Rover Science Autonomy: Probabilistic Planning for Science-Aware Exploration
Doctoral Consortium Thesis Summary

Trey Smith (advisor Reid Simmons)
Robotics Institute, Carnegie Mellon University
{trey,reids}@ri.cmu.edu

Abstract

Future Mars rovers will have the ability to autonomously navigate for distances of kilometers. In one sol a traverse may take a rover into unexplored areas beyond its local horizon. The rover can explore these areas more effectively if it is able to detect and react to science opportunities on its own, what we call *science autonomy*. We are studying science autonomy in two ways: first, by implementing a simple science autonomy system on a rover in the field, and second, by developing probabilistic planning technology that can enable more principled autonomous decision-making in future systems.

Introduction

Ongoing improvements in Mars rover mobility, power systems, and navigation provide both an opportunity and a challenge: we need new techniques for performing effective science while over the horizon and out of contact. This work deals with *science autonomy* (SA), the ability of a rover to reason about science goals and the science data it collects in order to make more effective exploration decisions.

SA extends the set of rover operational modes available to scientists, adding new modes that are more effective in over-the-horizon situations. For instance, in the *science on the fly* mode the rover would opportunistically sample interesting features observed during long traverses. In the *intelligent site survey* mode the rover would characterize a site, choosing its coverage and sampling strategies in order to assemble a useful summary of what is present.

The goal in SA is not to take control of the rover from scientists. Rather, it should provide them with new modes of operation to use or not as they see fit. The science team can begin to think of the rover as a robotic graduate student (Gulick *et al.* 2001) that interprets their commands intelligently, and can take some initiative when exploring new areas.

We are studying science autonomy in two ways: first, by implementing a simple science autonomy system on a rover in the field, and second, by developing probabilistic planning technology that can enable more principled autonomous decision-making in future systems.

The first part of our proposed work is to develop and field test a prototype SA system that we can use to get a better understanding of the problem. Some components of this work are to create (a) an overall system architecture that can support the kind of SA operational modes described earlier, (b) an interface for interacting with scientists, and (c) a prototype SA planning module.

The second part of our proposed work is to extend partially observable Markov decision process (POMDP) planning algorithms in ways that bridge the gap between the current state-of-art and realistically sized SA problems. POMDP planners can generate high-quality plans that take into account action and sensing uncertainty, but realistic problems in the SA domain are far beyond existing algorithms.

Prototype SA System Development

Our progress in SA system development has so far fallen into the three main areas of architecture, priority representation, and onboard planner development.

First, we have developed a software architecture and operations protocol that supports SA (fig. 1). One interesting aspect of the architecture is that the same planning module is employed both off-board (used interactively by scientists) and onboard (adapting the rover's plan to changing conditions at execution time). For more see (Smith 2004).

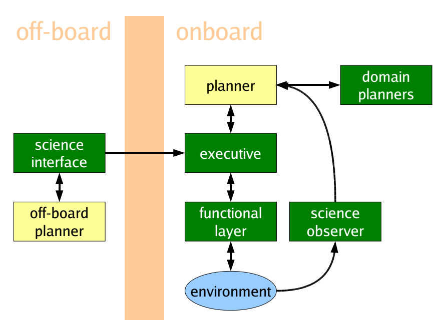

Figure 1: Science autonomy system architecture.

Second, in order to make effective decisions about science actions, the rover must understand scientist priorities. We have developed and tested a formal representation for these priorities. Abstractly, the priorities constitute a value function that scores possible returned data sets, allowing them to be ranked from most to least preferred. A set of target signatures forms the core of the function representation. Each time a sample in the data set matches a target signature, the value of the data set is incremented. More advanced ideas such as novelty detection and representative sampling are supported by a mechanism for automatically adding new target signatures to the set.

The representation was tested during the 2004 expedition of the Life in the Atacama project, during which the Zoë rover was deployed to the Atacama Desert of Chile for two months (fig. 2). We asked scientists to use the language to specify their priorities for rover study of a particular site. On the mission control side, scientists reported that overall the representation was sufficiently expressive, although they had difficulty setting some of the numerical values (such as those that determine the relative importance of different features of interest). On the rover side, lacking a full SA system for rigorous evaluation, we tested whether the representation could at least be interpreted by a human operator controlling the rover in real time using high-bandwidth telemetry images. The result was ambiguous: the human operator had difficulty controlling the rover, but mainly because its sensing configuration needs to be improved to better support SA. For more see (Smith *et al.* 2005).

Figure 2: Zoë rover in the Atacama desert, 2004.

Third, we have developed and tested a planner that supports the kind of priority trade-off reasoning required by SA, although it has not yet been integrated with a science data understanding module to form a closed-loop SA system. The planner made use of both science activity models and a domain expert planner for long-distance navigation. During the field expedition it was successfully used both in both the off-board and onboard planner roles of our architecture.

In future work, we plan to test a full SA system during the upcoming 2005 Atacama field expedition. Our priority representation will be supported with a user interface that allows scientists to specify and adjust their priorities during the course of the mission. Our planner will be integrated with an onboard science data understanding module and will have the ability to autonomously react to incoming science data.

Probabilistic Planning

The second part of our work is to extend partially observable Markov decision process (POMDP) planning algorithms in ways that bridge the gap between the current state-of-art and realistically sized SA problems. POMDP planners can generate high-quality plans that take into account action and sensing uncertainty, but realistic problems in the SA domain are far beyond existing algorithms.

The main motivation for considering a POMDP model for SA is the importance of *sensor planning*, i.e., modeling the information the rover expects to gain from future sensor readings and the effect it will have on subsequent rover decisions. For example, a rover may have a long-range sensor that returns data with both inherent interest to scientists and *informational* value: once the rover has a long-range reading, it can make a better decision about whether to approach the rock and use contact sensing.

How can the rover choose whether to apply the long-range sensor? Informational value depends on (1) the current level of uncertainty, (2) the importance of the potential rewards affected by the uncertainty, and (3) any constraints (e.g., limited time or obstacles) that may render the information irrelevant. This combination of factors makes it difficult to create useful heuristics for sensing. Planning in the POMDP model, if it can be made tractable, implicitly deals with all of them.

Our main progress to date in this area is development of the Heuristic Search Value Iteration (HSVI) algorithm, which incrementally calculates the POMDP value function by forward exploration from the starting state, using heuristic choices of action and observation to focus updates on parts of the value function that are most relevant (Smith & Simmons 2004). These focused updates allow HSVI to converge orders of magnitude faster on some benchmark problems than well-known algorithms such as Incremental Pruning (Cassandra, Littman, & Zhang 1997) or PBVI (Pineau, Gordon, & Thrun 2003). In future work, we will scale to larger problems by using more compact policy representations and fast replanning.

References

Cassandra, A.; Littman, M.; and Zhang, N. 1997. Incremental pruning: A simple, fast, exact method for partially observable Markov decision processes. In *Proc. of UAI*.

Gulick, V. C.; Morris, R. L.; Ruzon, M. A.; and Roush, T. L. 2001. Autonomous image analyses during the 1999 Marsokhod rover field test. *J. Geophys. Res.* 106(E4):7745.

Pineau, J.; Gordon, G.; and Thrun, S. 2003. Point-based value iteration: An anytime algorithm for POMDPs. In *Proc. of IJCAI*.

Smith, T., and Simmons, R. 2004. Heuristic search value iteration for POMDPs. In *Proc. of UAI*. (to appear; a companion technical report with more detailed theoretical analysis is currently in preparation).

Smith, T.; Niekum, S.; Thompson, D.; and Wettergreen, D. 2005. Concepts for science autonomy during robotic traverse and survey. In *Proc. of IEEE Aerospace Conference*. (to appear).

Smith, T. 2004. *Rover Science Autonomy: Probabilistic Planning for Science-Aware Exploration*. Ph.D. Thesis Proposal, Robotics Institute, Carnegie Mellon University.