

Predictive Exploration for Autonomous Science

David R. Thompson

The Robotics Institute, Carnegie Mellon University
Pittsburgh, PA 15206, drt@ri.cmu.edu

Abstract

Often remote investigations use autonomous agents to observe an environment on behalf of absent scientists. *Predictive exploration* improves these systems' efficiency with onboard data analysis. Agents can learn the structure of the environment and predict future observations, reducing the remote exploration problem to one of experimental design. In our formulation information gain over a map guides exploration decisions, while a similar criterion suggests the most informative data products for downlink. Ongoing work will develop appropriate models for surface exploration by planetary robots. Experiments will demonstrate these algorithms on kilometer-scale autonomous geology tasks.

On Remote Autonomous Science

In general today's planetary exploration robots do not travel beyond the previous day's imagery. However, advances in autonomous navigation will soon permit traverses of multiple kilometers. This promises significant benefits for planetary science: rovers can visit multiple sites and survey vast areas of terrain in a single command cycle.

Long autonomous traverses present new challenges for data collection (Gulick et al. 2001). These rovers will travel over their local horizon so that scientists will not be able to specify targets in advance. Moreover, energy and time constraints will continue to limit the number of measurements; sampling density will decrease as mobility improves. Finally, bandwidth limitations will preclude transmission of most collected data. These resource bottlenecks beg the question: is it possible to explore efficiently, with long traverses and sparse sampling, while preserving our understanding of the explored environment?

One strategy to improve sampling efficiency involves onboard data understanding (Pedersen 2001). Pattern recognition technologies enable robots to place instruments and take measurements without human supervision. These robots can autonomously choose the most important features to observe and transmit (Castaño et al. 2003).

This document suggests that these agents must learn and exploit structure in the explored environment. In

other words, they must be mapmakers, representing spatial structure (similarities from one locale to the next) and cross-sensor structure (correlations between different sensing modes). These models guide the agent's exploration to informative areas while minimizing redundant sampling. During downlink, the map can summarize hundreds of measurements in a bandwidth-efficient representation.

Predictive exploration uses these generative models to infer the probabilities of future measurements. One can then describe the exploration task in terms of established principles of experimental design. Information gain over the map guides exploration decisions; a similar criterion suggests the most informative data products for downlink.

Optimal Collection and Return

We treat the exploration process as an attempt to estimate a parameter of interest $\Theta \in \Omega$. This represents the locations of specific features or (more generally) any hidden parameters of the observed process. Optimal actions are those which minimize remote scientists' uncertainty over Θ as defined by the Shannon entropy. We describe any remaining "uninteresting" environmental unknowns with a parameter $\Phi \in \Gamma$.

For simplicity we assume that exploration occurs in two phases (Figure 1). First the remote agent explores the environment while collecting measurements. This *data collection* operation maps $\{\Theta, \Phi\}$ onto observations given by $X \in \mathcal{X}$. The set of collected data may be much larger than the allowable downlink budget. In the second phase the agent chooses a subset of observations $X' \subseteq X$ and sends these back to remote scientists. In other words, a *data return* operation maps X onto returned data X' . The optimal experimental design minimizes the entropy $H(\Theta|X')$ while respecting the constraints of available time, energy, and bandwidth.




Figure 1: Data collection and return.

We make the simplifying assumption that the exploration and return policy is deterministic; for a given policy the unknowns $\{\Phi, \Theta\}$ contain all of the entropy in the system. A value of $\{\Phi, \Theta\}$ fully determines X , which fully determines X' . Then multiple applications of the chain rule give the objective function (up to a constant) in terms of the following upper bound:

$$H(\Theta|X') < -H(\Phi, X) + H(X|X') \quad (1)$$

The first term $H(X, \Phi)$ is familiar from experimental design literature: it favors increasing the entropy of our observations X as in Maximum Entropy Sampling (Sebastiani and Wynn 2000). It also promotes uncertainty about the extra environment parameters Φ ; observations that do not decrease uncertainty in Φ must instead be communicating parameters of interest. The second term describes optimal data return. The data return procedure to minimize $H(X|X')$ is that which best reduces the receiver's expected uncertainty about the set of collected data.

Unfortunately both sampling and return terms involve observations X so the optimal policy must consider potential downlinks during data collection. This requires a difficult combinatorial optimization to evaluate *every* candidate data collection action. Alternatively one could maximize data collection and return terms independently. Here the agent chooses observations to optimize the data collection term through a sequential design procedure such as Bayesian Adaptive Exploration (Loredo and Chernoff 2002). Then it assembles an optimal downlink during the return phase using whatever collected data is available. This approximate solution only requires a single combinatorial optimization at downlink. Independent collection and return policies parallel the source and channel codes of noisy-channel communications; ongoing work aims to characterize their asymptotic error using existing results from information theory.

Inference and Action Selection

The form of the objective function in equation (1) depends on the data model relative to which we assess the entropy of future observations. In other words, the information-driven approach reduces a nebulous question of data's science value to the concrete issue of identifying parameters of interest and correlations among these parameters and data products. We can improve exploration efficiency with models that capture known conditional dependencies. One such correlation is the spatial structure of the environment — any observed smoothness that makes dense samples redundant. The model should also reflect correlations between sensing modes (for example, between images and spectra or between surface and orbital data).

Graphical models are a likely candidate. Previous work used Hidden Markov Models of rover traverse geology for autonomous science tasks like novelty detection in image sequences (Thompson 2006). Here spatial and cross-sensor correlations can define expected information gain for multiple sampling points along linear transects. General Markov Random Fields (Cressie 1991) could extend this notion to spatial maps in higher dimensions. We are also investigating nonparametric models involving local regression and kernel smoothing (Wasserman 2006).

Efficient inference for a particular model does not guarantee that the action selection problem is also tractable. Computing the expected entropy of a single data collection action requires a difficult integration over the unknown parameters and all possible observations (Loredo and Chernoff 2002). This inherent complexity may restrict action selection to greedy or short-planning-horizon solutions. Nevertheless, systems that can reason about information gain can significantly improve time and bandwidth efficiency. Future work will continue to develop spatial models for surface exploration together with appropriate inference and planning algorithms.

Forthcoming experiments in Fall 2007 will demonstrate these principles in a system for autonomous geology that explores planetary analogue environments on kilometer scales. A rover platform from the Carnegie Mellon Robotics Institute will survey terrain to map surface composition using a combination of surface imagery, a VisNIR reflectance spectrometer, and orbital data. Experiments will evaluate the fidelity of maps generated by predictive exploration against status quo periodic sampling strategies.

Acknowledgements

This research progresses with guidance from David Wettergreen, Jeff Schneider, Reid Simmons, Steve Chien, and Phil Christensen. It is supported by NASA grant NNG0-4GB66G.

References

- Castaño, R.; Judd, M.; Anderson, R. C. and Estlin, T. 2003. Machine Learning Challenges in Mars Rover Traverse Science. *ICML Workshop on Machine Learning Technologies for Autonomous Space*.
- Cressie, N. 1991. *Statistics for Spatial Data*. New York: Wiley.
- Gulick, V. C.; Morris, R. L.; Ruzon, M. A.; and Roush, T. L. 2001. Autonomous image analysis during the 1999 Marsokhod rover field test. *J. Geophysical Research*. 106(E4): 7745–7764.
- Pedersen, L. 2001. Autonomous characterization of unknown environments. In *Proceedings of the IEEE International Conference on Robotics and Automation*. 277 – 284.
- Sebastiani, P. and Wynn, H. P. 2000. Maximum entropy sampling and optimal Bayesian experimental design. *Journal of the Royal Statistical Society*. 62(1): 145–157.
- T. J. Loredo, T. J. and Chernoff, D. F. 2002. Bayesian Adaptive Exploration. *Statistical Challenges in Astronomy 275–297*. New York: Springer.
- Thompson, D. R.; Smith, T. and Wettergreen, D. 2006. Autonomous Detection of Novel Biologic and Geologic Features in Atacama Desert Rover Imagery. In *Proceedings of the Lunar and Planetary Sciences Conference*.
- Wasserman, L. 2006. *All of Nonparametric Statistics*. New York: Springer.