. Gao and Er (2003) proposed two criteria (i.e. system error and rule completeness) to generate fuzzy rules, and an error reduction ratio concept was utilized to prune redundant rules. Other rule base determination methods include similarity/degree measures [Kavli 1993; Chao et al 1996], and various integrated rule generation and pruning methods [Kim & Kasabov 1999], etc.

3.14.4. Fuzzy-GA Approaches

Genetic algorithms are robust optimization routines modeled after the mechanics of Darwinian Theory natural evolution [Goldberg 1989]. GA does not require gradient evaluation; hence they are applicable to solve a great range of optimization problems including determination of optimal parameters of a fuzzy logic rule-set. GA have demonstrated the coding ability to represent parameters of fuzzy knowledge domains such as fuzzy rule sets and membership functions in a genetic structure, and hence are applicable to optimization of fuzzy rule sets.

To understand the actual mechanism of GA, one may begin with its three most commonly used operators, namely: reproduction, crossover, and mutation. A member of a given population which has a higher fitness is given a higher chance to reproduce identical replicas of itself in an intermediate population. In this fashion the optimization routine facilitates reproduction of higher fit individuals and hampers the reproduction of lower fit individuals. After reproduction, crossover randomly *mates* two individuals from an intermediate population and crates offsprings which are made up of a random combination of their parent's genetic code. For each generation, the process of crossover is repeated for all individuals in the population. The population size is often a constant equal to the number of individuals in initial population. The operations of reproduction and crossover create an environment where every generation benefits from the best genetic codes of the previous generations. However, if the building blocks for the optimal genetic structure are not in the initial population these two genetic operators will be unable to find it. The last genetic operator, mutation, randomly mutates one or more of the values in the individual's genetic code in order to create diversity. The mutation operator allows for exploring new structures hence allowing the genetic optimization routine to invent new solution and finally locate the optimal solution even though the individuals in the initial population may not have contained the building blocks for the optimal solution [Linkens & Nyongesi 1996].

When applying GA to optimization of fuzzy rule sets, several questions arise. First is the design of the transformation function (in interpretation function) between the fuzzy knowledge domain (phenotype) and the GA coded domain (genotype). This is perhaps the most crucial stage of GA design and can significantly degrade the algorithm's performance if a poor or redundant set of parameters are chosen for a given optimization problem. Two important general categories of fuzzy expert knowledge consist of domain knowledge and meta knowledge. Meta knowledge is the knowledge used in evaluating rules such as fuzzification, rule inference and defuzzification methods. Most of recent research concentrate on optimize parameters of the domain knowledge, which consists of:

- Membership function: general shape (triangular, trapezoidal, sigmoidal, Gaussian, etc.) and defined points (center, max right, min left, etc)
- Rule base: fuzzy associative memory, disjunctive (OR) and conjunctive (AND) operations among antecedents in the rule base.

There are three common approaches for integrating fuzzy-GA applications. In one case, linguistic rules of conventional fuzzy controller are fixed and their membership functions are optimized [Akbarzadeh et al 2000, Meredith et al 1993, Karr & Gentry 1993, Park et al 1994]. In the second case, the membership functions of specified linguistic values are fixed and the GA is used to determine an optimal set of rules for the application [Karr et al 1991]. There is a third approach which combines both and in which rules and membership functions are adjusted simultaneously [Lee & Takegi 1993, Linkens & Nyongesa 1994]. The interdependence between linguistic variables and their membership functions suggests that both of them should be adjusted to achieve superior performance. Thus, the membership functions are adjusted individually for each rule and it has been shown that this approach can achieve better results than the other methods. However, such coding can have several potential difficulties. In such situation, in addition to the level of complexity and large number of optimization parameter, the problem of competing conventions may arise and the landscape may unnecessarily become multi-modal. This is an important problem since there are often several or many fuzzy rule sets which can represent a given nonlinear function. This means that there are more than one optimal solution to a given optimization problem which raises the issue of multi-modality for fuzzy logic systems. The associated drawback with this type of application is that the ability to interpret and explain the behavior of the fuzzy controller may subsequently be lost, since the membership functions are no longer associated with any one linguistic name. When designing the interpretation function, therefore, the coding needs to contain fewest possible parameters to avoid the problem of dual representation, and yet the coding needs to have enough complexity to contain all possible optimal or near optimal solutions. Evolutionary approaches such as Niched GA [Chambers 1995] are designed to search in complex multi-modal landscapes.

Genetic algorithms, and for that matter any stochastic search techniques, are not suited to real-time control problems because they take a long time to converge and may also inflict severe consequences on the process. However, there have a number of studies which have sought to gain the advantages of GA in this particular application. It has been suggested that when applied within the structure suitably stable controllers, such as a fuzzy controller, it is possible to apply GA to real-time control problems [Linkens & Nyongesa 1993]. In these studies reinforcement learning is used to adapt the rules and membership functions. A more general possibility for real-time application of genetic algorithms is provided by a machine learning system called a classifier system [Linkens & Nyongesa 1994]. A fuzzy classifier system learns by creating fuzzy rules which relate input variables to internal or output variables and, through performance assessment, credit assignment and reinforcement is able to achieve online improvement. A genetic algorithm is used to evolve the system [Akbarzadeh et al 2000].

3.14.5. Fuzzy-GP Approaches

Like GA, GP computationally simulates the Darwinian evolution process by applying fitness-based selection and genetic operators to a population of individuals. However, in this case each individual represents a computer program of a given programming language, and is a candidate solution to a particular problem. The programs are structured as hierarchical compositions of functions (in the set *F*) and terminals (function arguments in a set *T*). These individuals participate in a probabilistic evolutionary process wherein the population evolves over time in response to selective pressure induced by the relative fitness of the individuals for solving the problem.

For the purpose of evolving fuzzy rule base, the search space is contained in the set of all possible rule bases that can be composed recursively from *F* and *T*. The set, *F*, consists of components of the generic if-then rule and common fuzzy logic connectives, i.e. fuzzy intersection, rule inference, consequents, etc. The set, *T*, is made up of the input and output linguistic variables and the corresponding membership functions associated with the problem. A rule base that could potentially evolve from *F* and *T* can be expressed as a tree data structure with symbolic elements of *F* occupying internal nodes, and symbolic elements of *T* as leaf nodes of the tree. This tree structure of symbolic elements is the rule bases in the initial population are randomly created, but descendant populations are created primarily by reproduction and crossover operations on rule base tree structures [Akberzadeh et al 2000].

The symbolic data processing done by GP makes it particularly amenable to automatic evolution of fuzzy rules, which are comprised of symbols representing fuzzy sets and fuzzy inference. In the GP approach to evolution of fuzzy rule bases, the same fuzzy linguistic terms and operators that comprise the genes and chromosome of the rule base persist in the phenotype. Thus, the use of GP allows direct manipulation of the actual linguistic rule representation of fuzzy rule based systems. Furthermore, the dynamic variability of the tree representation permits populations with rule bases of various sizes and different numbers of rules. This enhances population diversity which is important for the success of a GP system, and any evolutionary algorithm for that matter. For example, Akbarzadeh *et al.* [2000] have applied GP to evolve fuzzy rule bases that are used to coordinate multiple fuzzy behaviors arranged in a hierarchical structure. This hierarchy of fuzzy behaviors has been employed for autonomous control of mobile robot.

3.14.6. Fuzzy-ANN-GA Approaches

The motivation for the integration of fuzzy control, ANN and GA is to enable the system to take advantage of the strengths of each methodology. Specifically, this should make use of the strong knowledge representation facilities inherent in fuzzy systems and the learning ability of neural and genetic systems. To take advantage of the representational capacity of fuzzy knowledge-based systems, and the learning ability of ANN, the structure of a fuzzy-ANN system should take the form of a fuzzified connectionist network, such as discussed earlier. The use of locally tuned type of ANN, such as RBF is therefore desirable. These networks respond to a narrow range of input and are hence suitable for faster learning. Thus the system is an ANN whose operations are carried out using fuzzy logic information processing. Such a structure provides for the application or the well known ANN learning algorithms and retains the rule-based structure of fuzzy controllers. On the other hand, it is not always possible to attach linguistic connotations to rules in such systems, which may have an if-then format without relating to any linguistic labels. The use of GA in fuzzy-ANN control has also been shown to enhance the performance of the system by overcoming some common bottlenecks, especially those requiring human intervention. In traditional neural applications the designer is usually responsible for determining the number of neurons and layers, the type of learning laws and transfer functions, and the learning rate in addition to the network parameters. Similarly, in a fuzzy system the designer has to specify the number of rules and their membership functions. GA can overcome these difficulties, for example, through being used for synthesis of optimal network structures or to generate optimal network parameters [Feldman 1993].

Learning in real-time situations is achievable more rapidly and efficiently by partitioning the input space in response to process data and using local learning rules. Knowledge acquisition and learning, on the other hand, is related to the form of available information. Unsupervised learning is suitable for classification or partitioning of online data, however, it provides no evaluation feedback or any indication of the correct response to a situation. Reinforcement learning is designed to learn the correct responses at particular instances using imprecise evaluation of the actions taken previously, through a process of credit assignment. It is usually the only alternative when precise training data cannot be obtained. One of the common problems with this learning approach is that it is generally associated with long periods of experimentation to converge to satisfactory performances, but can be alleviated by applying any *a priori* knowledge available. For example, an initial structure of a fuzzy controller can be assumed and reinforcement learning used to refine the rules and membership functions. A further difficulty is that in fuzzy systems rule spaces are often overlapping, which can present a dilemma to a credit assignment scheme when the rules happen to be antagonistic.

There are great opportunities for future integration of fuzzy, neural and genetic techniques. GA is powerful for offline synthesis of optimal structures and parameters of FNN. No plausible approaches have been suggested for their online applications. A generalized framework for integrating these techniques in a real-time control system combines the methodologies selectively to deal with specific design problems. Thus, a FIS is a connectionist network with fuzzified operations and functions; Genetic algorithms augment neural techniques to evolve optimal inference systems; unsupervised learning, for example, is used to partition the input space using feature mapping so as to enable local learning, while evaluation of competing strategies can be carried out using reinforcement learning on an online model of the process. Such a model can be obtained using supervised learning techniques with online data [Linkens & Nyongesa 1996].

3.15. GROUP BEHAVIOUR CONTROL

Communication is a mode of information transfer which in animal populations involves information transfer between individual organisms. Biological communications extensively utilises both optical and radio frequency signals for both communications and navigation. In biological communications systems, evolution has favoured virtually exclusively frequency modulation for analogue signals and pulse code modulation for digital signals, eg. sinusoidal wave emitting electric fish encode through FM while pulse emitting electric fish encode through PCM (pulse code modulation) [Richards 1985]. Furthermore, frequency hopping is used extensively in biological communications to minimise the information transfer to predators.

Animal communication may be regarded as manipulation of the receiver agent by the sender agent. All human language according to the Chomskian paradigm comprises a three layered structure – syntax which determines the rules for the construction of language structure, semantics which maps the vocabulary to objects (nouns) and events (verbs) of the world, and pragmatics which determines the function of linguistic communication in altering the behaviour of the listener – it is this last through which the capacity for language is believed to have evolved. We do not believe that there is significant amounts to be learned in considering animal and human communication further in terms of biomimicry. Much of the relevant context may be subsumed within 6500 in consideration of learning. We therefore suggest that beyond the cursory outline given above, we do not consider it further.

Communicative behaviour in animals is usually associated with courtship and aggression in which they emit information about their state. They usually involve bizarre behaviour patterns used as signals which reflect the animal's internal state. Precursors to such signals were characteristic of motivational conflict, eg. dithering due to incomplete or alternating intentional movements. These ancestral movements became ritualised through evolution as signals. Ritualisation involves standardisation of behaviour patterns with an emphasis on exaggeration, repetition and stereotypicism of the signals. This ensured that the recipient detects the signal with the minimum possibility of misinterpretation. However, contests between animals may be viewed as games in which each combatant attempts to maximise its payoffs. Individual interests do not coincide as each individual will follow a course of action that maximises its own self-interest at the expense of others. Ritualised signals increase ambiguity in that high stereotypical signals conceal the animal's real internal state. Communication is the process by which one animal exploits and manipulates another to its own benefit. Signals evolve because animals gain benefit through using give-away cues to predict the behaviour of other animals. In a contest, the ability to predict opponent's next move to allow appropriate evasive actions. Movements evolved into signals that predict future courses of action of individuals. The advantage of anticipating future reactions of individuals provides a selection pressure. Reactors would gain advantage in predicting the actor's future behaviour by reading the characteristic cues while the actor would gain advantage by concealing those cues. In co-evolution, the signal is ritualised to make the signal stronger, repetitive and more conspicuous but it often bears no objective relation to its function.

Animal communication may thus be regarded as manipulation of the receiver agent by the sender agent. An animal, humans in particular, use communication as the means by which to manipulate other agents. It is energetically more efficient to persuade another animal to perform an action than it is to perform it oneself – communication is the means by which an animal utilises the muscle power of another animal through manipulation of its nervous system. In animal communication, a distinction is made between signals, cues and signs. Signals are temporary sources of information concerning the state of the environment, such as vocalisations or emotional displays. Cues are permanent sources of information of the state of the environment such as visual adornments. Signs however provide information about characteristics of the individual. Signals are semantic if they provide listeners with information about objects and events in the environment. Signals are representational if a mental image of the object or event is elicited in the listener. Signals are referential if they are associated with objects and events in the world. Some animal species communicate about the location of food sources (such as honeybees), about the detection of predators (such as monkeys), or most commonly in social interactions [Maestripieri 1977]. In humans in particular, language evolved to establish, maintain and alter social relationships. Vocalisation in particular is common in arboreal lifestyles suggesting that vocalisation is an adaptation to poor visibility environments.

Natural language is the means of information coding for the transmission of complex information – human communication is a social adaptation which typically involves more than one modality. The phonetic function of human language is to provide vocal symbols. Consonants interspersed by vowel sounds increase the variety and hence information content of words and are phonetically universal. Complications occur when intonation alters the pitch at a syllable level – indeed, half of the languages in the world use pitch differences to differentiate words, eg. Mandarin-Chinese. Languages such as English use pitch to convey emotional rather than lexical meaning. The human vocal range allows around 85 phonemes and all human languages comprise of a limited sequence of 55 phonemes. English comprises 38 phonemes (16 vowels and 22 consonants) and words on average have ~10 phonemes per word. A man's vocal cords vibrate at ~100-150 Hz compared with a woman's pitch of 200-300 Hz. The tongue acts to block or constrict the airstream in forming consonants, but vowels are formed through an unrestricted airstream and are determined by variations in tongue position. Vowels form the central component of syllables while consonants typically mark the beginning and end bounds of the syllable. The use of digital syllables is how human and animal communication differs. While animal communication comprises of a fixed repertoire of invariant calls, human communication is based on syllables that are meaningless per se but can be used generatively and combined into an infinite number of meaningful sentences. A set of words covering a certain domain in one language is unlikely to correspond to those in another language. Word meanings are not just made up of a stock of basic components by classification since the components form only a small part of the overall meaning, eg. opposites may not be absolute but relative. For instance, the double negative does not correspond precisely to positive assertions as would be the case in logic. Complications can also arise when some words possess the same phonetic properties but differ in meaning, eg. homophones (different spelling). Speech contains more modes of symbol transmission than the written word (pitch, volume, pausing, etc). Intonation through pitch of the voice is used to convey attitude (in written language, punctuation provides this function). Intonation patterns in English also provide the basis for humour. Languages such as Chinese use voice pitch patterns for lexical functions in distinguishing different meanings of a word. One problem in speech recognition by machine is that humans use knowledge, contextual information and environmental cues to decode acoustically ambiguous and incomplete sounds. Meaning is ascribed by the surrounding context of words in that it is dependent on the relationship with other words in the sentence. This knowledge comes from numerous sources: semantic, syntactic, acoustic. Acoustic knowledge refers to the speech sounds themselves such as frequency. Syntactic knowledge comprises the grammatical rules that specify legal sequences and constructions, eg. an adjective is a qualification to a noun. Semantic knowledge is general knowledge about the domain of discourse. Syntactic knowledge alone helps to resolve ~30% of acoustic ambiguities in normal conversation, eg. homophones (words which sound alike but which are spelt differently) are often distinguishable at the syntactic level. Often it is difficult to decide whether the syntactic subject of a sentence refers to the agent or the instrument since the syntax may be identical in both cases. For sentence understanding, non-acoustic information is needed and this knowledge comprises a world model of syntactic and semantic information.Human communication also uses non-linguistic visual and vocal signals, which typically assign emotional context to language, but these are limited in repertoire. Much of language is coordinated and refined by nonverbal signals for communicating emotions and motivation in particular to the receiver(s). For body language illustrates the progression of communication by signalling attitude and motivation of the speaker. Facial expressions provide continuous emotional feedback for fine-tuning of communications. Eyes especially dominate in conversation with eye contact being maintained between 25-75% of the time. Gaze shifts are used to indicate start/finish of talking. Gestures accompany speech as markers to give emphasis to verbal conversation to visually illustrate words with emotional context. Expressions of emotional state are often innate and universal across cultures as they are often physiological reactions based on social relationships between the communicators. Although gestures are learned culturally, non-verbal signals are universal across cultures. Encoding these emotions and attitudes is mainly through facial expression and voice tone, but other factors are usually involved. Facial expression is regulated by some 20 muscles controlled from the motor cortex of the brain. As alluded to earlier, the voice is modulated by emotion through speed, pitch, etc. Gaze, touch and proximity determine social relations. Gestures are often used to illustrate utterances visually (eg. Shape, size, direction, etc), Head gestures offer yes/no responses with no being the shaking of the head in most of Europe. Gaze and tonal contour often provides feedback signals to synchronise conversational cues.

It has been stressed that visual input is an important component in language understanding – the image of the face provides information on the speaker's emotional state and motivations. Indeed, lip-reading deaf people watch the whole face [Massaro & Stork 1998]. The human brain, specifically the parietal region, combines both acoustic and visual information and that these speech modalities are complementary with little cross-talk – acoustically differentiated speech originates in the visually inaccessible vocal tract (like nasality) while visually differentiated speech originates from visually available regions due to articulation from the lips and tongue but which is difficult to distinguish acoustically. This suggests that any sophisticated automated language understanding system must employ a video camera capable of differentiating the position and shape of the mouth and lips. Language is not independent from the rest of cognition. Language is based on cognition. The cognitive model is constrained by the nature of physical reality and social interactions. The Sapir-Whorf hypothesis suggests that human thought is mediated by language and that our perception of the world is determined by language. We do not consider communications further here but note that modulation/demodulation circuits may be applicable to the evolutionary hardware paradigm.

Wasps, ants and termites are social insects. The majority of insects in an insect colony are sterile workers which never have their own offspring. However, all insects in the colony descend from the same mother. The queen is the only reproductive member who indulges in procreation, relying on workers for food and protection of the queen and the eggs. This is a basic form of division of labour. Workers are infertile males in termite nests and infertile females in all other insect societies. Workers are, however, closely related to the offspring. Hymenopteria colony comprises of one queen who made a single mating flight to store sperm earlier in life. Unfertilised eggs become male, so males have only a single set of chromosomes from the mother. Fertilised eggs become female, so they have a double set of chromosomes, one each from the mother and father. All males are genetically identical obtaining all his genes from the mother. Females share the same father and obtains all her genes from the identical father and the mother. However, the mother donates only half her genes to the offspring. Hence, male offspring share half their genes with the mother but sisters share three-quarters of their genes. Henbce, female hymnoptera are more closely related to susters than to her offspring of either sex. Females thus invest in the mother to produce sisters which are genetically closer to them than their own offspring. The sex ratio of females to males is this 3:1. Some ant species cultivate fungus gardens in their underground chambers for food. They also domesticate aphids for milking of nectar in return for protection against predators. The naked mole rat is the mammalian equivalent of social insects which live in underground colonies of burrows in East Africa. Only one female of the colony breeds having mated with only two or three males. All other members of one colony of both sexes are infertile and behave as workers.

Genetic evolution acts over generations of populations. Fundamentally, there is a random element at the genomic level introduced by mutation – such mutation is a necessary consequence of the second law of thermodynamics which suggests that no copying process can exhibit 100% fidelity. Selection however, acts at the level of the phenotype which is not random (in general, though there is an inevitable significant random element to the function of life). Sexual reproduction is commonly associated with multicellular organisms in which two separate genomes are mixed, affording a greater rapidity of evolutionary capitalisation than that afforded in asexual reproduction. However, as in much of multicellular processes (such as the forerunner to multicellularity in biolfilm colonies), sexual reproduction has its precursor in bacteria which often exchange sections of genetic material, i.e. plasmids. Sexual reproduction involves merging of components parts of genetic material from the parents through cross-over to create entirely new individuals. The prevailing rationale for the extended survival of this process is the protection it provides in resisting parasites by presenting new genomes with immunological novelty [Hamilton 1998]. Sexual reproduction provides organisms with long generation times the immunological novelty to resist invasion by short generational parasites which evolve rapidly. Hence, sexual reproduction provides for the evolution of novel solutions in offspring based on the previously successful solutions of their parents in response to environmental changes. Unsuccessful solutions will not procreate. Each generation is effectively optimised for survival *within the genetic resources available* [Ng & Leng 2002]. It is important to note this last point as evolution does not provide pure optimality, merely optimality within the inherited constraints (witness the mammalian eye in comparison to the octopus eye). There are three main types of evolutionary algorithm of interest – the genetic algorithm, genetic programs and classifier systems.

Sociality is the basis of intelligence in humans - sociality complicates the environment, so that survival depends on predicting the behaviour of others in the social group. The best model for achieving this prediction is by self-inspection to understand the motivation of others. Language evolved as a form of social manipulation. Soviality evolved despite individual selfishness. Mutual grooming is common in both birds and mammals. The benefits to social grouping include reduction in the probability of predation. Alarm calls are altruistic acts though they are selfish. The key is the ability to recognise and remember individuals for the evolution of reciprocal altruism. Sociality provides the basis for cultural transmission of behaviour but it is complex. Cultural transmission is analogous to genetic transmission but it evolves much faster. The unit of cultural transmission is the "meme" which propagates from mind to mind through language. The mechanism for meme survival is its psychological appeal. Meme transmission is subject to mutation and merging with other memes. They are in competition with each other to dominate attention of the human brain. Many human establishments are sets of co-adapted memes – church, law, art, philosophy, science, etc. If successful, such memes can influence human culture and survive genes. Indeed, memes can be self-destructive to the individual as well as self-perpetuating in themselves, eg. suicide. Culture generally refers to learned behaviour through imitation or explicit teaching which may be applied to many animals, eg. English blue tits, macaque monkeys, chimpanzees, etc. Culture in animals categorically excludes extended phenotypes which are inherited complex behaviours, eg. nest building in bower birds, dam building in beavers. In human terms, culture is generally understood in reference to technology (technology being defined in its broadest terms) mediated by language. There is little doubt that language as a symbolic means of communication is a powerful medium for the transfer of cultural information. The origins of human language are shrouded in mystery, but recent work in combined genetic and linguistic analysis suggests that fully modern syntactic language dates back to around 50,000 years ago coinciding with the explosive technological advances of the Cro-Magnon culture. However, rudimentary language must have existed earlier, and indeed, it appears that *Homo erectus* may have possessed a significant Broca's area (implicated in human syntactic processing). All human language according to the Chomskian paradigm comprises a three layered structure – syntax which determines the rules for the construction of language structure, semantics which maps the vocabulary to objects (nouns) and events (verbs) of the world, and pragmatics which determines the function of linguistic communication in altering the behaviour of the listener. It was through pragmatics – the effects of language within the social group – that the capacity of language must have evolved. Hence, the capacity for language is genetically inherited but the use of language creates culture. Post-Neolithic humans have taken language to its next level through the cultural invention of writing – a permanent codification of spoken language. Although cultural transmission is dominated by vertical mediation between parents and offspring, there is also a lateral component of behavioural consistency between peers in a social group. Cultural transmission of behaviour is Lamarkian in that acquired characteristics are communicated rather than inherited characteristics – this makes it rapid and accumulative within generations. It provides greater adaptivity to the environment (this environment is critically dominated by the social rather than physical environment). Memetics is a new discipline which is attempting to apply mathematical techniques borrowed from genetics to describe the behaviour of cultural transmission. Hence, vertical transmission of culture resembles traditional genetic transmission while horizontal transmission of culture resembles epimediological spread. The Baldwin effect provides the basis for inheritance of learning capacity though it is important to note that the Baldwin effect does not contravene Darwinian evolution and the fundamental paradigm of molecular biology (DNA-RNA-proteins).

Jennings & Compos (1997) suggested that decision-making agents should maximise their individual utility and also thise of the overall system so that their input does not diminish the performance of the overall system. This is the concept of social responsibility in constraining the problem-solving of individual agents. The principle of social rationality maximises the joint benefit to its agent members and the society of agents – if an agent member of an effective society can perform an action whose joint benefit is greater than its joint loss, it may select that action. Joint benefit represents a balance between the individual agent and the social community of agents. Members indulge in activities for which their member benefit is less than their member loss if the member benefit if society gains more in total than it loses. Members indulge in activities which provide personal benefit but which are detrimental to society if the member benefit is greater than the social loss. This places a limit on the degree of cooperation – a member will not indulge in activity which brings a personal loss greater than the benefit accrued to society. Rights indicate ways in which members are permitted to exploit social resources while duties are restrictions imposed on agents as members of that society. Social membership provides certain rights but at a cost of certain duties – this reciprocation is the basis of social commitment. Responsible agents which maximise joint benefits consistently perform better than selfish/selfless agents.Distributed decision-making in teams have five major components [Ho 1981]:

3. a vector of random variables $\xi = \xi_1 \ldots \xi_m$ with a distribution $p(\xi)$ representing the uncertainty of nature;
4. a set of observations (information) $z = z_1 \ldots z_n$ which are functions of ξ: $z_i = \eta_i(\xi_1 \ldots \xi_m)$ where η_i = information structure;
5. a set of decision variables $u = u_1 \ldots u_n$;
6. strategy (decision rule/control law) of the decision to map $\gamma_i : z_i \rightarrow u_i$ i.e.. $u_i = \gamma_i(z_i)$;
7. payoff criterion (loss), $L = (u_1 \ldots u_n, \xi_1 \ldots \xi_m) = L(\ldots u_c = \gamma_i(\eta_i(\xi)), \ldots, \ldots \xi_j, \ldots)$

The problem is to find γ_i that minimises the expected payoff functional $J =< L(u = \gamma(\eta(\xi)), \xi) >$. There is different but correlated information concerning the uncertainty available to each decision-maker and the need to coordinate with other decision makers to realise the payoff. The control variable that must be extended is the design of the information structure η - this essentially involves reducing the uncertainty by improving knowledge. Highly structured information enables minimisation of the payoff function.

Distributed information processing is fundamental and ubiquitous in natural and social systems as a form of complexity reduction through decomposition into modular chunks while providing redundancy for graceful degradation in the event of component failures [Chandresakaran 1981]. No single sequential processor can overcome complexity problems, but distributed systems with two-way information exchange can provide some measure of solution. No individual has complete knowledge of any complex task. Indeed distributed computing is in fact an example of Minsky's concept of "a society of mind". This notion involves symbol processing arising from a collection of primitive computational elements (neurons) at a subcognitive level, i.e.. an emergent epiphenomenon resulting from complex self-organisation. Cognitive behaviours emerge as a statistical property of distributed neural networks. There are no overall rules of thought deterministically manipulating symbols. Federations of neural assemblies implement a community of programs. Each of these special purpose networks is linked via a communications network linking various neural regions. Each program represents an independent agent, forming an interactional network. Intelligence arises from this society of agents by the interactions of these conflicting and competing agent-based processes without overall global control.

An example of distributed computing is the Worm program. Worm programs are computations which span over individual module boundaries [Shoch & Hupp 1982]. Programs in individual computing elements comprise segments of the worm. They remain in communication with each other - if a segment fails other pieces search for another machine. Worms have several components: an initialisation code runs the main program. Worms allow the dynamic allocation of resources from machine to machine. Loading the operating system initialises the worm. Worms can be used to load diagnostic test programs in distributed systems. By incorporating an independent self-destruct timer segments may be "killed" off to overcome the problem of uncontrolled replication from machine to machine.

Multiple agents that are the desired form of modular architecture are characterised by [Jennings & Wooldridge 1996]:

- autonomy without the interaction of humans;
- social ability so that agents interact with each other;
- responsiveness whereby agents perceive their environments and respond to it;
- proactiveness by exhibiting goal-directed and opportunistic behaviour.

Distributed artificial intelligence (DAI) is based on the concept of organising a heterarchical team of interacting decision makers who indulge in coordinated activity operating under definite strategies to generate and develop solutions to problems which cannot be solved by a single individual. DAI appears to offer advantages for the solution to wicked problems that are characterised by [Partridge 1981]:

- no definitive formulation;
- no halting rule;
- no absolute solution;
- each possible solution is unique.

Real world interactions are complex generating wicked problems as opposed to well-formed problems. The intelligent agent must use strategies to obtain feedback from the real world interactions in the minimum amount of time. Central to this is the model which represents the state of the relevant aspects of the world which may be used to generate consequences and so evaluate actions. To cope with (i), the problem is broken down into subproblems by imposing restrictions and assumptions to facilitate formalisation. To cope with (ii) and (iii) a referee is included to make decisions to terminate processes which lead to circularity, infinite regression, and deadlock. Property (iv) implies the non-existence of a widely applicable generalised structure and relies on the ability to make online discoveries during the problem solving process to aid in the selection of a possible solution. The DAI paradigm effectively partitions the planning or problem-solving search space into smaller localised search spaces. The amount of cooperation between agents (usually knowledge bases) can vary from fully cooperative to conflicting but most systems are cooperative to varying degrees. Klein (1991) considered competitive conflict situations where each node has its own benefit superordinate to any globally optimal goals. The central purpose in conflict resolution is "if two plans for achieving two different agents' goals conflict, then find an alternative way of achieving one goal that does not conflict with the other agent's plan for achieving its goal". This requires specialist heuristics for compromise. However Vamos (1983) suggested that cooperation between agents through information exchange as an essential condition for DAI. Distributed expert systems attempt to solve problems that cannot be decomposed into independent subproblems which can be solved in isolation by virtue of cooperative.

Natural complex systems are composed of basic single components which in large numbers act collectively together to perform complex tasks in a robust manner [Wolfram 1986]. This robustness is afforded by irreversible dissipation where only a few major features in behaviour are retained over time. Such phenomena have been investigated in chemical systems. Nonlinearities arising in complex spatial and temporal patterns are created in a chemical medium through the production of entropy. This allows dynamic structure-formation in the form of chemical gradients. Complex systems are often characterised by several dynamic basins of attraction which correspond to potential minima. Complex systems are often computationally irreducible such that their behaviour is in general not predictable (ie. NP hard). However, only approximate behaviour may be required. Minsky (1980) encapsulated these ideas through K-lines which connect to modular mental agents that are partially autonomous forming a society of such units as the basis of mental states. The society of minds hypothesis views intelligence as the product of numerous, diverse agents which interact nonlinearly to produce mental activity [Minsky 1991]. The nonlinearity in a large interacting system implies the impossibility of prediction without knowledge of the system's history. Each agent is a specialist and has two states – on or off. Total mental state is a specification of all agents that are active. K-lines represent the key to partial mental states and form into dynamically changing hierarchical nets of interacting partial mental states. Each K-line represents a relation among the connecting nodes and so the system forms a dynamic semantic hierarchical network. Others have expressed similar ideas that intelligence arises as the result of components (schemas) which trigger their activity states interactively determined by the plastic structure which in turn is affected by the activity [Winograd 1980]. Dependence on the past history of activity states allows "nonlogical" reasoning. Complex systems comprised of many interacting parts is fundamental to entropy-reducing self-organising capability under far from equilibrium conditions [Kiehn 1979]. This model of human cognition assumes that there are distributed multiple hardwired specialist units or modules which are invoked in parallel. Current activity is represented as a global working memory (blackboard) with a virtual machine implementation of the agent-based interaction. Cooperative self-organisation emerges spontaneously from phase transitions in many DOF via non-equilibrium dissipative dynamic convergence to attractor states of its components in interaction [Shoner & Kelso 1988].

Knowledge bases do not necessarily require global consistency so truth maintenance may be implemented at individual knowledge base level [Lenat 1990]. Cooperative problem solving exhibits the properties of "satisficing" problem solving in that it generates useful and acceptable results using a reasonable amount of processing rather than using extensive processing to generate optimal solutions. This overcomes the limit of bounded rationality in that each individual component is limited in problem solving capability. This is essentially an economic concept in that economics is concerned with the rational allocation of scarce resources among competing users [Simon 1978]. Computational cost increases search due to limited computational capacity with respect to the expected improvement in planning alternatives. In satisficing, search terminates when available solutions exceed the expected improvement. Complexity is a problem that can only be solved using finite resources by satisficing, i.e.. iterative approximations by local optima. Satisficing is optimal in that it trades off performance with cost. DAI systems are differentiated according to their control methods of organisation and their communications protocols [Decker 1987]. The distributed paradigm is also taking over from relational databases [Hsaio & Kennal 1989]. Both control and communication are distributed whereby agents acquire and disseminate knowledge as the computation evolves.

The amount of cooperation between agents can vary from fully cooperative (with high communication cost) to non-cooperative (no communication cost). Cooperation reduces control uncertasinty and can coordinate multiple expert systems with different knowledge domains. Cooperation introduces a communications overhead but it also helps to reduce uncertainty and incompleteness. Information exchange is necessary to allow the system as a whole to produce global responses but it is important that the majority of time is used in computation rather than communication. Protocols keep communication between components to a minimum which may vary from a shared global memory (blackboard) in a hierarchical network structure to message passing whereby agents as individual active entities pass messages between each other to communicate. Synchronisation may be achieved through a protocol such as the contract net protocol. Each agent may be an independent specialist expert with different knowledge using cooperative methods to minimise communications yet retaining multiply coherent behaviour to solve common problems. Each agent has limited expertise and resources in that no one of them has sufficient information to solve the entire problem and by cooperating they can increase the speed of problem solving. Other possibilities include communities of experts characterised not only by their knowledge but also by their problem solving techniques, eg. mixtures of rule based experts, frame based retrieval systems and model based causal reasoning systems. Communications between processors may imply message passing, marker passing or value passing [Wah et al 1989]. Message passing which is employed by distributed processes involve high communication costs, sharing of resources and synchronisation. Marker passing involves the transfer of single bit markers between processors which indicate the presence of a given property. Each processor is simple and can store a few distinct marker bits. If more than one marker arrives at the same destination they are ORed as the basic inference operation. Semantic networks use marker passing (eg. the Connection Machine). Value passing involves information passed as continuous quantities and only simple operations can be performed. If several values arrive at a destination simultaneously they are combined into a single output by a mathematical function. This is the system used in neural networks such as the Boltzmann machine. Synchronisation is not required for marker passing or value passing systems. Global coherence is important so that each agent acts in accord with the common goals of the entire system. This requires some form of organisation ranging from free teams to dominance hierarchies. Teams are totally free and are data driven adopting parallel problem solving without cooperation. Hierarchical approaches use some agents for global direction and dictate lower level nodes' actions - these systems are goal driven. They break down tasks into more specialised subtasks and a manager supervises the task execution while subordinate contractors perform the task execution. However a compromise with data driven low level agents and goal driven high level agents is possible when negotiation is permitted.

Some distributed systems are organisation-based and rely on Simon's principle of bounded rationality that the human and machine information processor has limited capacity for problem solving and as the complexity of a problem increases some form of partitioning is required. This partitioning is accomplished through the division of labour and resources [Fox 1981]. Hence complex organisations are needed to solve difficult problems that are subdivided into functional units. As the size of a group increases collective decision making becomes difficult and a multilevel hierarchy is required under a single decision maker coordinating and delegating the activities of the lower levels. Each level of the hierarchy acts as an information filter at successive levels of abstraction and representation. When multiple products are created competition for scarce resources ensues. To reduce this, a multidivisional hierarchy can be effected where the organisation is divided into different product lines with strategic control and resource allocation occurring at the higher levels. As the organisation grows coordination becomes difficult. Adam Smith's concept of self-stabilising price may be introduced in the form of contract negotiation with marginal pricing to maximally utilise resources.

There is a possibility for dynamic organisational control structures such that the organisation can be created at the start of the problem solving process or be allowed to adapt during the problem solving process as a metalevel control implementation [Lesser 1991]. This allows the system to reorganise itself in case of failure and to a particular task. There is no general all purpose organisational structure which is ideal for all problems. Such dynamic structuring requires the metalevel controller to improve coherent behaviour between the nodes, eg. contract net protocol. Alternatively metalevel control can be attained by allowing an agent to accept its organisational role as a company node or reject it as an entrepreneurial node - this allows effective restructuring of the framework to adapt to new circumstances.

Durfee & Montgomery (1991) suggested that small groups of agents may combine and abstract their activities as teams which in turn may combine at various levels of abstraction to form a functional hierarchy headed by a single entity. As increased cooperation and coordination imply the need for increased communication and that communication bandwidth is limited, communications methodologies are based on protocols. In general there is no requirement for global knowledge base consistency between agents since coordination is achieved through negotiation and compromise to resolve any conflicts and come to a mutually acceptable agreement. It is through the exchange of relevant information that agreement may be achieved in the reduction of inconsistency and uncertainty. Proposals are a basic mechanism for the solution to a problem relevant to all agents concerned. An agent receiving a proposal can respond with a critique of the proposal, or a counter-proposal which amends parts of the original proposal. An explanation can be provided as additional relevant meta-information. The process continues until agreement is reached. Message coherence is characterised by their relevance to the overall problem which dictates its consistency, timeliness which dictates its effect on the recipient towards the overall solution, and the fraction of completeness which dictates the solution state towards the overall solution. There is a trade-off between these qualities: a complete solution is unlikely to be timely but may very relevant. Agent awareness of these parameters allows effective local decisions to be made on message-passing effectiveness. The paradigm for communication may be either through a shared global memory, message passing between agents or a combination. Using both overcomes the disadvantages of each and utilises the advantages of each, i.e.. using shared memory for local processes with message passing between groups of agents. The shared global memory is exemplified by the "blackboard" model which agents post messages, partial results and requests. Message passing is based on the ACTOR model in which the actors are characterised by their actions to messages and their acquaintance of other actors. Agents in a message-passing system do not directly affect one another and remain independent objects. Hence this is a form of object oriented information processing whereby actors (objects) pass messages asynchronously. Communication protocols vary from selective where messages are targeted to specific receivers, on-demand where messages are sent only on request, or broadcast where messages are available to any agent. The blackboard is a broadcast mechanism. The message-based contract net protocol (CNP) is particularly suited to a distributed set of heterogeneous specialist knowledge based agents [Smith 1980; Smith & Davies 1981]. They cooperated by mutual sharing of information by exchange between loosely coupled agents, evaluation and mutual agreement. Firstly the problem is decomposed into subproblems until the non-decomposable kernels are reached. Next each kernel is solved allowing communication and cooperation between agents. Finally all answers are synthesised to provide an integrated solution to the overall problem. Negotiation provides the communications protocol mechanism to structure agent interactions so that relevant knowledge sources exchange information and come to agreement. This allows a more fluid relationship between nodes. Essentially tasks are dynamically shared so that agents cooperate in solving subtasks of the problem and share the computational load. In the contract net protocol, a hierarchical organisation is produced dynamically through a contracting and subcontracting process. The task-enumerating manager agent divides the complete problem into subtasks in top down fashion and negotiates well defined contracts with task receiving contractor agents. No agent is designated a priori as both roles of contractor and manager are taken on dynamically but the process is elaborated in a top down manner. The manager advertises by broadcasting task announcements with a task description and eligibility requirements across the system. Potential contractor agents submit bids based on their suitability and resources which the manager evaluate and then selects and makes contract awards based on the bids. The contract is an explicit agreement and links the manager and the contractor. Contractors may further partition subtasks and may subcontract tasks and award subcontracts to other nodes forming a contract net between agents linked by the contracts. Results are also shared such that agents can share partial results from different perspectives of the complete problem perhaps through a blackboard. Solutions from each agent are synchronised to provide a complete solution. The contact net protocol may be goal driven when lightly loaded or resource driven when heavily loaded. More time is spent on computation than in communication. The CNP has been applied to a computerised manufacturing system (YAMS) for the control of discrete parts manufacturing.

The functionally accurate/cooperative (FA/C) approach involves asynchronous, opportunistic knowledge sources cooperatively exchanging and integrating partial and tentative high level results which may be incomplete or inconsistent being based on local information in pursuit of complete global solutions in bottom up fashion [Lesser & Corkhill 1981]. By transmitting only high level, high impact information the communications overhead between knowledge sources is reduced. Communication typically takes more time than computation and it is desirable to maximise computation and reduce communication. Agents generate partial solutions using incomplete and uncertain data which may be mutually constraining between the sets of agents to maintain consistency and resolve ambiguities. Agents iteratively exchange information during intermediate steps to reduce waiting time. Asynchronous operation reduces agent idle time and introduces greater parallelism. Multiple consistency reinforces confidence in partial solutions to constrain complete solutions in the face of incomplete or uncertain data from diverse knowledge sources. The heart of the FA/C approach lies in the consistency checking of high level exchanged results to enhance global coherence. Inconsistent partial solutions are discarded and only consistent solutions are searched for. This bears resemblance to the scientific community metaphor with reference to paradigm-consistent theories and the generation of alternate concurrent partial plan bears resemblance to rival scientific theories. Partial global plans can be constructed from individual agents' long term node plans. Partial solutions which are consistent and have high confidence can guide further low level processing at a local level. When these goals are combined the global partial plan represents the intentions of all the agents working on the problem. By iterative exchanging node and partial plans organisational restructuring can be effected. Partial global plans are effective at coordinating problem solving agents and allow reordering of local plans based on global plans. They also decrease communications traffic to cope with the limited bandwidth by transmitting only high level partial solutions rather than raw data. These partial solutions enable maximum utilisation of limited computational power. They can also provide prediction information to generate expectations which may reduce search for global solutions by imposing constraints. High level planning provides a means of focussing attention and the dynamic allocation of resources. Planning also implies the need for a world model as the simulation of the effects of planned actions on the environment. Each node has a world model that is sufficient for local plan development. Changes are communicated to other relevant node planners for revision of their world models. High level partial global planning by individual nodes allows them to reason about its effects on local activities of node groups and provide a degree of local control. Both the FA/C and the CNP approaches are mutually consistent and may be combined.

Expectation driven communication involves agents having models of other agents so that agents may cooperate without routine communication. This is tacit bargaining [MacIntosh et al 1991]. Only if unexpected events occur are communications necessary. This concept was based on the assumed rationality of each of the agents who may otherwise be antagonistic due to conflicting goals. Good models allow intelligent predictions about other agents' future activities without exchange of information and helps to reduce redundant processing. This provides a good model of goal-oriented human dialogue which minimises the communications requirements. This process works best when agents' goals are not contradictory otherwise Prisoner's Dilemma two-person non-cooperative game theory scenarios can arise (i.e.. two prisoners are punished if neither or both confess but if one confesses then the other only will be punished). Lesser (1991) discussed the need for high level plan policy nodes to ensure global coherence and coordinated problem solving by providing a guiding mechanism to individual nodes similar to a blackboard. Such plan policies are required to provide expectations for top down processing based on high level goals in conjunction with bottom up data driven agent activity. Each node may also have metalevel information which may be exchanged with other nodes' metalevel information for strategic decision-making as well as tactical decision-making. Such metalevel control allows the exchange of agent goals to allow the agents to accommodate each others' requirements in a sophisticated manner. This does introduce increased communications in that high level coordination messages must be passed in addition to the other messages and this may not be feasible. If exchangeable local metalevel structures are rejected then this returns to the concept of metalevel plan policy nodes. This implies a trend from heterarchical organisation to a hierarchical organisation since the hierarchy is capable of focussing problem solving without a heavy communications burden.

MacIntosh et al (1991) investigated general issues concerning coordination in their multiagent loosely coupled distributed automated reasoning system. Each agent had local knowledge and they cooperated concurrently with each agent using resolution theorem-proving to solve local problems. No agent had a global view and agents only communicated in making requests for additional information from other agents and in replying to requests. This reduced the time spent in communication and maximised the time spent on computation. They found that the best performance occurred when the redundancy factor as a global measure of the amount of relevant knowledge shared amongst agents was low and when the number of agents was large, i.e.. each agent was a specialist with little overlap and the network of agents was large. This is effectively a solution to conflict between agents through the division of labour as a conflict resolution strategy.

Wesson et al (1981) discussed two distributed systems of limited perspective: an anarchic committee of cooperating specialised experts with message passing between agents and a hierarchical system of agents with differing levels of perspective but without intralayer communication. They found that the anarchic committee communicated sparingly and operated faster than the hierarchical system with consistently better performance. They suggested that only when signal processing reduces the abstract data in the higher levels from the lower levels should hierarchical systems be used. Further, allowing intralayer communication at lower levels may enhance data compression for abstraction at higher levels. The anarchic committee is limited by the communication requirement. The problem may require a trade-off between the number of cooperating experts and the hierarchical structure to cope with multiple environments. The addition of a global blackboard provided a hybrid two-layer hierarchical method to give a global view and this gave the best performance. Blackboard systems are appropriate when high communication rates exist between agents and when agents may be grouped into local patterns. Message passing is suitable for widely physically distributed systems that need to be failure tolerant. It is generally preferred to use a hybrid hierarchical/distributed system which combines the accuracy of the hierarchy with the speed of the distributed system. Asynchronous computations may be performed in parallel with one processor acting as an overall monitor to resolve conflicts. Such a monitor acts at the top level of the control hierarchy as an operating system in making global decisions about the computational capabilities of module resources. A blackboard may provide the organisational structure to provide control and connectivity between agents. The blackboard can provide broadcast communications with directed message traffic between nodes provided by the CNP. In general complexity favours a heterarchical organisational structure while uncertainty favours a hierarchical organisational structure. The committee based message passing system, the blackboard and information filtering hierarchy form an evolutionary sequence towards more complex structures for more complex problems. Ideally the organisational structure should be flexible enough to suit the task at hand [Durfee et al 1989].

The blackboard model is inspired by cognitive psychology in that it essentially models aspects of human cognition: knowledge sources represent permanent long term memory while the blackboard represents short term working memory. Blackboard systems use multiple knowledge sources to analyse different aspects of a complex problem. The blackboard control architecture provides a uniform system which integrates a number of diverse, specialised and independent knowledge sources which communicate through a common global database (blackboard). The essence of the blackboard system which uses a common memory between multiple knowledge sources is that it provides the means for results sharing when subproblems cannot be solved by independent non-communicating agents. Different specialists contribute to different aspects of the plan which are posted to the blackboard and incorporated into the plan. All domain specific data (including partial solutions) are posted onto the blackboard by each knowledge source for accessibility allowing the broadcast of events to all the knowledge sources via the blackboard. A master control program examines the blackboard and schedules the component subsolutions into composite solutions. The blackboard is a high level operating system that integrates the control of the distributed component knowledge sources according to its "master" plan. An executive program controls the activation of each knowledge source in turn when an event occurs. The system goals determine the plan which specifies the tasks to perform and in which order according to priorities and time constaints. No individual agent can solve the problem alone so the blackboard acts as a shared global database to allow communication and cooperation. The agents are knowledge sources that can react to and modify the blackboard data structures. Each knowledge source comprises of procedural condition-action production rules for event driven processing. The knowledge sources are independent and autonomous yet allow simultaneous indirect cooperation amongst the knowledge sources through information in the blackboard. Each knowledge source is guided by strategic plans to provide metalevel control. The blackboard is a centralised database which allows hypothesise-and-test processes. The system maintains a set of hypotheses based on its world model and uses these to focus processing effort on expected events. Each knowledge source applies either hypotheses or tests hypotheses on the blackboard. Knowledge sources can modify posted plans based on subsequent contributions from other sources. Executive knowledge sources apply inference techniques to generate solution elements on the blackboard and knowledge sources respond to, generate, and modify solution elements on the blackboard. The blackboard system is then both data and goal driven. A set of solution elements on the blackboard constitutes a partial plan. Complementary or alternative plans can coexist and one or more partial plans can be merged to form more complete plans. Each agent develops partial interpretations and hypotheses based on their incomplete data. Solutions can be constructed by the aggregation of mutually constraining partial solutions on a blackboard using hypothesise and test strategies. Such results sharing based on different perspectives is data directed. Partial hypotheses are proposed and tested for plausibility at each stage of the processing. Results sharing facilitates solutions to problems that cannot be subdivided into subtasks. Results from one agent influences and constrains results from other nodes.

The blackboard model of an expert system comprises a conference table around which the chairman convenes a meeting of a number of specialist agents. The chairman controls the blackboard activity by managing the project agenda and work schedules. The purpose of the blackboard is to find the optimal $D_i(t)$ for all N agents such that J(t) is minimised [Silverman et al 1989]:

$$J(t) = <\theta(t)>$$

where:

$\theta_j(t)$ = *vector of states of nature*

$p_j(\theta(t))$ = *probability distribution of each state of nature*

$z_j^{obs}(t)$ = *observation vector of states of nature j by each of the i agents*

$D_i(t) = \{WP_{ij}(Z(t))\}$ = *decision variable of agent based on observations of agent*

Agents decisions are multiple conflicting $p(D_i(t))$ and $p(z_{ij}(t))$ which must be optimised as min.$J_j(t)$. Since agents' observations also depend on other agents' observations and decisions agents require information from the other agents. Distributed processing allows parallel computations of concurrent activities at different levels of granularity. Large granularity may involve multiple copies of the blackboard over several machines. Medium granularity involves the distribution of individual agents as knowledge bases. Each specialist attempts to achieve its goals according to the blackboard requirements. Problem-solving may be characterised as a search through a large problem space which exhibits exponential explosion in computational complexity. Partitioning the search space and using multiple agents to independently search each partition is one approach. However, it is possible to generate a superlinear speedup and computational increase with polynomial performance. Cooperation through a central blackboard accessible by all agents alters the search problem by introducing massive speedup over using non-cooperative agents [Clearwater et al 1991]. Hence, the blackboard is an optimal solution.

The blackboard may be partitioned into a hierarchy of levels for different representations or abstraction levels of the problem. Data may be accessed at the required level. The hierarchical structure of multiple levels enables lower levels of the hierarchy to hold data that will change more rapidly than those in the higher levels. Furthermore knowledge at higher levels may be used to predict future events and to confirm current conclusions. Successive refinement from higher to lower levels provides optimum resource utilisation. Actual events or facts are then used by all the knowledge sources to revise plans and hypotheses. At the lowest level of stimulus/response, rapid activity for reflexive behaviour is required. A network of virtual addresses denoting the agents permits agent communication. A separate domain independent goal blackboard with similar levels of abstraction as the domain dependent data blackboard improves cooperation control by allowing goal transmission between agents so enhancing coherence [Hayes Roth 1985]. The control blackboard contains data associated with strategy selection, allocation of resources and the coordination of knowledge sources. This is metaknowledge that is knowledge concerning the domain knowledge such as the reliability of data and the possibilities of events. Rather than adding more domain knowledge, more metaknowledge increases the effective use of current domain knowledge. It also provides for the formulation of belief and default reasoning. A metalevel architecture entails a fixed repertoire of control heuristics since it does not reason about control actions. It distinguishes domain actions and meta-level control actions. It is generally preferable to dynamically construct and modify situation-specific control plans from modular control heuristics and so allow reasoning about the problem solving process itself. Adaptability in the control of problem solving behaviour is characteristic of intelligent agents such that control is used to adapt planning behaviour to solve the problem situation. There are preferences encoded in control heuristics that focus on tasks or actions that enable such decisions to be made. Different control heuristics are suitable for different problem situations and changing circumstances. To maintain flexibility a least commitment strategy with partial planning is required. Several control heuristics may be operative simultaneously. Control knowledge sources respond to, generate and modify solution elements on the control blackboard via a scheduling mechanism. They are specialists in selecting the strategy to be used to complete the required mission from the current situation. The scheduling mechanism comprises 3 control knowledge sources: Update-to-do, Choose-knowledge-source and Integrate-knowledge-source. These knowledge sources manage triggering scheduling and interpreting problem solving. They create and assign weights to simultaneously conflicting or complementary heuristics to resolve them to the control blackboard. The control blackboard has 4 levels of abstraction: Problem, Strategy, Focus and Priority levels. The Problem on the blackboard initiates problem solving by triggering knowledge sources with the appropriate rules. Strategy decisions establish general sequential plans which guide the rest of the process. Focus decisions establish local problem solving objectives to execute particular rules and so influence scheduling. Priority decisions establish global scheduling criteria favouring particular rules. The blackboard system plans strategic actions sequences using the control knowledge which are then implemented as a series of focus actions. The dynamic problem situation may require interrupts at the focus level. Once the alternative actions are completed the system resumes the interrupted strategic action sequence. Some events may however require strategic replanning and so the system can back up and redo parts of the strategic plan.

Single blackboards are vulnerable to failure suggesting that a hierarchical network of blackboards which use a form of message passing may be advantageous. The blackboard architecture may implement a version of the contract net protocol. Agents may contract for work packages by sending proposals (specialist activation requests, SAR) to the chairman. If accepted work authorisation is issued, specialist work packages are then initiated. SAR's are packeted messages which identify the specialist and the work package and the information from the blackboard required to complete the planned process. The chair comprises a schedule controller to provide conflict resolution and to select work activities and agents using metalevel monitoring. Furthermore agents and work schedules may be cloned for distributed and parallel implementation. Some applications are highly parallel in nature while others are highly serial.

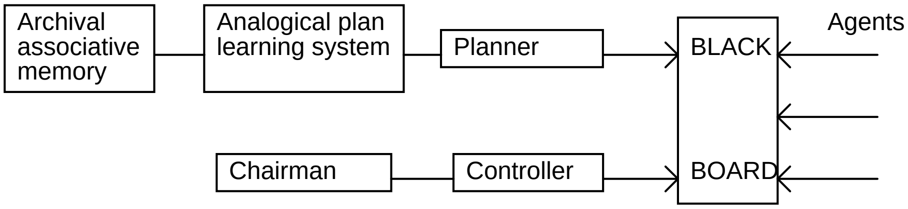

Figure 217 – Blackboard architecture

Distributed artificial intelligence includes swarm control such as bird flocking (boids) and fish shoals. Traditional approaches such as the contract net protocol based on economic payback models (similar to the bucket-brigade algorithm used in classifier systems) are not biologically-inspired. The most appropriate biological analogy for distributed intelligence is the termite colony [Deneubourg et al 1986, Franks 1989] – individual termites are simple processing units with limited behavioural repertoires but coordinated through pheromones are capable of building termite mounds of high complexity through the self-organising emergence of complex higher-order behaviour. Loosely coordinated groups can perform many tasks in parallel and offer the potential for reconfigurability of role allocations, making such techniques ideally suited to space exploration [Brooks & Flynn 1989]. Ants communicate information between individuals using paths marked out by pheromones. Initially, a single ant will move randomly until it detects a pheromone trail which it will follow. A moving ant deposits variable amounts of pheromones on the ground which creates a trail. Another ant in the vicinity will follow the trail with high probability secreting its pheromones. This will reinforce that trail, attracting more ants to it based on the strength of the pheromone trail. Generally, the trail has a limited lifetime as pheromones evaporate over time. Dorigo et al (1996) have simulated this process to solve the travelling salesman problem (TSP) on the basis that the quantity of pheromone deposited on a shorter trail is much greater than that deposited on a longer trail. The algorithm has a complexity of $O(n^3)$ for a single cycle, better than most TSP solvers. Distributed learning through imitation is an unsolved problem – for instance, path imitation involves passing through a sequence of dynamic states which are effectively hidden from the observer.

Hewitt (1985) proposed an open system of multiple processors which allows continuous evolution, decentralised decision making, inconsistency between knowledge bases, negotiation between agents or "actors", and openness to the outside world. It was similar to the scientific community metaphor but the central thrust was the open systems concept as a necessary prerequisite of intelligence. Whereas AI is based on searching a problem space comprising symbolic structures to find a path through the sequence of states from an initial state to a goal state, the open system approach involves parallel exploration beyond searching so that goals evolve with flexibility as the exploration proceeds. All human knowledge is highly inconsistent and depends highly on context and indeed all systems exhibit inconsistency due to unavoidable error (Hewitt's inconsistency conjecture). This is in fact a version of Godel's incompleteness theorem that states that all human knowledge will forever be inconsistent. Logic cannot accommodate inconsistency and so logical deductive inferencing is unsuitable for the foundation of intelligent machines due to the indeterminacy arising from the use of shared systems generating conflict through deductive indecision. Any form of local concurrent and interdependent activity will generate conflict and indeterminacy. Inconsistent beliefs are not meaningless since they are based on empirical knowledge which involves categorisation and interaction with the real world. Message passing regards the meaning of a message to be the effect it has on the subsequent behaviour of the recipient, i.e.. semantics is based on interaction through communication rather than internal closed world logic. Independent, simultaneous active objects (actors) which communicate concurrently by sending messages to each other can model any control structure and provide a means for modelling causal relations through message passing. Central to the open system was overcoming the limitations of Turing's halting problem that plagues all recursive sequential systems. PLANNER is a goal directed problem solving system that functions without an externally specified problem space. It is based on the distributed concurrent ACTOR model in which each actor is an object-oriented primitive. It is implemented in the object oriented programming language SMALLTALK based on Simula for the sequential simulation of concurrent systems. They communicate cooperatively by message passing and subsequent actions of the actor are determined by the local messages and its current behaviour. An actor may make local decisions, send communications or alter its behaviour. They can also make new objects and so can create organisational structures composed of multiple objects organised dynamically to suit the required process. PLANNER used pattern directed invocation of procedures by advertising them as goals, i.e.. it specified directly what procedures to use to solve the goal. It is more indirect than procedure directed invocation and allows the exploration among alternatives. It used both forward and backward chaining of production rules to create subgoals. Furthermore, theorems of the form (implies A B) which permit logical rules of inference are used as declarative sentences when used as data and as imperatives when executed.

Global emergent behaviour from local interactions of units is exemplified by the ant or termite colony which comprises simple behaviours with complex interactions. Each ant responds to the behaviour of all other ants in its colony collectively. Ants appear to indulge in dynamic programming processes whereby each step of its behaviour is completed and its result defines the program for the next step without the need for overall control. A small change in the individual behaviour patterns will be amplified globally (the multiplier effect), e.g. nest building in termites. Colonies of “virtual” ants (vants) have been simulated with global societal behaviour emerging from local interactions between individuals simulating pheromone trails. Each vant’s behaviour is determined by its environment, but the environment is also the result of its past behaviour. However, due to the influence of other vants, each vant also responds to the past behaviour of other vants collectively, i.e. vants interact indirectly through their effects on the environment.

4. CONCLUSIONS

In this report, we have provided a snapshot survey of the field of biomimetics around the world. There are arguments that suggest that traditional engineering approaches are most appropriate for developing engineering solutions, the commonly quoted example being fixed wing aircraft rather than flapping wings. This argument however does not apply to autonomous robotic systems – engineering solutions have yet to be found so biological organisms which exhibit autonomous behaviour *par excellence* provide both a existence proof and an excellent model of autonomous robotic agents. We have emphasised application with respect to robotics in anticipation of future potential application for robotic space exploration missions. We anticipate that the next generation of space explorers will emphasise in-situ exploration with mobility capability. This is particularly the case as space exploration missions give way from global reconnaissance to more focussed *in-situ* investigation requiring enhanced capabilities. To that end, we have emphasised those aspects of biomimetics which have application to robotic rovers. Biomimicry applied to space engineering promises the prospect of high miniaturisation, high integration and packaging efficiency, high energy efficiency, and high autonomy and robustness. Such robustness and adaptability is particularly critical in space exploration as the environments to be explored are typically unknown with unknown dynamics and variability. We highlight five main broad biological principles that are particularly relevant to space systems:

1. The principle of autonomous behaviour to minimise reliance on ground systems
2. The principle of integration of mechanics, control, software and electronic systems
3. The principle of robustness for survival under widely variable conditions
4. The principle of compliant, multi-functional structures as part of a control system
5. The principle of neurally-inspired control systems with selectionism (e.g. neural Darwinism).

Control aspects were the integrative theme as control is the mechanism for autonomous intelligent functionality. In this report, we have considered twelve aspects of biomimetics applied to robotics and related issues:

1. Biological modularity and pleiotropy – biological approaches to modularity and hierarchical structures were considered and were relevant to the organisation of intelligent control systems
2. Biomimetic materials, structures and mechanisms – biological materials by virtue of their hierarchical organisation are composites and incorporate sensors and actuators which may be emulated by smart materials and multi-functional structures, especially with the introduction of compliance
3. Biomimetic methods of propulsion – the incidence of walking, flying and slithering in animals offer robust modes of locomotion and with human manipulation capabilities are applicable to robotics
4. Biomimetic energy generation – artificial photosynthesis and bacterial fuel cells offer potential autonomous power raising but are likely in the near future to be uncompetitive with more conventional techniques
5. Immunological approaches to self-defense – immunological analogues for computer protection and fault tolerance are applicable to computer systems including those associated with robotic control processing

6. Biological behaviour control and navigation – biological extraction of animal behaviours provides a useful model for robotic applications, particularly reflexes and navigation required for robotic explorers
7. Biological learning and memory – learning is a fundamental component of memory and for generating sophisticated robotic behaviours on-line, particularly CMAC-based neural networks
8. Genetic & neural approaches to learning – exploitation of genetics-based neural networks provides a powerful methodology for robust behaviours and navigation in robotic rovers
9. Biomimetic sensing – biological sensing, particularly active vision, electronic noses and taction, offer the possibility of enhanced scientific instrumentation and manipulative capabilities respectively in robotic rovers
10. High-level cognition – expert systems and artificial intelligence methods attempt to model higher level human cognition and have application in robotics for planning future behaviours
11. Hybrid approaches to intelligent control – hybrid methods attempt to combine the advantages of behaviour-based, neural-based and symbol-based approaches to intelligent control in robotics
12. Group behaviour – group behaviours for coordinating teams of robotic rovers offers a distributed approach to planetary exploration

Generally, biological approaches naturally lend themselves to robotics as animals provide excellent models for autonomous robotic agents. Indeed, robotics is fundamentally a biologically-inspired engineering discipline based on the emulation of human manipulative and locomotory capabilities. Autonomous agents will be essential for the next generation of robotic explorers. Each of these areas has potential application in space systems, robotic explorer missions in particular. We consider such applications in detail in TN3.