Science Autonomy in the Atacama

Trey Smith TREY@RI.CMU.EDU
Robotics Institute, Carnegie Mellon University, 5000 Forbes Avenue, Pittsburgh, PA 15213 USA

Abstract

We have recently embarked on a three-year project, funded by the NASA ASTEP[1] program, to develop robotic astrobiology in the process of learning the limits of life in the Atacama desert of Chile. We see this as an opportunity to develop a more science-aware rover: one that, on encountering a new area, can select interesting features, perform initial experiments, and selectively return relevant data, all before receiving feedback from the science team. Several components of the proposed science autonomy system can make use of classifiers (is this the kind of rock we are looking for?) and clustering algorithms (is this rock like anything we have already sampled?). The unknown character of unexplored areas motivates use of on-line learning techniques.

1. Introduction

We have recently embarked on a three-year project, funded by the NASA ASTEP program, to develop robotic astrobiology in the process of learning the limits of life in the Atacama desert of Chile. The project completed its first expedition to the Atacama this April. The primary goal of this preliminary visit was to learn about the effects of the Atacama environment by testing the Hyperion rover (Figure 1) and sensors (Figure 2) prior to extensive integration in year two. The expedition was highly successful, achieving all 12 of its experimental goals, including measurements of wind, insolation, rover motion power requirements, and instrument validation.

By the third year, the rover will perform science operations under Mars-like communications restrictions. It will integrate fluorescence-based sensors for detection of specific organic molecules, panoramic imagers,

Figure 1. Hyperion in the Atacama.

microscopic imagers, spectrometers, as well as mechanisms for shallow subsurface access.

The measurement and exploration technique produced by this investigation combines long traverses, sampling measurements on a regional scale and detailed measurements of individual features, in pursuit of three primary science objectives:

1. Seek Life: Seek and characterize biota surviving in the Atacama and analyze microhabitats. We will question the hypothesis that the most arid regions of the Atacama represent an absolute desert.
2. Understand Habitat: Determine the physical and environmental conditions associated with identified past and current habitats, including the search for structural fossils, the monitoring of current biological oases and microorganic communities, and learning how these organisms have contributed to the modification of their environment.
3. Relevant Science: Develop, integrate, and field test a suite of science instruments that form a complete payload relevant to the NASA Mars Exploration Program.

We see this project as an opportunity to develop a more science-aware rover: one that, on encountering a

[1] Astrobiology Science Technology for Exploring the Planets

Figure 2. (above) colored LEDs provide illumination for a nighttime fluorescence imager test; (below) a backpackable fluorescence microscope was used for ground truth

new area, can select interesting features, perform initial experiments, and selectively return relevant data, all before receiving feedback from the science team. This ability will dramatically increase the amount of science return per communication cycle.

2. Related Work

This new science autonomy effort will build on capabilities developed for the Robotic Antarctic Meteorite Search with the Nomad robot (Wagner et al., 2001). Nomad autonomously executed coverage patterns, identified rocks on the ice, and classified them as meteorites or terrestrial material.

The Multi-Rover Integrated Science Understanding System (MISUS) project at JPL developed a multi-rover system in simulation (Estlin et al., 1999). The system's top-level goal was to characterize the rock distribution in a region, and it was able to generate sub-goals (new locations to study) based on the frequency of each rock type in different areas. Rover sensors returned an IR spectrum, a point in a 14-dimensional feature space. The rock types that guided sub-goal generation were identified by an online k-means like clustering algorithm.

Our system will operate in a similar rover domain, and also incorporates the concept of automatically generating science sub-goals during the mission. But because we must satisfy the demands of the science team on a real expedition, our system will need to accomodate changing scientist preferences, and competing mission objectives, such as requests for traverse towards distant sites identified from orbit.

The Autonomous Sciencecraft Experiment (ASE) will demonstrate onboard autonomous science on the Techsat-21 Air Force satellite constellation (Davies et al., 2001). ASE has two main focuses. It will use onboard data analysis to increase science return given a fixed downlink budget, reporting summary statistics rather than entire raw data sets. It will also perform detailed observations of interesting events (such as erupting volcanoes) based on pre-defined trigger conditions.

We believe that science autonomy aboard orbiters and planetary rovers present somewhat different problems. Orbiters typically have long mission lifetimes and repeatedly view the same science features. Science team interaction is for the most part not time-critical, and it can be relatively infrequent after an initial shake-down period.

In contrast, planetary rovers have short mission lifetimes and may only visit a particular site once. Because sites vary widely and in unpredictable ways, it is difficult to envision giving the rover detailed science goals and then allowing it to proceed on its own. A more likely scenario involves daily updates to the goals as the science team's initial ideas about an area are refined by rover data. These updates are time-critical since they are relevant to a particular area that may not be visited again. Thus interacting with the science autonomy system will be an important part of the command generation cycle in mission operations, and the associated user interface issues need to be considered carefully.

3. Approach

Figure 3 shows a proposed architecture for our science autonomy system. In this architecture, sensor data is first passed to several feature detectors, each looking for a specific kind of feature. Once features have been identified and segmented out of the background, they are passed to one of several classifiers, depending on feature type. For instance, if the feature is a rock, and an infrared spectrum is available, the system will try to infer the rock's mineral composition. The experiment generator uses this detailed feature information, along with scientist preferences, to generate a list of potential experiments (readings to take), each tagged with a priority level. The science-aware planner generates a plan using these proposed experiments and any explicit commands. The executive and functional layers execute the resulting plans, and enable the rover to interact with the environment at a more detailed level than is modeled by the planner; a discussion of these modules is beyond the scope of this paper.

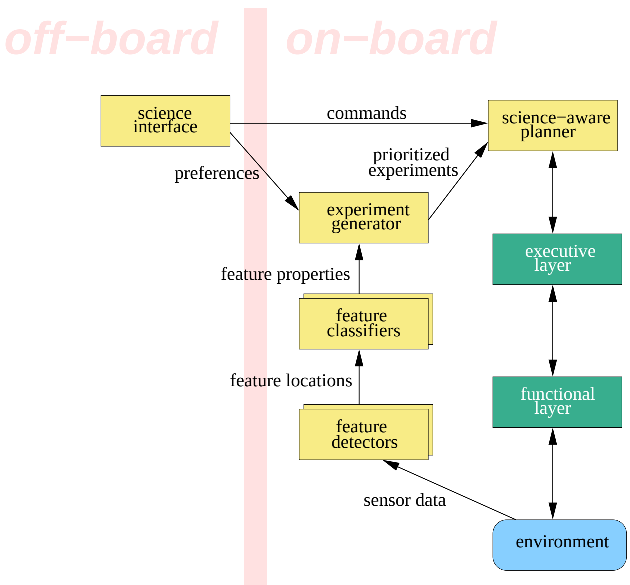

Figure 3. Science autonomy architecture.

Now we will discuss each component in more detail.

3.1. Feature detectors

Some previous work has focused on segmenting rocks from background soil and detecting sedimentary layering (Gulick et al., 2001), both useful capabilities for our purposes. A potential thrust for us is detecting medium-scale features such as outcrops and dry streambeds, using a combination of images and 3D stereo data. Sharp transitions in soil type, or in the type and density of overlying detritus, can be used to identify the boundaries between geological units. This capability is useful not only for opportunistic science (e.g., return an image of any layering you see), but also for target recognition upon approaching targets selected from orbital imagery.

3.2. Feature classifiers

Several kinds of classification could be useful. One obvious example is determining the mineral composition of rocks from spectra and other cues (Gazis & Roush, 2001). This is a well-studied problem in the context of remote sensing, and has also begun to be addressed for rovers. (Pedersen et al., 2001) used a Bayes classifier with graph structure determined by a domain expert, and off-line supervised learning of parameters. Other classifier applications include identifying how weathered a surface is based on spectra or texture cues, and various indications of possible life, such as spiky texture, lumps or ring shapes, and presence of characteristic colors (Figure 4).

On-line learning for rover science is relatively unexplored, but could be useful. During our recent trip to the Atacama, we learned that many of the green "lichens" we had seen were most likely sand grain sized sediments of the green mineral chlorite, cemented to salt. Similarly, a rover could learn that green color is not strongly correlated to other signals of the presence of life, and begin to use other cues.

Figure 4. Example photosynthetic life in the desert environment: (above) presumably lichens; (below) less clear.

3.3. Experiment generator

This module essentially implements a preference function that maps feature type and sensor type to the value scientists place on the resulting data set. As new features are detected, new experiments (sensor readings to take) are generated and prioritized. Our proposed approach relies on a flexible set of rules. Each rule fires when it matches a given feature and type of sensor reading, indicating either increased or decreased interest.

A more sophisticated preference function would also take into account feature context, allowing scientists to prioritize features that are anomalous relative to their surroundings, distant from prior sensor readings, or in a special location (just below an outcrop, on the edge of a salar, on the windward side of a hill, etc.). If there are a few rock types present in an area, it might be desirable to return representative sensor readings for one rock of each type. Both anomaly detection and representative sampling would require on-line clustering (Cheeseman & Stutz, 1996).

3.4. Science interface

Existing rover command interfaces are geared towards generating precise action specifications relative to the known local surroundings of the rover. Some steps have also been taken towards "over the horizon" navigation. However, specifying preferences for what to do when the rover gets over the horizon (not knowing what is there) is a relatively open problem.

The science team should be able to tune the experiment generator rule set in a variety of ways, such as:

- Specifying that a particular data set was interesting, implicitly adding a rule to take the same sensor reading in the future when we see features similar to the one specified. Clustering similar features would help enable this capability. Negative preferences could be added in the same way.
- If rules have parameters, such as relative weight, the interface could help to tune the parameters by presenting the science team with a series of what-if scenarios. This would be particularly helpful if there are rules that fire for more abstract reasons, such as those that prioritize anomalies.

We should note that the ability to command the rover to specific features will not be removed from the interface. Explicit commands and preferences over not-yet-seen features will both be available.

Our project is already collaborating with the Carnegie Mellon Studio for Creative Inquiry to extend their EventScope educational interface so that it can be used for rover science operations (Coppin et al., 2002), (Figure 5). This system was tested successfully over six days of operations during our first expedition. We anticipate that preference specification will be integrated into the EventScope interface.

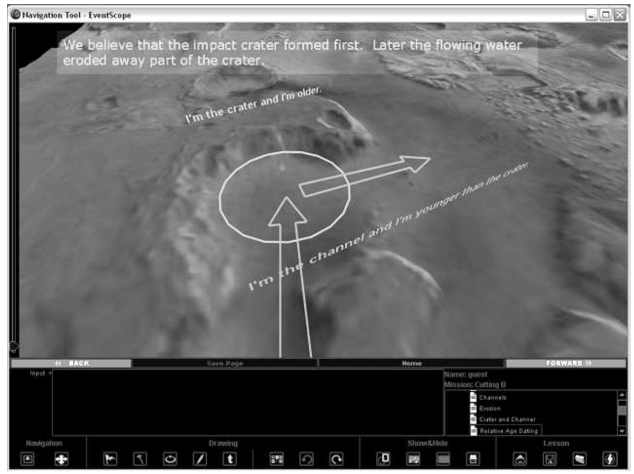

Figure 5. The EventScope 3D educational interface.

3.5. Science-Aware planner

The rover should be able to weigh proposed experiments and explicit commands against time, energy, and other constraints to arrive at a plan. An important point is that we are planning to gain information. Sometimes the rover may take sensor readings not because they are explicitly prioritized by the scientists, but because they are cheap to perform and return more information about whether the expensive, valuable experiments are justified. This kind of planning problem can be formulated as a partially observable Markov decision process (POMDP) (Cassandra et al., 1994; Murphy, 2000).

Unfortunately, POMDP planning algorithms tend to be extremely complex both in theory and in practice. We will employ several techniques to help make them tractable for this domain. Many POMDP algorithms try to come up with a long-term policy at the beginning, then follow that policy throughout execution. We will instead rely on frequent replanning, focusing during each planning episode on approximating the best next action to take. This avoids excessive "what-if" reasoning about situations the rover may never encounter. We are encouraged by a new class of value iteration algorithms that solve large problems by focusing on reachable belief states (Pineau et al., 2003).

Other technical challenges for the planner include:

- Exploration vs. exploitation: The planner will need to decide whether to perform detailed study of known features or look over the hill for new ones. An important question is how to model the potential benefits of moving to a new area. Previous field experiments have tended to overemphasize detailed study, which is understandable, given that the science team wants to look at everything. We believe that faster-paced operations can lead to much greater science return. In the summer of 1997, the Nomad rover performed one week of science operations during its Atacama desert trek. During one day of this week, the science team was constrained to use only 25% of rover time on sensing, the remaining 75% being spent on moving from site to site. This "science on the fly" operational scheme led to the only fossil find of the week, and proved in general to be very productive (Cabrol et al., 2001).
- Plan modification: As new features are detected, new experiments will be proposed by the experiment generator. The planner must be able to efficiently integrate resulting changes into its plan, as some existing planners do (Chien et al., 1999).
- Integration with existing system: The TEMPEST planner was used on Hyperion for the recent field expedition (Tompkins et al., 2002). TEMPEST considers time, terrain, and solar ephemeris to estimate power consumption and production, generating resource-cognizant plans for navigation that also include pre-specified science operations. Integrating science awareness into this already complex planning framework may prove to be a difficult challenge.

4. Summary

As we have seen, several components of the proposed science autonomy system can benefit from use of machine learning techniques, from feature classifiers to clustering algorithms for anomaly detection. Different types of training will be possible depending on the application. Here are some examples:

- Pre-mission training (possibly supervised): Generate initial classifiers, e.g., for mineral composition, based on data from the lab or from an initial shake-out period in the field.
- Daily training during mission (possibly supervised): Retrain parameters in classifiers based on the frequency of different feature types observed in the local environment.
- Training on the fly (unsupervised): Cluster features observed at a new site so that sensor readings for one representative feature of each type can be returned.

In order to support planning to gain information, we plan to focus on classifiers that explicitly model priors and output probability distributions.

We are still in the early planning phases of our science autonomy development, and the scope we will eventually tackle remains to be defined. We hope for continued input from the machine learning community as we proceed.

Acknowledgements

Many of the ideas presented here came out of a series of discussions with David Wettergreen and Michael Wagner.

References

Cabrol, N. A., Chong-Diaz, G., Stoker, C. R., Gulick, V. C., Landheim, R., Lee, P., Roush, T. L., Zent, A. P., Lameli, C. H., Iglesia, A. J., Arrerondo, M. P., Dohm, J. M., Keaten, R., Wettergreen, D., Sims, M., Pedersen, L., Bettis, A., Thomas, G., & Witzke, B. (2001). Nomad rover field experiment, Atacama Desert, Chile, 1, science results overview. *J. Geophys. Res.*, *106*, 7785.

Cassandra, A. R., Kaelbling, L. P., & Littman, M. L. (1994). Acting optimally in partially observable stochastic domains. *Proceedings of the Twelfth National Conference on Artificial Intelligence (AAAI-94)* (pp. 1023–1028). Seattle, Washington, USA: AAAI Press/MIT Press.

Cheeseman, P., & Stutz, J. (1996). Bayesian classification (AutoClass): Theory and results. In U. M. Fayyad, G. Piatetsky-Shapiro, P. Smyth and R. Uthurusamy (Eds.), *Advances in knowledge discovery and data mining*. AAAI Press/MIT Press.

Chien, S., Knight, R., Stechert, A., Sherwood, R., & Rabideau, G. (1999). Using iterative repair to increase the responsiveness of planning and scheduling for autonomous spacecraft. *IJCAI99 Workshop on Scheduling and Planning meet Real-time Monitoring in a Dynamic and Uncertain World*. Stockholm, Sweden.

Coppin, P., Wagner, M. D., & EventScope (2002). EventScope: A telescience interface for internet-based education. *Workshop on Telepresence for Education, International Conference on Robotics and Automation*.

Davies, A., Greeley, R., Williams, K., Baker, V., Dohm, J., Burl, M., Mjolsness, E., Castano, R., Stough, T., Roden, J., Chien, S., & Sherwood, R. (2001). *ASC science report* (Technical Report). Jet Propulsion Laboratory.

Estlin, T. A., Mann, T. P., Gray, A. G., Rabideau, G., Castano, R., Chien, S., & Mjolsness, E. D. (1999). An integrated system for multi-rover scientific exploration. *AAAI*.

Gazis, P. R., & Roush, T. (2001). Autonomous identification of carbonates using near-IR reflectance spectra during the February 1999 Marsokhod field tests. *J. Geophys. Res.*, *106*, 7765.

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

Murphy, K. (2000). *A survey of POMDP solution techniques.* http://www.ai.mit.edu/~murphyk/Papers/pomdp.ps.gz: (informal technical report).

Pedersen, L., Wagner, M. D., Apostolopoulos, D., & Whittaker, W. L. (2001). Autonomous robotic meteorite identification in Antarctica. *International Conference on Robotics and Automation* (pp. 4158–4165).

Pineau, J., Gordon, G., & Thrun, S. (2003). Point-based value iteration: an anytime algorithm for POMDPs. *IJCAI*.

Tompkins, P., Stentz, A., & Whittaker, W. L. (2002). Mission planning for the sun-synchronous navigation field experiment. *International Conference on Robotics and Automation*.

Wagner, M. D., Apostolopoulos, D., Shillcutt, K., Shamah, B., Simmons, R., & Whittaker, W. L. (2001). The science autonomy system of the Nomad robot. *International Conference on Robotics and Automation* (pp. 1742–1749).