‘Scent of Science’: Autonomous Source Localization for Exploration (Ref: 11-6301)

Type of activity: Standard study (25 k€)

1 Background & Motivation

1.1 Introduction

Intelligent decision-making in space systems is increasingly used to overcome the limitations derived from communication delays and overloaded bandwidths. An application often considered is that of exploring the surface of planets without direct human supervision, which is required whenever remote control is impractical [Bajracharya et al., 2008, Joudrier et al., 2009]. This applies for instance to planetary rovers, unmanned aerial/underwater vehicles, or survey satellites. To this day, however, autonomy in navigation and exploration (with a few noteworthy exceptions [Smith et al., 2007, Estlin et al., 2007, Estlin et al., 2009]) has been limited to providing basic behaviours –e.g. in rovers, obstacle-avoidance or detection of traversable areas. Higher-level behaviours that require scientific expertise (e.g. tracking chemicals, looking for water, or seeking footprints of geological activity) are still human controlled. Automating also these operations would imply that long-term plans and paths toward promising areas are to be identified by the agent itself.

The problem of autonomously locating areas of rich scientific content is compared to that of odour-source localisation, i.e. the sought locations (initially unknown) are treated as ‘sources’ that spread information (or ‘clues’) in their surrounding; navigation under such conditions is based on reasoning upon those hints, using them to trace back the path to their origin [Murlis et al., 1992]. This problem is complex when dealing with environments where clues are sparse and subject to randomness. It is a frequent situation in biology, and is solved daily by animals looking for food or mates. An illustrative example is that of a moth tracking a female in an open space; only seldom pheromones encountered along the way are at its disposal to guide the search in a world ruled by turbulent winds. Scent-tracking is one of the most remarkable examples of source localisation in nature. It has been widely studied and mimicked, in part due to its utility to solve problems for which engineering approaches usually show poor performance (e.g. gas-leak or pollutant detection), and has led to the development of efficient computational strategies [Kowadlo, 2008]. Unlike purely reactive control designs, these strategies utilize knowledge of the source environmental footprint (i.e. a model of the expected distribution of ‘clues’ around an odour-source) that helps infer where to move.

This study aims at investigating the use of similar strategies in connection to the localisation of different types of sources, and in situations where additional constraints arise (such as the limited measuring capabilities typically faced in planet exploration, and the high energy costs and time delays involved in satellites’ trajectory adjustments). The conjecture is that some scientifically interesting aspects can be modelled like sources, which spread the equivalent of a ‘scent-plume’ according to a recognizable dynamics model that helps in localizing the source.

1.2 Basic principles of plume-model based scent tracking

Odour-source localisation problems face the difficulty of finding a source by way of the clues it spreads into a highly dynamic surrounding: chemical particles (which give vital clues as to the source location) exhibit a chaotically evolving structure, where patches of odour are surrounded by wide voids of clean air, and gradients do not always point toward the source. Information is thus sparse, intermittent and unreliable, making it difficult to infer anything. A physical model of the scent plume may be helpful for interpreting the encounters and deciding where to move. Such a model would convey the dynamical properties underlying the distribution of particles in the environment, and thereby provides a definition of 'what to expect' given 'what is perceived'.

Since the instantaneous distribution of the plume around an odour-source is far too complex and uncertain to be described in computational models, stochastic approaches have been put forward to capture the uncertainty in the world. They create a belief (or source likelihood) map, computed through an inference process that combines encountered clues and an approximated model of how they spread in the environment. The belief map is dynamical: it is iteratively updated as the agent navigates and accumulates information. This belief is then employed in the decision-making process to guide the search. Navigation algorithms based on probabilistic internal models have been the focus of past research. Examples include [Balkovsky and Schraiman, 2002], [Vergassola et al., 2007], [Farrel and Pang, 2003], [Pang and Farrel, 2006], [Ferri et al. 2007], and [Kowadlo and Russel, 2009].

1.3 Sampling strategy

The application of the above ideas to space agents requires accounting for additional requirements that have not been considered in prior studies. Typically in space missions (e.g. [Vago et al., 2006]), obtaining measurements requires some sort of 'sampling' (e.g. drilling the ground to measure concentrations, visual processing and transmission, etc.). Exploitable 'clues' are therefore associated with a cost (e.g. in terms of energy expenditure, time and maximum number of attempts), and are available at sampling locations only, discrete in time and space. This is in contrast with scent-tracking cases, where measurements of odour concentrations in the air could be performed constantly and without cost as the agent navigates. It will be thus a requirement that recordings be undertaken at carefully chosen spots only, according to their expected 'impact' between what is needed and what it costs. A planning module will need to decide where and when sampling is to be carried out, accounting for these costs.

2 Study Description

2.1 Main objectives

This study attempts to assess the potential of using odour-source localization as a paradigm (i.e. use of a model about the effect of the source on the environment, and employ it to reason about the world as clues are encountered) to design the path-planning layer of an autonomous space agent (rover, UAV, survey satellite etc).

Therefore, the main objectives of this study are:

1. Development of the model

a) Propose one or more possible ‘sources’ related to space exploration. A ‘source’ is a scientifically-interesting area on the planetary surface, that spreads ‘clues’ in the surroundings in a way that can be statistically modelled. The main example on which the ACT has performed preliminary experiments [Moraud and Martinez, 2010] is the localization of chemical substances in the air. The algorithms tested can for example be used to localise methane sources on planets with an atmosphere. However, the application and the study is not restricted only to odour localisation, but research teams are encouraged to propose and argue for more suitable and interesting applications. Examples of such different sources and clues may include: the presence of a cyclone which may be detected by the cloud patterns in the surrounding areas, observable in spectral satellite images; microbial life on Europa’s under-ice oceans identifiable through biochemical traces diffused in the water; an area with high concentration of rare minerals on an asteroid, most suitable for mining, identifiable through chemical analysis of soil samples obtained by drilling; underground water or microbial life, possibly identifiable through ground properties such as temperature and consistency – assuming the availability of a mole-like drilling robot.
b) For the proposed source, suggest a model of its footprint on the surroundings (dynamic or static), i.e. how the clues spread around the source, and list the parameters needed to build such a model.
c) Additionally, since the retrieval of samples will imply different costs (e.g., costs of drilling, taking an image, spectroscopic measurements, etc), the measurement procedure and cost associated to sampling retrieval for a given source will need to be quantified.

2. Use this model to create and update the source-uncertainty map of the source location as the agent proceeds along its path and encounters clues. Examples could include probabilistic / Bayesian reasoning, stochastic representations (e.g. diffusion processes), or cognitive-inspired architectures for uncertainty representation. One requirement is that the area of interest is not to be confined in advance, excluding the use of fixed rectangular maps with a relatively small number of cells.

3. Use the proposed uncertainty map to guide the search. The proposal should already contain an intended strategy that relies on the probabilistic belief to guide the agent, and to decide where to sample next.

4. Assess the performance of the navigation strategy (as a function of the cost of retrieving clues, the success of localizing the source and the time to reach the goal have to be considered as criteria), and compare it with respect to simpler search-models that do not employ the internal model (e.g. stochastic gradient). Finally, evaluate the effect of varying the cost in the strategies adopted.

Proposals should contain the description of approaches for all of these main objectives. In particular the proposals should provide detailed answers to the first 2 points, which constitute the main work-packages of this study.

3. Collaboration with the Advanced Concepts Team of ESA

The project is mainly addressed at research groups working in the fields of probabilistic modelling, artificial intelligence or robotics with application to autonomous navigation, teamed up with a scientist (e.g. geologist, biologist, climatologist, etc) who will provide expertise in the characterization of the physical processes to be modelled. The project will be conducted in close cooperation with ACT researchers: algorithmic evaluations will be run in parallel for points 3 and 4, as well as algorithmic benchmarking. In addition, ACT researchers will be leading the selection of test cases, in particular the definition of the arenas to be used to test the efficacy of the proposed algorithm.

References

B.I. Schraiman and E.D. Siggia (2000), Scalar Turbulence, Nature, 405:639-646

E. Balkovsky and B.I. Schraiman (2002), Olfactory search at high Reynold numbers, PNAS, 22(20):12589-12593

M. Vergassola, E. Villermaux and B.I. Shraiman (2007), Infotaxis as a strategy for searching without gradients, Nature. 445:406-409

E.M. Moraud and D. Martinez (2010), Effectiveness and robustness of robot infotaxis for searching in dilute conditions, Frontiers in Neurorobotics, 4:1

J. A. Farrell and S. Pang (2003), Plume Mapping via Hidden Markov Methods, IEEE Transactions on Systems, Man and Cybernetics - part B: Cyb, 33(6):850-863

S. Pang and J. A. Farrell (2006), Chemical Plume Source Localization, IEEE Transactions on Systems, Man and Cybernetics - part B: Cyb, 36(5):1068-1080

G. Kowadlo and R.A. Russell (2006), Using naive physics for odor source localization in a cluttered environment, Autonomous Robots, 20:215-230

G. Kowadlo and R.A. Russell (2009), Improving the robustness of naive physics airflow mapping, using Bayesian reasoning on a multiple hypothesis tree, Robotics and Autonomous Systems, 57:723-737

G. Ferri, M.V. Jakuba, E. Caselli, V. Mattoli, B. Mazzolai, D.R. Yoerger, and P. Dario (2007), Localizing multiple gas / odor sources in an indoor environment using Bayesian occupancy grid mapping. Proceedings of ICRA 2007.

Complementary Reading

Murlis, J., Elkinton, J. S., and Cardé, R. T. (1992). Odor plumes and how insects use them. Annu. Rev. Entomol. 37, 479–503.

G. Kowadlo (2008), Robot Odor Localization: A Taxonomy and Survey, International Journal of Robotics Research, 27(8):869-894

M. Bajracharya, M.W. Maimone and D. Helmick (2008), Autonomy for Mars Rovers: Past, Present, and Future, Computer, 41(12): 44-50

J. Vago et al. (2006), ExoMars: Searching for Life on the Red Planet, ESA Bulletin 126.

L. Joudrier et al. (2009), Challenges of the ExoMars Rover Control, AIAA Infotech@Aerospace, Seattle, WA, USA.

T. Estlin et al. (2007), Increased Mars Rover Autonomy using AI Planning, Scheduling and Execution, IEEE International Conference on Robotics and Automation ICRA’07, Rome, Italy.

T. Estlin et al. (2009), Automated targeting for the MER Rovers, AIAA Infotech@Aerospace, Seattle, WA, USA.

T. Smith et al. (2007), Life in the Atacama: Science autonomy for improving data quality, Journal of Geophysical Research, 112, G04S03