Probabilistic Computing for Efficient Robotic Vision in Space

Type of activity: Standard study

1. Introduction

1.1 General Goal

In space robotics, the computational efficiency of algorithms is at a prime: tasks such as autonomous landing or rover navigation require a quick and efficient determination of actions, given an amount of energy and processing capacity that is much more restricted than in earth-based scenarios.

In this study we focus on a novel integrated software and hardware approach for reducing the computational effort and energy expenditure of computer vision algorithms. It is based on local sampling on the software side and probabilistic computing on the hardware side.

Specifically, on the software-side, the Advanced Concepts Team of the European Space Agency will study random local sampling techniques that reduce the number of necessary computations at the cost of a slightly lower accuracy. On the hardware-side, a research group with expertise in probabilistic computing will study such computing as a means of saving energy at the cost of occasional errors in the calculations. Both aspects are further explained below.

1.2 Software – efficient and robust algorithms

Visual attention models could form a successful strategy for reducing the computational effort of computer vision algorithms. The central idea of visual attention models is that they devote most of the processing to the most relevant parts of the image. Surely, the main challenge is to automatically determine what is `relevant'. Several models have been devised that can determine relevant parts of an image, but many of them extensively process the entire image for determining regions of interest, inadvertently still leading to a large computational burden [1, 2, 3].

In order to achieve a higher computational efficiency, many attention models perform visual tasks on the basis of a restricted number of local samples [4,5,6,7]. In particular, these studies focus on informed sampling, in which the information from the current sample is used to select the next. This can lead to large computational efficiencies, but also often creates a challenging Partially Observable Markov Decision Problem (POMDP). Such a POMDP is currently difficult to solve, and the mentioned studies either make strong assumptions on the task [4] or have to train a model for each different task [6].

In this study the more straightforward strategy of random sampling is employed. It has been demonstrated that in some cases random sampling suffices to obtain speed-ups in the order of a 100 times with an unnoticeable cost in accuracy [8,9,10,11]. Below a figure illustrates this for an application of pitch and roll estimation of an Unmanned Air Vehicle that processes local image samples to find the horizon line. The y-axis represents the absolute error in each estimate, the x-axis the number of samples, with full sampling on the far right. It can be observed that most performance is gained after a 1000 samples the performance does not noticeably improve.

Extracting only 1000 samples results in a speed-up of 14 times for a small image of 160 x 120 pixels (a border of 10 pixels is used in these specific experiments).

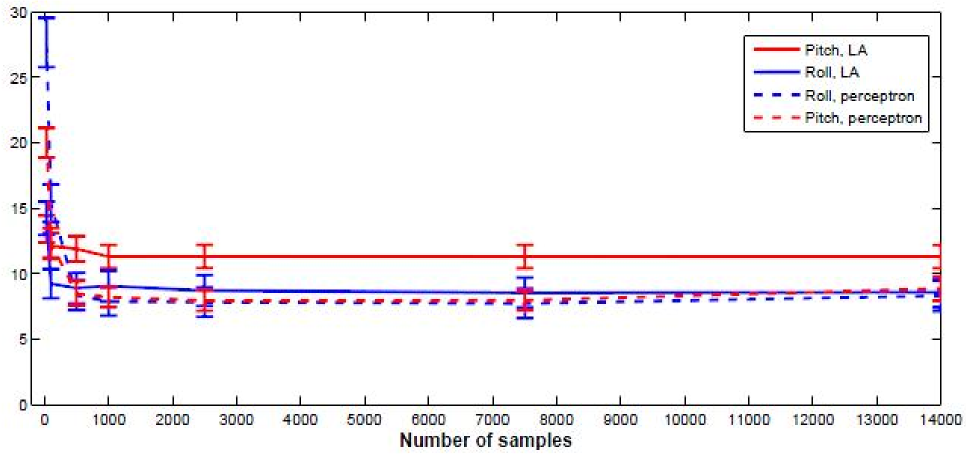

Figure 1: Number of samples extracted from an image versus average absolute errors in the estimates of the pitch and roll angle. The solid lines are for the results in which a linear algebraic solution is used for estimating the horizon (LA), the dashed lines are for the results with a linear perceptron.

1.3 Hardware – probabilistic computing

If one accepts that vision algorithms can gain in efficiency by using only part of the information in an image, one can extend this reasoning to the hardware involved in the calculations to save even more time / energy. In typical processor hardware considerable amounts of energy are spent on obtaining correct calculation results, e.g., for adding or multiplying numbers. In probabilistic computing [12,13,14], the energy spent by the processing units is lowered, resulting in an increase of the probability that some operations might go wrong. Fortunately, it has been shown that the amount of energy saved is significantly larger than the amount of probability traded in [12] - see the figure below. As a consequence, in theory, large energy gains (in the order of 5 times) can be obtained at a minimal cost in calculation errors.

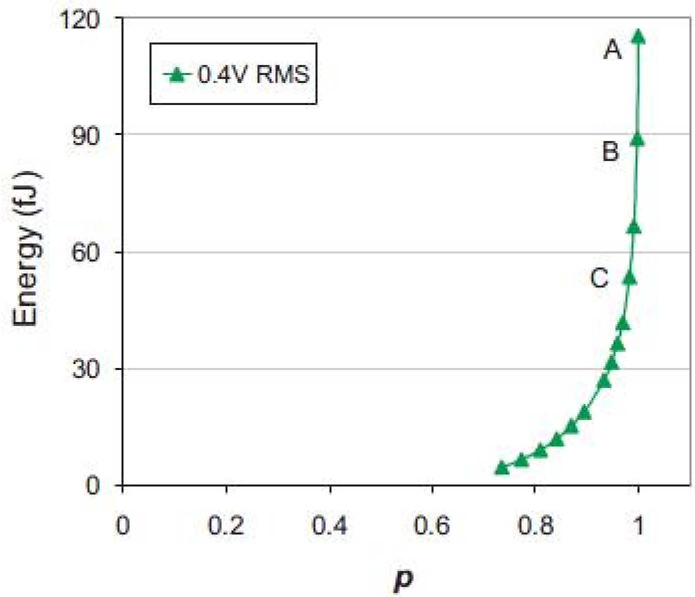

Figure 2: Energy expenditure versus the probability p that calculations will be incorrect. Figure adopted from [12].

Different applications of probabilistic computing have been investigated. The study in [12] focused on the application of probabilistic computing to the calculation of the Fourier transform, showing that for human subjects the probabilistic reconstruction is visually identical to the error-free one. Another study involved the application of probabilistic computing to Continuous Restricted Boltzmann Machines (CBRM) [13]. Such CBRMs normally require the generation of pseudo-random numbers. Instead, in [13] probabilistic computing naturally provided random numbers.

This study proposes to perform research on probabilistic computing in a different context, namely that of a vision algorithm for robotics. The main idea behind the application is that a slight loss in accuracy in the computations of the vision algorithms can still lead to robust and successful retrieval of information on the environment.

2 Research approach

The research approach is to study the application of probabilistic computing in the context of a specific vision algorithm. In Section 2.1 we will explain the stereo vision algorithm that will serve as a starting point for the research. In Section 2.2 we will address the research activities for the probabilistic computing.

2.1 Specific vision algorithm: stereo vision

The extraction of information on the three-dimensional structure of the world is important to almost all robotic applications. The current Mars Rovers mainly rely on stereo vision: the extraction of depth from two images made by two cameras positioned at a known distance from each other.

In the last decade, there has been a considerable progress in stereo vision algorithms. On the one hand, there have been more successful applications of robotic stereo vision, for example [15,16].

On the other hand, many techniques have been introduced to improve the accuracy of stereo vision algorithms, cf. [17,18].

The most basic stereo vision algorithm takes a window from image 1 at position (x,y) and compares it with windows from image 2 at positions (x+Δx,y) for various Δx. The comparison typically involves the sum of the squared differences between (grey-scale) pixel values:

$$C = \sum_{i,j \in W} (I_1(x+i, y+j) - I_2(x+\Delta x+i, y+j))^2 \qquad , \qquad (1)$$

where i and j are coordinates within a window W, I_1 contains the grey-values of the first image and I_2 contains the grey-values of the right image. The Δx leading to the lowest C-value represents the disparity at (x,y) in image 1 and is a measure for the distance of the object seen at that image location. Most of the recently introduced more accurate methods combine the so-obtained local disparities with global constraints. One common method [18] is to segment image 1 (on the basis of color and / or texture) and assume that each segment conforms to a certain model. For example, one can assume that each segment conforms to a plane. The disparities in the segment then serve as data for a plane fitting process, which smoothes out many of the local errors.

In this project, we propose to speed up the calculation of stereo vision with the following method: instead of determining a disparity for every image location in image 1, only a subset of the locations is processed. In order to fit a plane through a segment, at least three disparities are necessary. To get a robust estimate, more samples are advisable, yet it is not necessary to determine the disparities for all pixels in the segment. Determining fewer disparities leads to fewer computations, but results in a less accurate estimate of the plane parameters. As in other cases of (random) sampling, we expect the trade-off between the number of samples and the accuracy of the estimates to be rather gentle: allowing significant speed-ups at a negligible cost in accuracy.

2.2 Probabilistic processing architecture

The bulk of the stereo vision processing consists of calculating equation 1. The idea is to implement the related calculations in a specifically designed probabilistic processing architecture.

The proposals of the groups applying for performing this part of the research should include argumented choices concerning the following issues:

- Different designs are possible of the components that will perform addition / subtraction / multiplication. One example is that in an adder, all bits can be provided with the same voltages, or higher voltages can be applied to the higher order bits (as in the BIVOS approach [12]). Another example is that instead of variable voltages and a traditional representation of integers, a different representation of numbers can be employed in combination with a single voltage for all bits.

- The proposal should clarify which methodology they want to employ. In particular, it should address the connection between the probabilistic computing architecture and the stereo vision algorithm. The stereo vision algorithm will be implemented in MATLAB. In a first possible methodology the work is split in two independent parts. The hardware group then identifies for each operation in equation 1 the probabilities and types of errors that can occur at different voltages. These errors can then be modelled in the MATLAB scripts. A second possible methodology would be to couple the MATLAB stereo vision script with hardware simulator software: calculations would then be done directly by the simulated hardware. A third possible methodology is to have actual hardware in the loop. The first methodology is the baseline requirement for the study. In any case, the chosen methodology should provide for a way to compare the amount of energy consumed in a probabilistic and a non-probabilistic scenario.

- The design should address how the dedicated probabilistic computing parts for addition / subtraction / multiplication fit within a further non-probabilistic CPU / DSP. How would they interface and does sending numbers to these dedicated probabilistic processing components incur additional costs?

3 Study Objectives

The main objectives of this Ariadna study on efficient computing would be:

1. To propose a probabilistic processing architecture that would best fit the stereo vision algorithm. This requires answering questions like whether the processing structure should involve both the additions and multiplications involved in the stereo vision algorithm, etc.
2. To evaluate the energy efficiency gained by employing the probabilistic processing structure, and analyse its relation to the performance (in terms of the probability and size of the calculation errors and in terms of the 3D reconstruction).
3. To perform a preliminary investigation on the influence of radiation particles on the calculations performed at lower voltages.

4 Collaboration with the Advanced Concepts Team of ESA

The study is addressed at a research group with knowledge of probabilistic computing. The research will be conducted in close cooperation with ACT researchers.

The ACT will contribute to the research project the knowledge of stereo vision algorithms, including the necessary code. This will constitute the basis for exploring the way in which probabilistic computing can be applied. It is expected that analysis from a probabilistic computing point of view could lead to adjustments of the stereo vision algorithm. The end goal of the interaction should be a stereo vision algorithm that in combination with a probabilistic computing architecture should exploit the trade-off between energy and reconstruction result as well as possible.

5 References

[1] L. Itti, C. Koch, E. Niebur, A model of saliency-based visual attention for rapid scene analysis, IEEE Transactions on Pattern Analysis and Machine Intelligence 20 (11) (1998) 1254 - 1259.

[2] U. Rajashekar, L. Cormack, A. Bovik, Image features that draw ¯x- ations, in: 10th IEEE International Conference on Image Processing (ICIP 2003), Barcelona, Spain, Vol. 3, IEEE Computer Society, 2003,pp. 313-316.

[3] A. Torralba, A. Oliva, M. Castelhano, J. M. Henderson, Contextual guidance of eye movements and attention in real-world scenes: the role of global features in object search, Psychological Review 113 (4) (2006) 766-786.

[4] J. Denzler, C. Brown, Information theoretic sensor data selection for active object recognition and state estimation, IEEE Transactions on Pattern Analysis and Machine Intelligence 24 (2) (2002) 145-157.

[5] N. Sprague, D. Ballard, Eye movements for reward maximization, in: S. Thrun, L. Saul, B. SchÄolkopf (Eds.), Advances in Neural Information Processing Systems 16, MIT Press, Cambridge, MA, 2004.

[6] G. de Croon, E. Postma, H. van den Herik, A situated model for sensory-motor coordination in gaze control, Pattern Recognition Letters 27 (11) (2006) 1181-1190.

[7] S. Jodogne, J. Piater, Closed-loop learning of visual control policies, Journal of Artificial Intelligence Research 28 (2007) 349-391.

[8] L. Xu, E. Oja, A new curve detection method: Randomized hough transform, Pattern Recognition Letters 11 (1990) 331 - 338.

[9] J. Shotton, J. Winn, C. Rother, A. Criminisi, Textonboost: Joint appearance, shape and context modeling for multi-class object recognition and segmentation, in: ECCV 2006, 2006.

[10] C. Barnes, E. Shechtman, A. Finkelstein, D. Goldman, Patchmatch: A randomized correspondence algorithm for structural image editing, in: ACM Transactions on Graphics (Proc. SIGGRAPH), 2009.

[11] de Croon, G.C.H.E., de Weerdt. E., de Wagter, C., Remes, B.D.W. "The appearance variation cue for obstacle avoidance.", in the IEEE conference on Robotics and Biomimetics 2010 (ROBIO 2010).

[12] Probabilistic Arithmetic and Energy Efficient Embedded Signal Processing (Received Best Paper Award.) Jason George, Bo Marr, Bilge E. S. Akgul and Krishna V. Palem. Proceedings of the Intl. Conference on Compilers, Architecture and Synthesis for Embedded Systems (CASES), Seoul, Korea, October 23-25, 2006.

[13] Probabilistic Computing with Future Deep Sub- Micrometer Devices : A Modelling Approach, Hamid, Nor H, Murray, Alan F, Laurenson, David, Roy, Scott
Cheng, Binjie, Symposium A Quarterly Journal In Modern Foreign Literatures, 2005, pages: 2510-2513

[14] Error Immune Logic for Low-Power Probabilistic Computing, Marr, Bo, George, Jason, Degnan, Brian, Anderson, David V., Hasler, Paul, VLSI Design, 2010, pages 1-10

[15] Stereo based navigation in unstructured environments, Mark, W. van der, Groen, F.C.A., & Heuvel, J.C. van den (2001). In IEEE Instrumentation and Measurements Conference (pp. 2038-2042). Budapest, Hungary.

[16] Learning long-range vision for autonomous off-road driving, R. Hadsell, P. Sermanet, J. Ben, A. Erkan, M. Scoffier, K. Kavukcuoglu, U. Muller, and Y. LeCunn

[17] Advances in computational stereo, M.Z. Brown, D. Burschka, and G.D. Hager. In IEEE Transactions on Pattern Analysis and Machine Intelligence, vol. 25, no. 8, 2003.
[18] A region based stereo matching algorithm using cooperative optimization, Zeng-Fu Wang, Zhi-Gang Zheng, 2008 IEEE Conference on Computer Vision and Pattern Recognition, pp.1-8.