

# Intelligent Mapping for Autonomous Robotic Survey

David R. Thompson

CMU-RI-TR-08-33

*Submitted in partial fulfillment of the  
requirements for the degree of  
Doctor of Philosophy in Robotics*

Robotics Institute  
Carnegie Mellon University  
Pittsburgh, Pennsylvania 15213

Thesis Committee:  
David Wettergreen, Chair  
Reid Simmons  
Jeff Schneider  
Steve Chien, Jet Propulsion Laboratory  
Phil Christensen, Arizona State University  
©David Thompson



## Acknowledgments

This thesis was only possible through guidance and support of countless others. Thanks are due to all of my friends at Carnegie Mellon who helped out with field experiments. Dom Jonak's knowledge of the rover, its hardware and software systems proved vital to our effort. Likewise, Francisco Calderón supplied his gifted intuition for hardware issues and a tireless work ethic. Even more than their technical abilities, their steadfast, easy-going companionship brightened our field expedition. James Teza lent a great deal of his personal time and effort in preparations for our field tests.

Our experiments relied on the expert knowledge of a group of Geologists at Arizona State University. Phil Christensen, Ron Greeley, and Shelby Cave were all vital as consultants in choosing and evaluating our field site. They humored us through a long site selection process and tolerated the eccentric requirements of the rover experiments. They supplied us with materials, advice, and technical information about multiple sites in the American Southwest. I feel strongly that in the future, the success of the autonomous science endeavor will depend on these close collaborations between researchers from the two sides of algorithm development and science applications.

The Bureau of Land Management (California Needles office) went through a great deal of trouble on our behalf to facilitate an expedited permit for work at our field site. My thanks go out to Rodney Mouton and his staff.

Thanks to Reid Simmons and Jeff Schneider, the committee members at Carnegie Mellon with whom I consulted throughout the past year. Their wisdom was vital to the progress of the thesis. Despite their busy schedules, they always found time to chat about my work, answer my deluge of questions, or provide advice.

Staff at the Caltech Jet Propulsion Laboratory offered lots of help. They were important, both as a source of new ideas and validation techniques, but also as a check to ensure that my work remained relevant to flight projects. Steve Chien has been extremely helpful and supportive in this regard. Many of the rock detection experiments were conducted under the supervision and support of Rebecca Castaño from the Machine Learning and Instrument Autonomy Group. My work with proxy features paired us with scientist consultant Bob Anderson. Kiri Wagstaff was a valuable sounding board for discussing science autonomy ideas.

Thanks to the other students at the Robotics Institute. Trey Smith was an invaluable mentor during my first few years at the institute. He taught me much of what I know about the research method — both the methodology, and the daily practice — by his example. My officemates Nathaniel Fairfield and Gil Jones were a great source of companionship during the thesis-writing process.

The benevolent mastermind behind the curtain is certainly my advisor, David Wettergreen. Throughout my studies and this thesis he has selflessly devoted himself to my success as a student and my development as a researcher. He's kept me challenged with new problems and well-supplied with the tools to solve them. I greatly admire David's vision for robotics and feel lucky to have served under his watch.

Finally, I owe a great gratitude to my family. My parents have been a sturdy support through the best and the worst episodes. It wouldn't have happened without them.

## II

This work was made possible through funding provided by a NASA ASTEP grant NNG0-4GB66G, “Science on the Fly,” and a JPL/Caltech Strategic University Partnership Grant.

## Abstract

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. These long traverses present new challenges to science-driven exploration. Rovers will travel over their local horizon so that scientists will not be able to specify targets in advance. Moreover, energy and time shortages will continue to limit the number of measurements so that sampling density will decrease as mobility improves. Finally, constraints on communications bandwidth will preclude transmitting most of the collected data. These issues raise the question: is it possible to explore efficiently, with long traverses and sparse sampling, without sacrificing our understanding of the visited terrain?

“Science autonomy” addresses the optimal sampling problem through onboard data understanding. Pattern recognition, learning and planning technologies can enable robots to place instruments and take measurements without human supervision. These robots can autonomously choose the most important features to observe and transmit.

This document argues that these agents should learn and exploit structure in the explored environment. In other words, they must be mapmakers. We advocate an *intelligent mapping* approach in which onboard predictive models represent spatial structure (similarities from one locale to the next) and cross-sensor structure (correlations between different sensing scales). These models guide the agent’s exploration to informative areas while minimizing redundant sampling.

The generative model allows us to formulate the exploration problem in terms of established principles of experimental design. Spatial experimental design criteria guide exploration decisions and suggest the best data products for downlink. In addition the map itself functions as a bandwidth-efficient representation of data gathered during the traverse.

This work bridges the gap between Bayesian experimental design, robotic mapping and their application in autonomous surficial geology. We develop generative data models that are appropriate for geologic mapping and site survey by planetary rovers. We present algorithms for learning map parameters on the fly and leveraging these contextual cues to choose optimal data collection and return actions. Finally we implement and test adaptive exploration schemes for kilometer-scale site survey tasks.



# Contents

<b>1</b>	<b>Introduction</b>	<b>1</b>
1.1	The Remote Exploration Problem . . . . .	2
1.2	Intelligent Mapping for Site Survey . . . . .	5
1.3	Contributions . . . . .	8
1.4	Thesis Overview . . . . .	9
<b>2</b>	<b>Feature Detection and Classification</b>	<b>11</b>
2.1	Previous Work in Geologic Feature Detection and Classification . . . . .	12
2.1.1	Previous Work in Geologic Feature Detection . . . . .	13
2.1.2	Previous work in Feature Classification . . . . .	17
2.2	Feature Detection . . . . .	19
2.2.1	Feature Detection Approach . . . . .	20
2.2.2	Feature Detection Performance . . . . .	25
2.3	Feature Tracking and Instrument Deployment . . . . .	31
2.3.1	Rover Hardware . . . . .	32
2.3.2	Feature Tracking Approach . . . . .	33
2.3.3	Visual Servo . . . . .	40
2.3.4	Feature Tracking Performance . . . . .	42
<b>3</b>	<b>Spatial Models</b>	<b>49</b>
3.1	Previous work in Mapmaking for Robotic Exploration . . . . .	50
3.2	Gaussian Process Background . . . . .	52
3.2.1	Inference with Covariance Functions . . . . .	55
3.2.2	Learning Hyperparameters . . . . .	57
3.3	Stationary and Nonstationary Environments . . . . .	58
3.3.1	Nonstationary Covariance Functions . . . . .	59
3.3.2	Augmented Input Spaces . . . . .	61
3.3.3	Evaluation on Synthetic Data . . . . .	64
3.4	Field Tests . . . . .	67
3.4.1	The Amboy Crater Lava Field . . . . .	68

3.4.2	Rover Instruments . . . . .	71
3.4.3	Procedure . . . . .	72
3.4.4	Results . . . . .	73
<b>4</b>	<b>Adaptive Data Collection</b>	<b>83</b>
4.1	Objective Functions for Remote Exploration . . . . .	84
4.1.1	Coupled and Decoupled Solutions . . . . .	85
4.2	Previous Work in Adaptive Data Collection . . . . .	87
4.2.1	Periodic Sampling . . . . .	88
4.2.2	Target Signatures . . . . .	89
4.2.3	Action Selection with Target Signatures . . . . .	91
4.2.4	Representative Sampling . . . . .	93
4.2.5	Novelty Detection . . . . .	93
4.2.6	Learning Mixed Reward Criteria . . . . .	94
4.2.7	Information-Driven Reward Functions . . . . .	95
4.2.8	Spatial Design . . . . .	98
4.2.9	Discussion . . . . .	100
4.3	Adaptive Sampling from a Transect . . . . .	102
4.4	The Corridor Exploration Task . . . . .	107
4.4.1	Problem Formulation . . . . .	109
4.4.2	Software Architecture and Execution Strategy . . . . .	115
4.4.3	Field Test Procedure . . . . .	116
4.4.4	Field Test Results . . . . .	119
<b>5</b>	<b>Selective Transmission with Proxy Features</b>	<b>125</b>
5.1	Proxy Features as Covariates in a Regression Model . . . . .	128
5.1.1	Shared Hyperparameter Learning . . . . .	129
5.1.2	Selective Return of Image Sequences . . . . .	134
5.2	Proxy Features as Noisy Observations . . . . .	142
5.3	Boundary detection with HMMs . . . . .	146
<b>6</b>	<b>Conclusions</b>	<b>153</b>
6.1	Contributions . . . . .	153
6.2	Future Work . . . . .	154
6.3	Maps and Exploration . . . . .	156

# Chapter 1

## Introduction

Since nature is too complex and varied to be explored at random, [the] map is as essential as observation and experiment to science's continuing development. - T. S. Kuhn, *The Structure of Scientific Revolutions* [1]

Maps are as old as exploration. For each culture that ventures beyond the safety of the village we find some distinctive new map: an assembly of sticks to represent ocean currents, a trail portrayed by scratches in bark, or a coastline sculpted in miniature [2]. Today the frontiers of exploration are the planetary bodies of our solar system. Much of what we know and theorize about them is best expressed by maps.

Maps are important to planetary scientists because they reveal the underlying structure of disparate observations. Figure 1.1 (left) illustrates this role. Here a geologic map shows several square kilometers of the Gusev crater region on Mars, identifying spatial trends and boundaries that indicate formative processes [3]. The map identifies regions of homogeneity and change, suggesting areas that warrant further exploration. The science map of Figure 1.1 (right) shows individual rocks within 10m of the Pathfinder mission landing site [4]. Its scale is radically different but its function is similar: to represent spatial relationships among simple features that reveal larger patterns and imply richer scientific interpretations.

Mapmaking, as the discovery and representation of structure in spatial data, is a vital tool for scientists and an apt metaphor for scientific exploration in general [1, 2]. Scientists seeking to understand a remote environment do not evaluate individual features in isolation but instead build a holistic account of many measurements and their distribution in space. The map embodies this process by situating each new observation in its relationship to the whole. Its coverage establishes the area of exploration, and the features themselves define the phenomena and interrelationships that the analysis will consider [5].

This document demonstrates that autonomous explorer agents can also be mapmakers. Explorer robots performing site survey face significant resource constraints; they must characterize large remote environments using very few observations. Onboard science maps form a framework by which these agents can learn, represent and exploit the structure of their environment to ex-

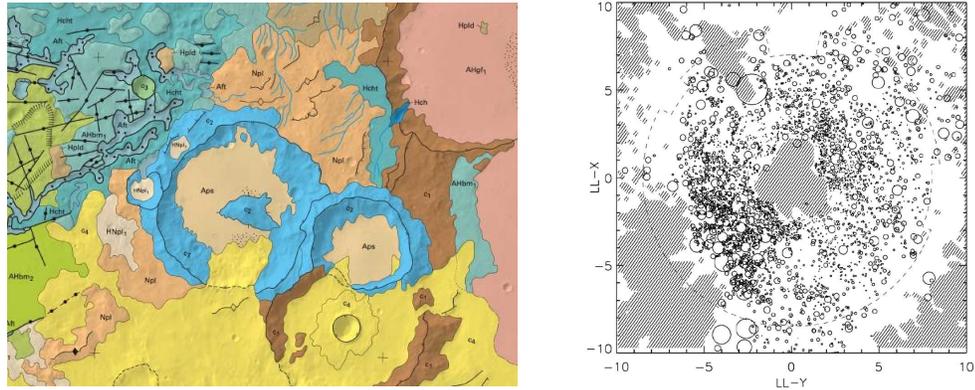


Figure 1.1: Left: A portion of the Gusev crater region on Mars. The colors and markings in this geologic map indicate predicted units of different surface material. Image courtesy USGS [3]. Right: Rocks at the Pathfinder landing site. Circle diameters indicate rock sizes. Shaded areas are unobserved due to occlusion. Image from Haldemann and Golombek [4].

plore more efficiently. Mapmaker agents can leverage trends in collected data to guide navigation, identify informative features for measurement, and minimize redundant science observations.

## 1.1 The Remote Exploration Problem

The investigation into autonomous mapmaking is motivated by continuing improvements in spacecraft lifespan and mobility. The recent history of Mars exploration exemplifies this trend (Figure 1.2). The immobile Viking lander preceded the 1997 Pathfinder mission that could travel within the field of view of the landing craft. By 2003 the Mars Exploration Rovers could travel tens of meters per day. The Mars Science Laboratory will build on these improvements with power and size to travel even greater distances. These rover missions generally travel within the terrain visible from the previous day's imagery.

New navigation technologies permit even longer single-cycle traverses of multiple kilometers. Recent tests have demonstrated robust kilometer-scale navigation at terrestrial analog sites [6, 7]. This portends a future generation of rovers that surpass the visible horizon on a single command, surveying terrain beyond what scientists have yet seen. New over-the-horizon operational modes will allow rovers to characterize large terrain areas and visit multiple geologic units [8]. They promise significant improvements in our understanding of Mars and the other planetary bodies of the solar system.

However, even as traverse distances increase, rovers' time and energy for taking measurements will remain mostly constant [9, 10]. In other words, robots may be able to survey large areas but they must still characterize the environment using just a few observations. Most of the terrain they visit may never be seen by humans. The survivability-driven concerns of the traditional lander

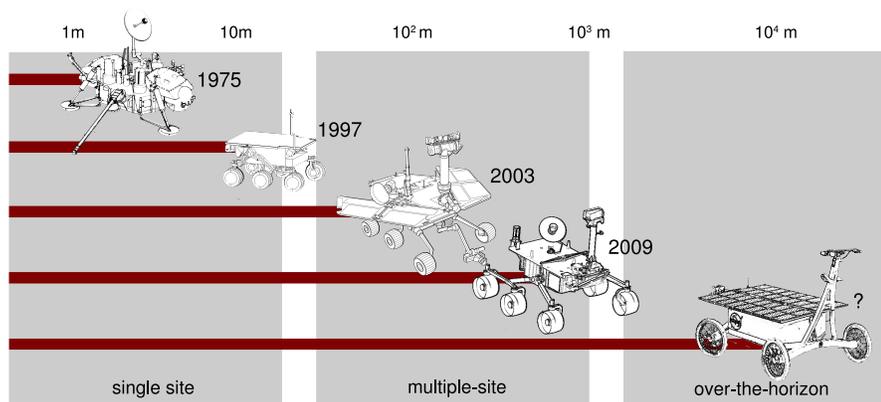


Figure 1.2: Operational modes for past and future Mars rovers are constrained by platform mobility. Improvements in autonomous navigation will soon permit new over-the-horizon traverse strategies in which rovers travel many kilometers in a single command cycle.

model, e.g. “how do we get there safely?” are complicated by new resource allocation challenges of mobile exploration: “where should we go?” and “what sites should we measure?” The long term trend in exploration missions show a steady advance from orbital flybys to stationary landers to survivable, mobile platforms [11]. Planetary exploration is entering a new era in which missions’ science yield is determined not just by the instruments we can deploy but also by our ability to allocate scarce data collection resources [10].

The communications downlink to Earth is another significant resource bottleneck. Figure 1.3 shows projected data transfer requirements of future NASA exploration missions. The current capabilities of the Deep Space Network (DSN) are inadequate for the projected volume of data. Even with “best-case” DSN upgrades, improvements in data acquisition and traverse length will continue to outpace bandwidth [11, 10, 9]. Therefore only a handful of the potential collected data products can be transmitted to scientists. Bandwidth limits already dominate science operations by the Mars Exploration Rovers; these platforms can now gather far more data than the daily downlink can accommodate [12]. Bandwidth concerns influence missions even to the design stage where presumed bottlenecks exclude valuable instruments like high-resolution cameras. Ultimately, time, energy and bandwidth constraints are not simply matters of degree or of improving the performance of an otherwise operable system. Rather, efficient resource allocation is a fundamental enabling technology — a basic requirement for the next generation of instruments and operational modes.

The challenges of resource-constrained exploration have created interest in onboard data analysis. Agents that understand the science content of collected data can adapt their behavior to explore more efficiently. They can recognize significant new science features and deploy sensors opportunistically. They can adapt to observed data, changing their exploration path on the fly to validate new trends or focus on sites that prove unusually interesting. They can selectively target the most

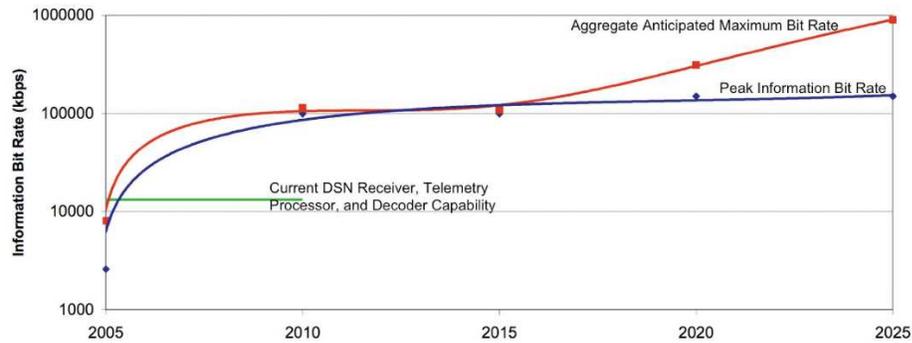


Figure 1.3: Projected requirements for data transmission on the Deep Space Network in future exploration missions. Image from [11].

time-expensive instruments. Finally, at the end of the command cycle they can make intelligent decisions about what data products to include in the downlink to Earth. Current rovers spend much of their time waiting for the next communications opportunity. In contrast, autonomous scientist agents could explore continuously in order to better utilize their finite operational lifespans.

Improvements in flight-ready computers suggest that improved onboard data understanding will be feasible for flight hardware. The performance of radiation-hardened electronics required for spaceflight lags behind terrestrial processors, but there is strong evidence that flight hardware follows a similar “Moore’s Law” development trend [13]. Extrapolating from the previous 20 years, flight hardware designers expect commensurate improvements in computing resources such as processor cycles, volatile and nonvolatile memory [13]. Designers can exploit these improvements with more sophisticated autonomous data analysis, trading computation for the time, energy, and bandwidth resources whose development follows more conservative trends.

Researchers have responded to this opportunity with a wide range of autonomous science systems. The simplest involve new software modules or procedures that serve singular well-defined objectives like automatic instrument placement [14, 15]. In these cases it is easy to identify appropriate behaviors, define the system’s role with respect to mission goals, and determine when it is functioning properly. At the opposite extreme, researchers speculate about a “robot geologist on Mars” that replicates some behavior of a human field geologist [16, 17]. Validation strategies for these more flexible autonomy systems are still unclear.

This work will focus on a case study involving geologic site survey by a surface rover. Here the issue of science value is already quite complicated. Even when a survey’s goals are enumerable (using suitably vague language), it may be impossible to score the value of particular data products. Scientists often change opinions about old data in light of new observations, methodologies, and even social factors [1]. They can only consider a small handful of working theories out of the infinite range of possible explanations for observations, so they can never say with certainty that some new data product is irrelevant. Indeed, scientists often disagree over the value of collected datasets [18].

These factors lead many to conclude that the only true standard for data's value is a subjective consensus of the community [1]. In light of this ambiguity, how do we specify appropriate data collection and transmission during autonomous site survey operations? Or more broadly, how can we quantify science value for autonomous resource allocation? Our solution is to restrict our definition of "science value," addressing a limited class of site survey tasks for which we can define specific quantifiable science goals.

## 1.2 Intelligent Mapping for Site Survey

We formulate exploration as a *mapping problem* guided by principles of spatial experimental design. Our objective is to discover the true state of the environment — to make a spatial model, or map, that reveals the ground-truth at the site with respect to some predefined phenomena of interest. It follows that the agent should collect and transmit observations that are most informative with respect to the spatial model. Exploration reduces to an active learning problem where data's value derives from the information it provides about the map. Remote scientists' reconstruction of the ground-truth phenomena from returned data comprises an objective standard to evaluate the efficacy of remote exploration.

Of course, a performance standard based on mapping accuracy ignores some aspects of remote exploration. Since the phenomena to be mapped are predefined we disregard the potential of novel features to inspire new interpretations or completely invalidate the model. In addition, our objective cannot account for situations in which scientists are searching in "treasure hunt" fashion for a single highly-desirable feature.

However, exploration-as-mapping does capture fundamental characteristics of geologic site survey. First, geologic samples' interpretations are sensitive to spatial context. Rather than consider each sample independently, field scientists incorporate each new observation into an evolving narrative of formation and alteration processes. This holistic interpretation involves large-scale spatial phenomena and the physical context of each sample is essential. For example, a single rock sample might be uninformative or revolutionary depending on the location at which it is found. It follows that robots performing site survey should represent the spatial structure of the explored environment. The model should capture phenomena like geologic units, borders and the general tendency toward spatial continuity in neighboring locations (Figure 1.4). By representing samples' locations in a map, the rover can capture this structure to identify regions of homogeneity and change. Sampling and transmission decisions can adapt appropriately to any spatial anomalies.

In addition, geologic interpretations are dependent on context provided by other sensing modes. Rather than considering only visual features, or compositional cues derived from spectroscopy, geologists' interpretations incorporate data from many sensors and multiple scales. A single site survey might consider *in situ* visual inspection, but also field and laboratory spectroscopy and large-area remote sensing data. It follows that the exploration agent's map must also represent cross-sensor relationships in the data. It should permit inference between orbital and surface data products. By incorporating these dependencies into the model we ensure that information value

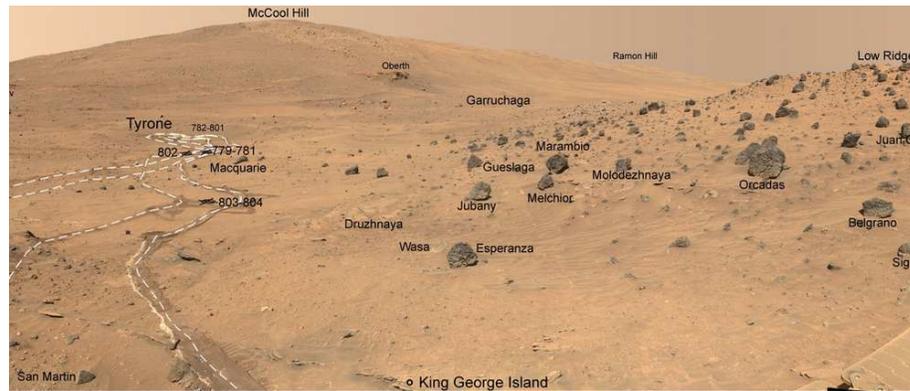


Figure 1.4: A typical MER panorama shows an environment with significant spatial trends and structure. Image courtesy NASA JPL / Caltech / Cornell.

judgments reflect the actual utility of a data product as judged by human geologists.

Ultimately we find that the mapping metaphor is a good match for most site survey tasks. Moreover, it distills a nebulous question of science value to the concrete issues of choosing the phenomena of interest and an appropriate statistical model relating them to observations. It yields a clear division of labor between the scientist users (who choose the phenomena of interest and devise a spatial model) and the remote agent (that tries mechanistically to reduce uncertainty over free parameters). We advocate a strategy in which the agent learns model parameters on the fly and exploits predictions from previous data to inform adaptive exploration decisions. The onboard model makes predictions about unobserved locations, so we describe our approach with the term *intelligent mapping*.

To summarize, our strategy treats exploration as adaptive experimental design with a predictive map that can incorporate multiple sensors and scales. The agent learns this model on the fly and uses it to compute the probability of new observations at candidate locations in the explored or unexplored space. The map can represent spatial structure (similarities from one locale to the next - Figure 1.5 left) and inter-sensor structure (correlations between different sensing modes - Figure 1.5 right). Optimal data collection actions maximize an experimental design criterion: in general, the expected information gain with respect to the map parameters of interest. We maximize this reward subject to cost penalties for collecting data or hard constraints on the allowable observations. A similar criterion identifies the most informative data products for downlink to scientists.

Intelligent maps offer both theoretical and practical advantages for site survey. The predictive model lets us quantify science value in terms of experimental design principles so that preferred actions follow logically from the variables of interest, prior observations and modeling assumptions. It promises testable performance improvements relative to mapping strategies based on coverage metrics. Maps alleviate resource bottlenecks by improving exploration robots' ability to collect science data autonomously. Maps can identify high value features for opportunistic data collection

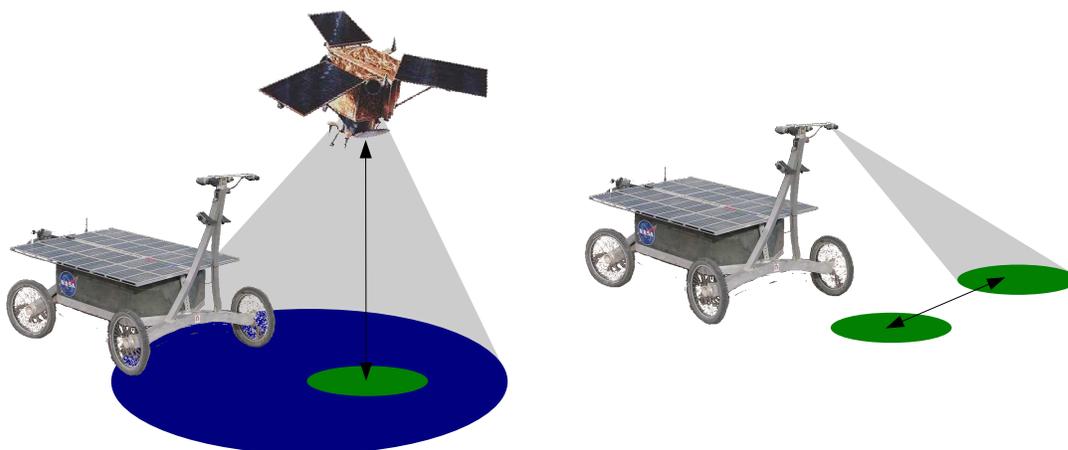


Figure 1.5: Left: A cartoon illustration of cross-sensor inference. A surface rover considering a new measurement can infer likely observations by exploiting learned correlations with preexisting orbital images. Right: A rover uses spatial inference, leveraging the continuity among neighbors to identify redundant samples.

and guide the explorer adaptively toward informative areas. At the end of the command cycle, the map helps the agent select the best data products for transmission to Earth. The map itself is a valuable data product; it summarizes hundreds or thousands of features in a highly compressed form. For a given resource budget, intelligent maps can improve the fidelity of the onboard map and the reconstruction inferred by remote scientists. We advocate the following conclusion:

*Autonomous mapping of science phenomena can benefit remote site survey under resource constraints.*

We support this thesis with experiments demonstrating specific site survey tasks in which autonomous maps of science phenomena improve performance. Under these circumstances, generative spatial models — and mapping more generally — benefit remote investigations.

First, we show new map-based data products that are informative, bandwidth-efficient descriptions of long traverses. We enable this operations scenario with detection and classification techniques for identifying geologic features of interest, and by demonstrating tracking of multiple features to form maps of features' distributions. This shows that mapping discrete science phenomena autonomously is feasible for real-time traverse. Ultimately these maps can provide more informative descriptions of the environment for a given bandwidth allowance.

Additionally, we show how adaptive decision making based on predictive spatial models can improve reconstructions of the ground-truth science phenomena. We present generative models of spatial data that synthesize observations from images and spectroscopy, and devise algorithms for computing favorable data collection and transmission actions based on these models. Then, we test a rover exploration system where the remote agent learns this map of the environment on

the fly. Experiments demonstrate improved reconstruction of ground-truth measurements based on the returned data products. This work demonstrates predictive mapping for autonomous geology applications, and validates it as a method for more general exploration problems.

### 1.3 Contributions

The thesis will demonstrate adaptive data collection by an autonomous remote explorer, showing that automatic feature detection and instrument targeting is possible for long traverse site survey operations. Here we describe a rover system using a common instrument designed for geologic field work - a Visible Near-Infrared (VIS/NIR) reflectance spectrometer. The rover is capable of detecting rocks in the environment and collecting spectral measurements from distant targets without any human intervention. The rover can resolve the correspondence problem to identify unique rocks and locate their absolute positions in a map of the explored environment. The result is a novel data product — the “spectrometer profile” — relevant to exploration plans that mix autonomous and scripted activities.

The thesis will refine existing spatial models to make them relevant to planetary site survey applications. We pay special attention to a classic model from the geostatistics literature: the Kriging model, also known as a Gaussian process [19, 20]. We present modifications of the classic Gaussian process formulation to take advantage of the spatial and cross-sensor correlations in geologic site survey, and present efficient methods for learning their parameters on the fly. In addition we will evaluate the fidelity of various models for mapping geologic phenomena in field data consisting of both surface and orbital measurements.

The Gaussian process describes similarity relationships between different locations in the explored environment. By formalizing the adaptive exploration problem in terms of active learning, we demonstrate how these relationships (the “covariance structure”) influence the agent’s adaptive sampling and return behavior. This favors different sampling strategies based on prior assumptions about the model, the environment itself, and the fidelity of feature detection. In this manner we unify a wide range of existing heuristic approaches to remote exploration. Ad hoc strategies like periodic and representative sampling, boundary following and coverage patterns, are revealed to be special cases of our broader formulation. The heuristic methods imply specific, testable assumptions about the explored environment and the fidelity of the agent’s onboard data analysis. By formalizing them we create guidelines for the application of such heuristics in future systems.

We show how adaptive exploration can affect exploration efficiency in field conditions. A series of tests at Amboy Crater, California demonstrate intelligent maps for kilometer-scale site survey by a mobile rover platform. Here the robot constructs geologic maps automatically based on remote sensing data and surface measurements from a Visible Near Infrared (VIS/NIR) Spectrometer. To our knowledge, the Amboy experiments yield the first real-time autonomously constructed geologic surficial maps. Our approach introduces geologic mapping using multiscale data, as well as efficient on-line learning of spatial and cross-sensor correlations.

Finally, we present a novel approach to selective data transmission that leverages the intelligent

map to produce informative downlinks even when the agent cannot explicitly detect the features of scientific interest. The use of “proxy features” subsumes a range of heuristic alternatives that prevail in contemporary science autonomy literature.

## 1.4 Thesis Overview

There are three conceptual components implicit in the design of an intelligent map system (Figure 1.6). The *Feature detection and classification* component detects discrete semantically-meaningful science features in raw input data. The second component, the *spatial model*, is a keystone of the intelligent map approach. It describes spatial and cross-sensor relationships between science features and permits inference about unobserved locations. Finally, the *action selection* component leverages model inference to improve data collection and transmission decisions.

There are several reasons we bother with this tripartite description. Conventionally each stage employs a distinct set of tools, such as computer vision and pattern recognition for feature detection and planning for action selection. These component systems may be designed and evaluated separately. Another reason for the three-part description is the convention of the data product (i.e. individual image or spectrum) as the basic element of collected data. Data products are scheduled, collected, and transmitted as atomic units; they are usually encoded as independent files in the rover’s computer. This suggests a natural distinction between the feature detection stage that analyzes data products independently, the modeling stage that discovers relationships between sets of data products, and action selection that supervises collection of new data products. In this work we will treat the components separately whenever such a division is reasonable.

Chapter 2 introduces the feature detection and classification component, with special focus on the issue of rock detection. We will review previous work in the area and present a new method for rock detection that facilitates real-time recognition of rocks by a rover in field terrain and lighting conditions. Finally, we will leverage rock detection to facilitate automatic spectra collection. The result is a real-time map of rock spectra encountered during a traverse.

Chapter 3 discusses spatial models, moving from characterizing independent features to modeling whole environments. We describe Gaussian process models in greater detail and evaluate their effectiveness at modeling synthetic and field data.

Chapter 4 deals with action selection. We begin by discussing the issues associated with quantifying samples’ science value. We review heuristic strategies like periodic sampling and more complex adaptive strategies based on set preferences and information gain metrics. We propose the use of data collection criteria based on principles of information theory. Finally, we present Amboy Crater experiments demonstrating information-driven data collection in the field. A rover exploring under time constraints demonstrates adaptive traverse planning in response to orbital data and in-situ spectroscopy.

These preliminary investigations presume that the agent enjoys the same perception capability as the scientist. Chapter 5 describes selective data return in more realistic remote science scenarios where artificial agent has an incomplete or noisy view of science phenomena. We address the

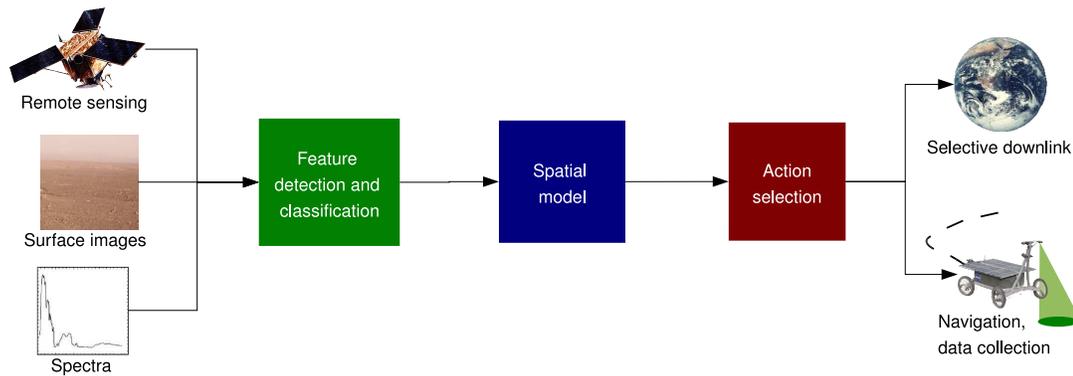


Figure 1.6: Three components of the proposed system for geologic site survey.

discrepancy with explicit *proxy features* — correlated, noisy observations that stand in for the science content of interest. Here, the intelligent map permits agents to identify favorable downlink transmissions without detecting science content directly.

Finally, chapter 6 closes with a discussion of other applications of the intelligent map approach and suggests domain areas for future research.

## Chapter 2

# Feature Detection and Classification

Another precept is that everything relating both to bodies and virtues in nature be set forth (as far as may be) numbered, weighed, measured, defined. - Francis Bacon, *Preparative toward a Natural and Experimental History*, 1620 [21]

The first stage of an intelligent mapping system recognizes discrete features of interest in the explored environment. The *feature detection* step analyzes raw instrument data to identify discrete, semantically-meaningful science phenomena. Feature detection can contribute in several ways to the science return of a command cycle. The map of features itself constitutes a bandwidth-efficient summary of the traverse — a more complete description of the environment than might be gleaned from sparse subsets of the raw data products. Moreover, because the features become atomic elements of the map, feature detection facilitates the spatial models that are central to our adaptive resource allocation.

Figure 2.1 portrays feature detection's place in the broader intelligent mapping approach. The component analyzes raw data products and produces a list of unique features along with associated positions and geologic attributes. Mapping these science features poses several technical challenges. The most interesting phenomena are invariably subtle, lying at the limits of detectability. Thus, detecting these science features in the first place is a difficult *pattern recognition problem*. In addition, the agent must also localize detected features physically in the environment. This requires solving a tracking or *correspondence problem*; the agent must recognize repeat instances of features that it has seen previously. The agent must solve a *sensor fusion problem* of associating data from multiple sensors. Finally, the agent may need to process this raw data to derive a science interpretation. For example, it could interpret spectra in order to infer a rock's origins and material composition. This constitutes a *classification problem* in which the agent categorizes features using their associated attributes.

Here we demonstrate a rover system that exhibits these foundational capabilities. Our objective is to demonstrate the feasibility of feature detection and classification for geologic site survey on a field platform. We focus on a specific case study involving mapping rocks during geologic site survey. We demonstrate a rover performing autonomous feature detection and classification

using image analysis, targeting rocks for further analysis with a visible Near-Infrared (VIS/NIR) Spectrometer.

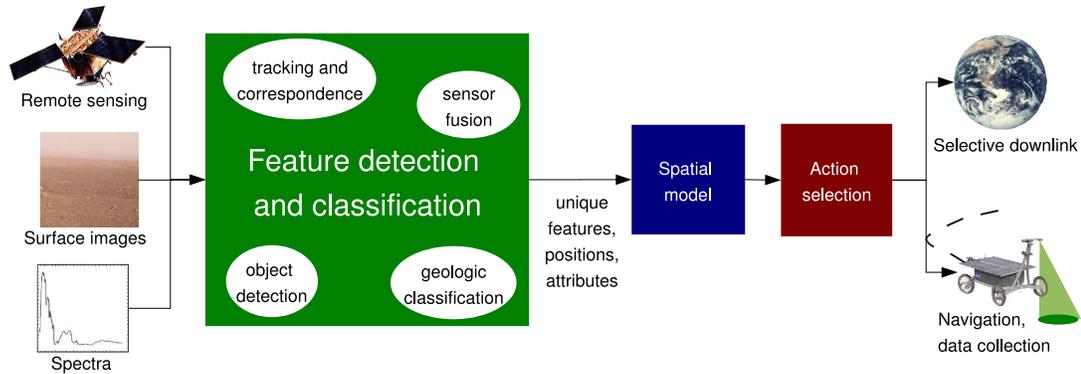


Figure 2.1: Feature detection and classification involves identifying and classifying unique science features of the explored environment.

We will begin by surveying previous work in autonomous rover geology, with specific focus on the example of rock detection and classification. Previous research has produced wide range of previous rock detection methods with distinctive performance characteristics. We will present a new method for rock detection in images: an illumination-sensitive method based on modifications to the classic “Viola-Jones” object recognition strategy. We will evaluate its performance against other real-time rock detection algorithms, showing that it exhibits performance that is competitive with contemporary alternatives.

Next, field tests will integrate rock detection with autonomous targeting of a Visible Near-Infrared Spectrometer. We present a new operational mode - a “spectrometer transect profile” that resolves the data fusion and correspondence problems to facilitate autonomous mapping involving dozens of target features. The spectroscopy system demonstrates the ability to detect, place, and classify large numbers of features on the fly. Compared with previous efforts, it offers significant improvements in the quantity of spectra collected per unit time.

## 2.1 Previous Work in Geologic Feature Detection and Classification

During the last decade onboard geologic data analysis has evolved from a hypothetical possibility into a flight-proven strategy for improving science data return. Tests have validated automatic science data analysis against ever-larger datasets; both analog studies and actual missions have shown improvements in returned data’s science value.

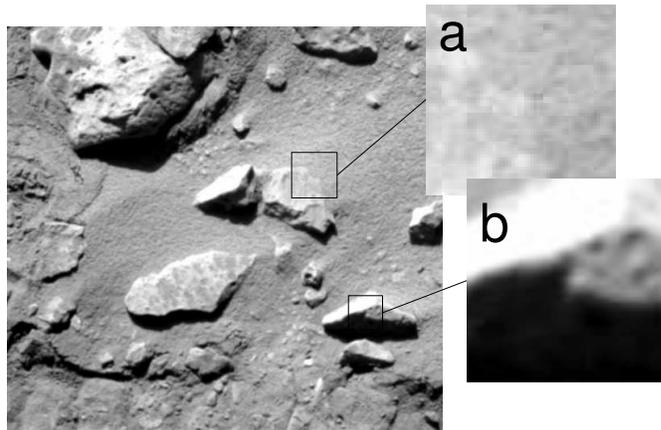


Figure 2.2: Rocks’ diverse textures, variable morphologies, and directional illumination all complicate the segmentation problem. Local cues are insufficient: real borders are often weak (a), while the rocks themselves are non-uniform with false interior borders (b). Image from the Mars Exploration Rover “Spirit,” courtesy of NASA / JPL Caltech / Cornell.

### 2.1.1 Previous Work in Geologic Feature Detection

In 1999 Gulick *et al* conducted tests of the “Marsokhod” rover at Silver Lake, California. Analyzing the data from these images, they broached issues such as selective collection and return that would pervade future science autonomy efforts [22, 9]. The researchers recognized that mobile autonomous rovers would benefit from onboard data understanding. They deferred the issues of appropriate science value functions and focused instead on the initial problem of geologic feature detection in images.

The Marsokhod tests demonstrate autonomous detection of three specific features of geologic interest: horizons, layers, and rocks. Horizon detection is useful for geology because a rover could exclude the sky from its terrain analysis routines. In addition, it could selectively compress the sky at a higher rate for bandwidth savings. The Marsokhod horizon detector uses an adaptive contour algorithm, similar to the classic “snakes” method for contour identification, that relaxes an initial guess in order to minimize an energy function. Over several iterations the relaxation aligns the detected horizon to edge pixels in the image.

The Marsokhod tests also demonstrate detection of layering in rock strata. Layers are significant because they suggest geologic history; it follows that images containing layering are likely to be especially interesting to scientists. The Marsokhod layer detector relies on an edge detection procedure followed by the identification of parallel lines in the image. Filtering operations recover contiguous regions with a common edge orientation.

Rock detection deserves special attention. The identification of observed rocks is an important task in geologic interpretation of a site. Rocks are particularly informative features: their shapes

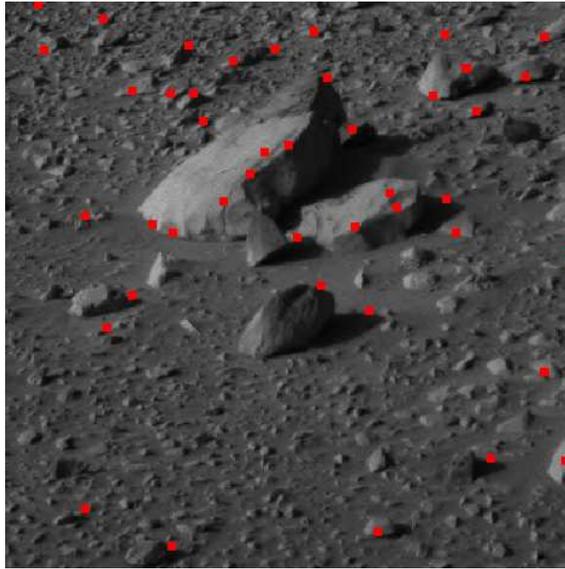


Figure 2.3: Rock detection on a Mars Exploration Rover Pancam image using the Marsokhod algorithm. The algorithm identifies points on the terminator lines between shaded and sunlit portions of the rock. Image from Thompson and Castaño [24]

carry information about physical processes in the environment and they are useful as targets for spectroscopic analysis. They are pervasive in planetary exploration environments. They are also a useful test case for developing and testing science autonomy: as discrete targets, rocks offer a well-defined detection and classification task that simplifies validation.

Segmenting rocks is difficult because they are non-uniform: their texture, color and albedo varies across their surfaces and from one rock to the next [23]. Planetary analog environments often exhibit strong directional lighting with cast shadows and highlights that violate the uniformity assumption. One must infer weak boundary edges from context (Figure 2.2). The Marsokhod algorithm finds rocks by searching for terminator lines between light and shaded areas. Edge detection combined with a heuristic color comparison identifies terminator lines. Then a spherical illumination model infers a rock's location from each detected shadow (Figure 2.3). These early tests are generally off-line, qualitative analyses; they do not yet integrate the demonstrated feature detection with onboard decision-making, nor do they attempt to quantify detection performance. However, the initial work suggests that a rover might autonomously identify some of the same visual phenomena geologists consider during field studies.

The following year the Robotic Antarctic Meteorite Search (RAMS) developed a robot that autonomously searched for meteorites on an Antarctic plateau [25, 26]. The Elephant Moraine is an ice plateau covered sparsely by rocks, a substantial fraction of which are meteorites. Here the Nomad platform uses cameras to segment rocks from the background ice [26]. Here the ice

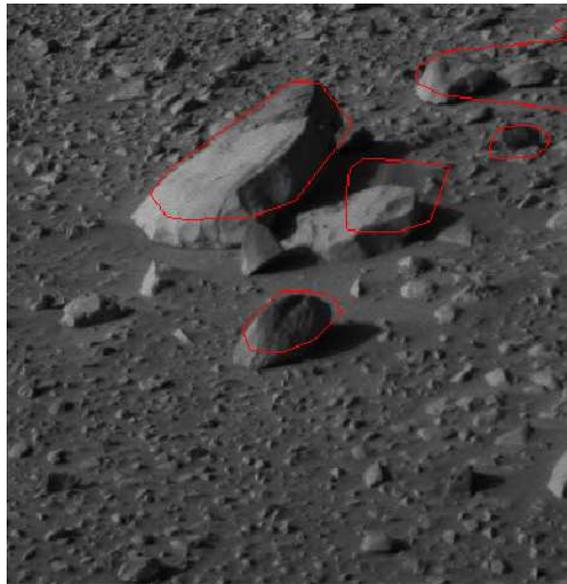


Figure 2.4: Rock detection on a Mars Exploration Rover Pancam image using stereo. Stereo methods excel in detecting the largest rocks on flat terrain. Image from Thompson and Castaño [24]

surface exposed the dark rocks and simplified the pattern recognition task. Nevertheless, the system provided another important field demonstration of autonomous feature detection.

Later work demonstrates still other rock detection methods, including stereo range-based techniques of Fox *et al.* and Pedersen *et al.* [27, 28]. These fit a simple parametric model to the terrain in order to identify large rocks that protrude above the surface. The algorithm of Fox *et al.* is typical of these approaches (Figure 2.4). It first aligns a planar model to the ground plane using a robust (outlier-resistant) least-squares regression method. Identifying the ground plane transforms a standard stereo range map into a height map; local maxima correspond to peaks of protruding objects.

The Fox algorithm utilizes a region-growing method to expand these peaks and grow pixel regions until they reach a minimum height. The result is a series of pixel regions that correspond to rocks. An advantage of these methods is their reliability when the ground is highly planar and the rocks are well-separated. Because they utilize 3D geometry, they are robust to shadows and other changes in lighting. However, practical exploration scenarios often involve undulating terrain or partially-embedded rocks.

Other detection strategies look for closed contours in the image itself. *Rockfinder* is an edge-based rock detection algorithm developed by the OASIS autonomous science group at the Jet Propulsion Laboratory [29, 12, 30] (Figure 2.5). The detection technique uses a series of denoising and bilateral filter operations to remove spurious texture. It applies an edge detector (the Sobel or

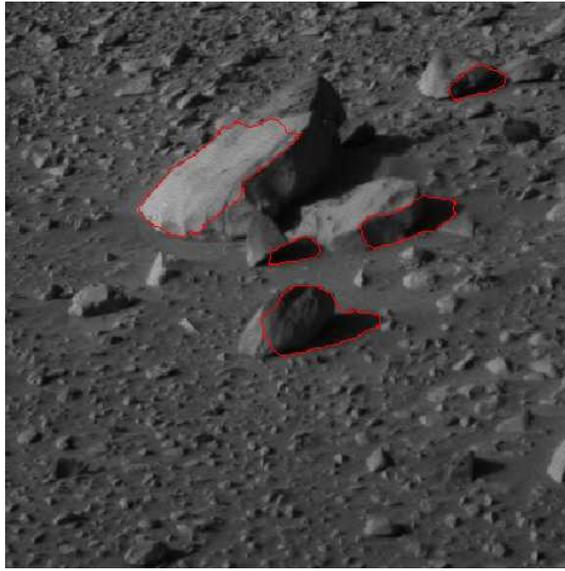


Figure 2.5: Rock detection on a Mars Exploration Rover Pancam image using the OASIS Rockfinder algorithm. Image from Thompson and Castaño [24]

Canny algorithm) to find edge fragments and traces the resulting pieces into closed contours. These closed shapes are unlikely to appear coincidentally in the Mars environment, and the vast majority of closed edge shapes lie on the surface of rocks or are rocks themselves. The algorithm runs at multiple levels of a multi-scale image pyramid in order to detect rocks of many sizes.

*Rockster* [12], an edge-based algorithm originally developed for synthetic images, has shown promise for many natural datasets. It identifies and removes image pixels above the horizon, and then detects edges using a method similar to the Canny algorithm. A sequence of edge cleaning operations finds closed contours. These steps include breaking edges at points of high curvature, connecting endpoints with nearby edges to form T-junctions, and a gap-filling stage to bridge missing contour segments. The final result is a set of closed shapes which can be extracted using a flood-fill technique (Figure 2.6). *Rockster* was successfully demonstrated on the FIDO platform in 2007 when it detected rocks during an integrated Mars Yard test of autonomous science operations.

Thompson *et al.* demonstrate learning-based approaches to the object recognition problem. A region-growing algorithm segments the scene into many fragments. Then a probabilistic classification model based on a Bayes network [31] distinguishes rock from non-rock regions [32, 33]. The Bayes network is a graphical model representing probability relationships between features of regions and the classification as rock or non-rock. These relationships are sparse, permitting efficient inference of new classes based on relationships discovered in hand-labeled training data. Another learning-based approach appears in Dunlop *et al.*. Here a support vector machine trained on lighting, texture, and shape features evaluates candidate segmentations of the scene [23]. Candidates are

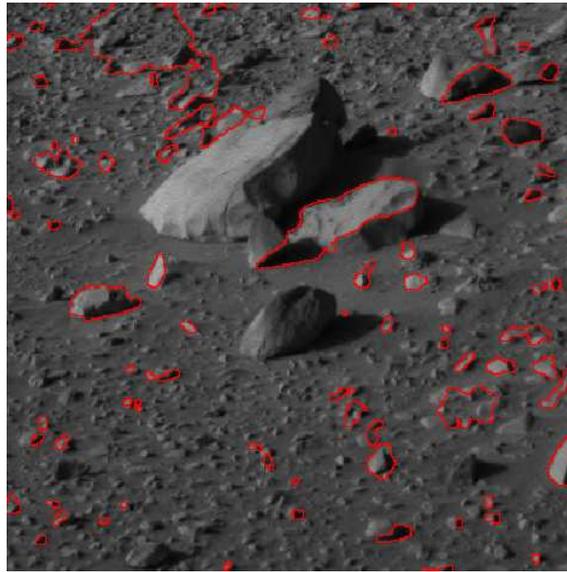


Figure 2.6: Rock detection on a Mars Exploration Rover Pancam image using the rockster algorithm. Image from Thompson and Castaño [24]

produced by assembling configurations of hundreds of image subregions called “superpixels.” A search through the space of combinations finds high-scoring segmentations that match the learned appearances of rocks and soils. Because this search is computationally intensive, the algorithm is not feasible for real-time use.

### 2.1.2 Previous work in Feature Classification

“Feature classification” involves extracting quantitative descriptors that characterize the features and permit meaningful science decisions. These descriptors range from simple visual attributes such as albedo to formal geologic classifications like calibrated spectroscopic analysis for determining composition.

An archetypal example from field geology is the analysis of rock particles’ shape and morphology. Here, “particles” can refer to rocks of any size — from large boulders down to the grains that compose the base sediment. In each case shape is important for tracking statistical trends in post-emplacement processes such as weathering and transport. Shape is an important complement to spectrometer-driven compositional studies because it can characterize these post-emplacement processes. Many previous geomorphological studies focus on the distribution of the sizes, shapes, and textures of individual particles. Particles’ edges become more eroded during transport as they collide against each other. Therefore angularity is a strong indicator of the processes that produced and transported them, and morphology is important in determining their plausible origins and the impact of fluvial, aeolian, or volcanic processes to their history. The distribution of particle shapes

and sizes in a given area of terrain constitutes a signature that can distinguish it from regions shaped by alternative processes.

Terrestrial analog studies reveal the value of particle size and shape analysis. Field studies by Craddock *et al.* investigate a number of Hawaiian lava flows modified by known processes such as chemical weathering, aggradational processes, and depositing by phreatic eruptions (analogous to impact cratering) [34]. Analysis of these Mars-analog sites indicates strong correlations between the dominant modification process at a site and the observed distributions of particle sizes and cross-sectional particle areas. These distinctions rely on large amounts of quantitative particle size and shape data. Numerous studies have employed particle morphology as a cue to reveal subtle trends in Mars Exploration Rover imagery [35, 36, 37]. Automatic methods for determining shape are useful not just for recognizing outlier samples during autonomous traverse, but also for assisting geologists with analysis of returned images.

Several efforts have attempted to automate this process through quantitative measures of particle morphology. Fox *et al.* investigate several visual features using quantitative shape measures for rocks' contours [27]. These shape measures include descriptors of size, sphericity and the contour's deviation from a best-fit ellipse. Dunlop expands this analysis with computational estimates of roundness and sphericity [38]. Her analysis includes not only traditional shape descriptors from computational geometry, but also geologic shape and roundness descriptors such as the Diepenbrock roundness and Crofts chart comparison [39]. She compares the estimates of human geologists viewing images to the ground-truth measurements produced by computational techniques. Given an accurate segmentation these roundness estimates are more accurate and consistent than those of human geologists (Figure 2.7).

Particle morphology is only one of many possible geologic classifications. Another vitally important characteristic is rocks' composition and mineralogy. The Robotic Antarctic Meteorite Search used vision and spectroscopy data to perform a formal geologic classification of mineral samples. After detecting a rock the rover approached and deployed an arm-mounted spectrometer using visual servoing. This measured each sample along several preselected wavelengths. The robot also carried a second instrument to characterize rocks using their magnetic signature. A Bayesian belief network trained on samples from a spectrum library incorporated both sensors to classify rocks and identify meteorites.

Additional studies beyond the RAMS project demonstrate autonomous interpretation of spectral signatures. Bornstein *et al.* [40] describe an automatic neural network carbonate detection system for planetary rovers. Automatic pixel-wise classification of multispectral images is common in remote sensing studies; of particular interest is the Earth Observer One, or EO-1 autonomous science project that brought automatic feature classification to an orbital platform [41]. EO-1's on-board data analysis classified pixel intensities from a collection of spectral bands using a Support Vector Machine classifier; it can respond to features such as thermal anomalies, clouds, flooding, and changes between images caused by transient geophysical phenomena [41]. These classification decisions trigger dynamic selective data return and retargeting actions.

The diversity of different feature detection and classification algorithms poses a difficulty for building integrated science autonomy systems. Estlin *et al.* [30] address the problem with a soft-

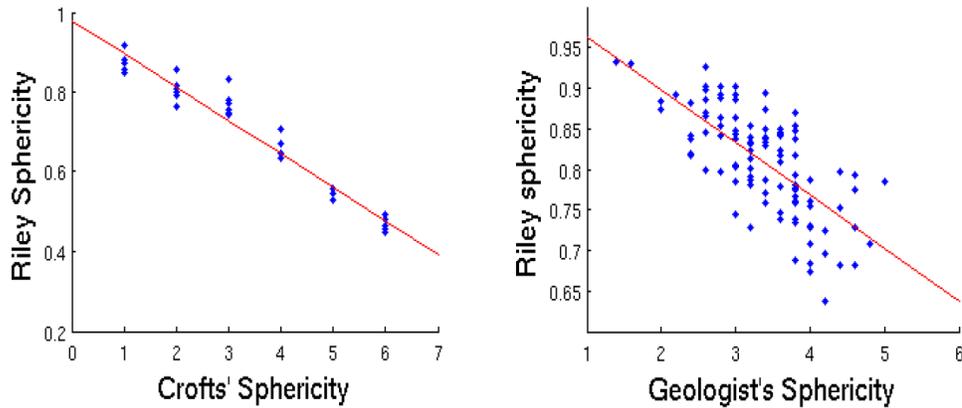


Figure 2.7: In Dunlop’s tests, geologists’ estimates of rock sphericity from inspecting images are less consistent in general than computational estimates. The Riley sphericity measure, which can be extracted automatically from the images, correlates strongly with ground-truth Crofts’ sphericity values. Human geologists’ Crofts value estimates exhibit greater variance. Image from Dunlop [38].

ware architecture that makes it easier to add additional, specific feature detectors for new phenomena of interest. This method takes a modular approach with many different feature detection and classification modules that can be used interchangeably in an autonomous science architecture. They propose a “pipe and filter” architecture with different modules for feature detection, classification, and target signature matching. A common interface for these services eases the development burden for new science features.

In summary, autonomous geology research has investigated a wide range of feature detection techniques. Most draw from research in computer vision and pattern recognition. Research into feature detection strategies is slowly maturing from proof-of-concept work to an established discipline characterized by formal comparative benchmarks and validation on large test sets. Meanwhile autonomous image analysis has seen its first use onboard planetary exploration platforms.

## 2.2 Feature Detection

Here we introduce a new feature detection mechanism, a *Modified Viola/Jones* (MVJ) detector using a template-based approach to rock detection. The algorithm utilizes a modified version of the template cascade detection strategy developed by Viola and Jones [42, 43]. A supervised learning scheme builds a sequence of classifier modules linked together in a series known as a “cascade.” This sequence of simple decisions has the net effect of distinguishing template windows containing rocks from those that do not. The surviving template windows constitute approximate bounding boxes that can then be used for spectrometer target selection or additional image pro-

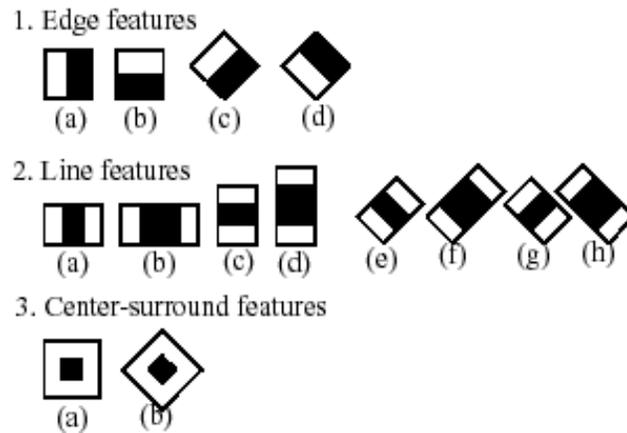


Figure 2.8: Haar-like wavelets used to construct the Adaboost filter cascade. From Bradski *et al* [43].

cessing. Our modified version uses several classifier cascades to account for different potential lighting conditions. The result is a system that is discriminative yet fast enough for real-time rover deployment. Performance evaluations on Mars Exploration Rover datasets suggest that the MVJ detection method is competitive with existing approaches.

### 2.2.1 Feature Detection Approach

The Viola/Jones detection method [42] is an appearance-based approach that places bounding boxes around the desired objects. It evaluates all potential bounding boxes in the image at various scales. For each candidate it classifies the bounding box according to whether or not it is likely to contain a rock. The exhaustive search requires that the evaluation procedure be computationally efficient; a major advance in the Viola/Jones approach is the use of advanced caching techniques to speed computation for large images.

The detector first convolves each image with a preselected cascade of Haar-wavelet-like filters (Figure 2.8). This yields a rich feature vector of filter responses with size  $n \times m$ , where  $n$  is the number of filters and  $m$  is the number of pixel positions within the candidate bounding box. We identify a small discriminative subset of these values to form a “weak classifier” that discriminates between rock and nonrock bounding boxes. In our case the weak classifier consists of a binary decision tree containing anywhere from 1 to 20 levels [43].

One trains the classifier using supervised learning on a set of manually labeled data (Figure 2.9). The algorithm places bounding boxes in these training images to produce sets of “object” and “non-object” regions. A learning algorithm chooses filters, bounding box locations, and response thresholds that best distinguish the objects in the training set.

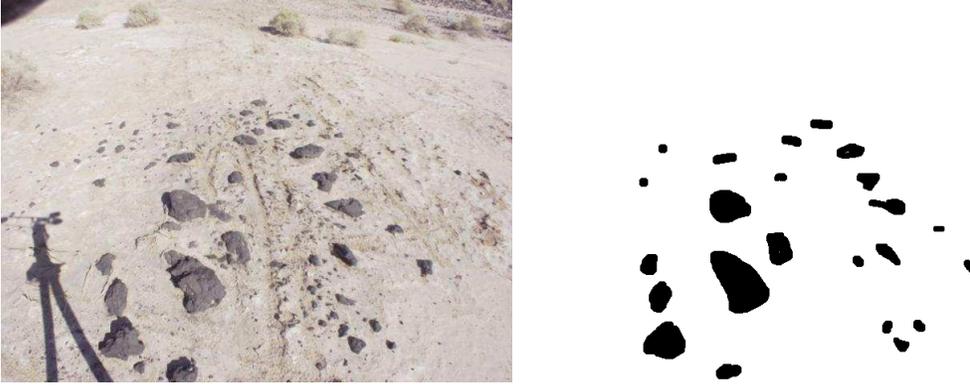


Figure 2.9: Typical rock detector training image and hand-labeled object masks.

The traditional Viola / Jones algorithm proceeds as follows. Presume a set of weighted training points with binary labels  $y_i \in \{0, 1\}$  and associated with an importance weight  $w_i \in [0, 1]$ . The error for a weak classifier  $h_j$  on a candidate bounding box  $x_i$  is given by the sum of weighted classification errors  $h_j(x_i)$  over the training set [44]:

$$\text{err}(h_j) = \sum_i w_i |h_j(x_i) - y_i| \quad (2.1)$$

Each filter alone is a poor discriminator, but one can recover subtle distinctions by using multiple filters simultaneously. The Adaboost training algorithm is a natural choice for this feature selection task [45]. Adaboost chooses the best weak classifier to add as a new feature, and re-weights the training set according to the following update rule:

$$w_{t+1,i} = w_{t,i} \beta_t^{1-e_i} \quad (2.2)$$

Here  $e_i = 0$  for correctly-classified datapoints. For incorrectly-classified datapoints,  $e_i = 1$  and  $\beta_t = \frac{\text{err}_t}{1-\text{err}_t}$ . Weights are normalized to yield a probability distribution, and the next feature is added to the feature set. The features combine to yield a final “strong classifier” according to the decision rule:

$$h(x) = \begin{cases} 1 & \text{if } \sum_t \log \frac{1}{\beta_t} h_t(x) \geq \frac{1}{2} \sum_t \log \frac{1}{\beta_t} \\ 0 & \text{otherwise} \end{cases} \quad (2.3)$$

The strong classifier exhibits better performance than any individual filter.

Viola and Jones achieve a considerable speed improvement by chaining together multiple strong classifiers together into an “attentional cascade” that requires only a few filter evaluations at each step (Figure 2.10). Each stage in the sequence has the option of rejecting the bounding box as a non-object. Since each boosted classifier rejects the majority of candidate bounding boxes it only passes on a small subset to the next stage for further processing by the other classifiers.

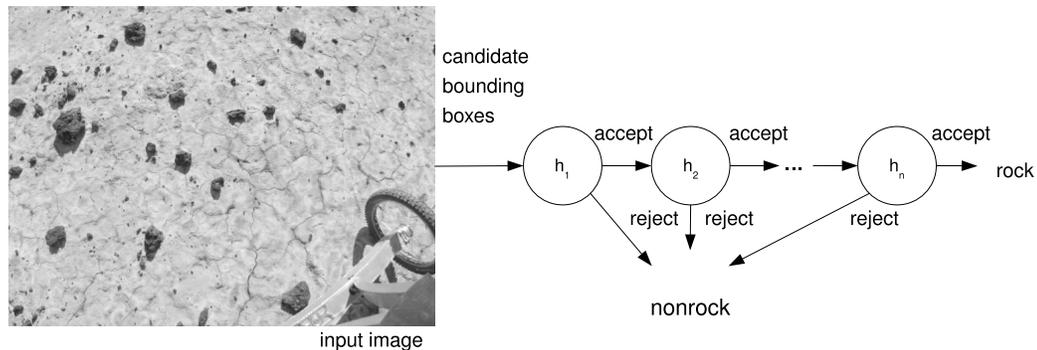


Figure 2.10: The original Viola / Jones (MVJ) detector. It consists of a cascade of boosted “strong learners”, each of which rejects a portion of candidate bounding boxes as “non-object.” Detections consist of bounding boxes that survive to the end of the cascade.

The detection result becomes progressively stricter with each additional stage that is added to the cascade (Figure 2.11).

Candidate objects that survive the decision pass to the next stage in the chain; the detection results are those bounding boxes that survive rejection by all the classifiers. Our cascade uses a sequence of 25 stages, so each individual stage must be quite conservative and reject very few true positives. We set the decision threshold to ensure a training set false-negative rate of no more than 99.5% for each classifier in the sequence. The training algorithm adds new classifiers to the sequence in a “greedy” ordering, favoring those filters that best discriminate among the candidates passed on from the preceding stage.

While classifying all the potential bounding boxes in the image seems a daunting task, several speed advantages of the Viola/Jones method make real-time operations possible. First, the filter values themselves are simple to compute using a caching trick known as the “integral image.” The integral image is an image transform that sums the pixel values to the left and top of each pixel in the image; one can compute it quickly in constant time. The integral image makes it easy to compute the convolution result of any uniform block aligned with the pixel directions. Such operations reduce to integer addition operations using the four corner values of the convolution mask in the integral image [42]. Since our filters are composed of these uniform blocks, computing filter responses is a fast operation. We derive another speed advantage from the attentional cascade; the first few stages of the filter cascade reject the large majority of bounding boxes, so we pay little computational cost for a long detector cascade.

Assuming the training data captured this diversity, the traditional Viola / Jones approach would fit a single decision boundary that separates all possible rock bounding boxes from all possible nonrock bounding boxes. However, this naïve approach ignores important information about image

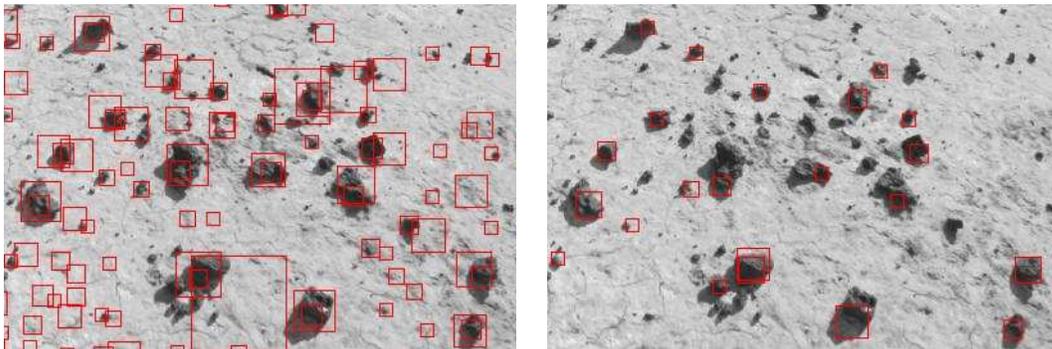


Figure 2.11: The cascade of classifiers rejects a portion of the remaining candidates at each stage. These images show lenient detections (Left) resulting from a depth-5 cascade. A depth-25 cascade results in a much stricter detection result (Right).

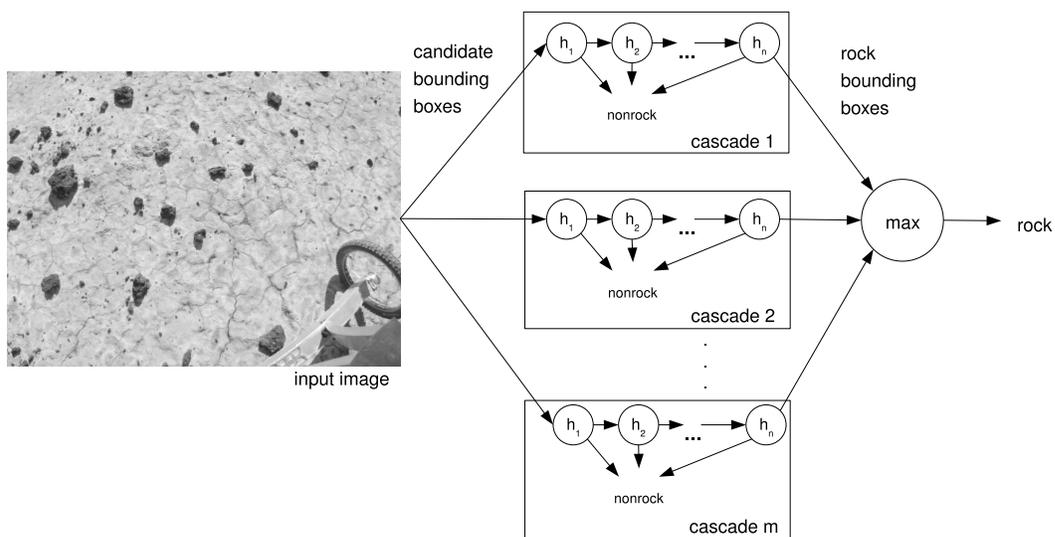


Figure 2.12: The proposed Modified Viola / Jones (MVJ) detector. It consists of a parallel series of filter cascades, each of which uses a chain-structured decision tree to identify rocks from the set of all candidate bounding boxes in the image. However, each cascade presumes a different lighting angle.

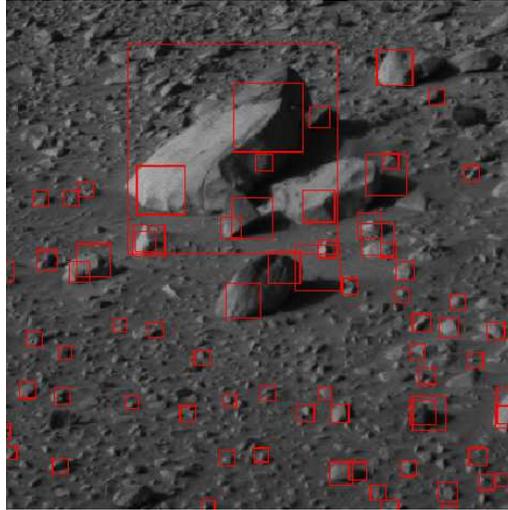


Figure 2.13: The MVJ detector applied to the example Mars Exploration Rover image. The detector successfully identifies the left-lit model as the correct lighting to explain the scene; the detection result reflects the characteristic asymmetry of highlight and shadow on the rocks' surface. From Thompson and Castaño [24].

illumination. Figure 2.2 evidences that features' appearances can vary significantly depending on scene lighting. Thus the distribution of rock appearances is highly multimodal, making the decision surface far more complicated than for a single lighting direction. Detecting all lighting directions simultaneously ignores the useful constraint that the scene lighting be constant throughout the image.

Our proposed MVJ detector addresses these shortcomings by modifying the classic Viola/Jones strategy. Rather than train a single filter cascade we use multiple cascades to recognize rocks under different lighting directions. We manually group training data into categories based on the lighting angle in the scene. The number of categories can be quite small: using two sides or four quadrants both work well in practice. We train a separate filter cascade on each lighting direction. At runtime the multiple cascades compete to interpret the scene. The system detects rocks once using each cascade. We take the cascade which returns the most detections to be that which best explains the scene illumination, and use its detection result as the final solution for the novel image.

The cascades learn to identify the characteristic shading patterns of a rock lit from a particular direction. These shading patterns seldom occur by accident so detection cascades associated with incorrect lighting conditions return few detections. The system uses only the results from the cascade that returns the greatest number of detections (Figure 2.12). In this manner a multiple cascade strategy identifies an approximate lighting direction even when the sun angle is not known in advance. In the trials that follow we use two detection cascades to account for rocks lit from the left and right sides respectively. As an added advantage we can effectively double the quantity

of training data by flipping each of the training images horizontally. Note that when the lighting angle is known with certainty, either through ephemeris calculations based on rover pose or direct observation through a sun-tracking instrument [46], the rover can use this information to choose the appropriate cascade directly.

The concept of multiple parallel object detectors may have implications beyond variable lighting conditions. In principle they can improve detection performance for any confounding phenomenon that changes from one image to the next but is constant throughout each image. For instance, it could improve performance on tasks with varying illumination levels and soil conditions.

Note that the detector cascades will occasionally identify more than one candidate bounding box on the surface of a single rock. We compensate by eliminating “double detections” detections which overlap by more than 50% of the smaller bounding boxes’ area. In cases where this overlap occurs we disregard the smaller bounding box.

### 2.2.2 Feature Detection Performance

Here we present a performance comparison of the Modified Viola-Jones detector against the range of existing detection mechanisms we presented in the first section. We consider system performance on a set of representative images from the Mars Exploration Rover catalog. The first test set contained of 104 panoramic camera images taken by the Spirit Mars Exploration Rover (MER). We selected a range of mast elevation angles but favored image content within a 3-10m range of the rover. This range was deemed the most relevant for autonomous geology operations. Together the images comprised a dataset containing over 35,000 hand-labeled rocks. They also included a selection of typical non-rock content such as rover tracks, shadows and sediment features. We drew images from two sites: the “Mission Success” panorama from Spirit’s landing site and the “Legacy” panorama imaged during Sols 59 through 61. The images were acquired using Spirit’s 753nm bandpass filter.

Particle size analysis usually excludes the smallest rocks that are too numerous for accurate labeling. Similarly, we limited our analysis to those rocks greater than 4cm in length because they could be consistently identified across all datasets. We reserved an additional set of 12 Pancam images from Spirit rover imagery to serve as training data. These contained several thousand rocks from different locations along the Spirit mission path; like the test set they exhibited a range of terrain and lighting conditions.

In addition to the panoramic camera imagery we considered 13 MER navigation camera images from the Spirit rover. These images contained over 13,000 hand-labeled rocks. We used Navcam images from three sites: initial images from the landing site, a panorama at sol 50, and a second panorama from sol 118. These monochrome images provided a wider field of view including both near-field objects and the horizon. They also showed a wider range of lighting conditions than the panoramic imagery; several were taken late in the Martian day with low-angle illumination and significant cast shadows. Because relatively few labeled Navcam images were available we did not reserve any as a training set. Instead, we used the Pancam training set mentioned above for all tests

involving Mars rover images. This compromise worked well for the specific detectors and datasets we considered. Again, we ignored rocks smaller than  $4\text{cm}$  in size.

It is common to evaluate pattern recognition systems according to a precision/recall profile, where *precision* refers to the fraction of detections that are actually true rocks and *recall* describes the fraction of rocks in the scene that are detected. The designer can favor increased precision or recall by altering the confidence threshold for accepting a detected rock. In our case many of the algorithms — such as the *Rockster*, *Rockfinder*, and shadow methods — lack a scalar measure of detection confidence that would define a precision-recall curve. Instead, these algorithms usually exhibit an obvious performance peak that strongly suggests a particular parameter setting [24]. Parameter tuning can modify this behavior to move outside this “performance knee,” but the resulting behavior does not automatically result in higher precision or recall. More often these changes simply hurt overall performance and stability by all performance measures. Thus, instead of producing a complete precision-recall *curve* we compare each method’s behavior in the parameter envelope where it would be used in practice. We use a training set to manually tune each algorithm to the knee of the performance curve, and use these parameter settings on the other datasets in the performance tests.

The algorithms exhibit markedly different performance profiles. Some favor large rocks, others small; some attempt to provide accurate boundary contours for detected rocks, while others simply provide bounding boxes. The appropriate evaluation metric ultimately depends on the functions that are most important for a given autonomous science application. We will consider several performance measures to evaluate test set performance for a range of autonomous geology tasks. First we evaluate a detector’s ability to perform target selection for applications like autonomous spectroscopy. This is tantamount to finding pixels that are most likely to lie on the surface of a rock. For algorithms that provided an outline or bounding box we used the centroid of the detected region as the target estimate. The relevant performance characteristic is accuracy - the fraction of chosen target pixels that actually lie on rock regions.

Next we evaluate each detector’s detection precision and recall. This score is more relevant to those mapping tasks where we hope to identify *all* the unique rocks in the environment with as few false positives as possible. To compute this score we should find a unique association so that each ground-truth rock matches with only one detection result. We find the area of overlap between detected and ground-truth rock regions, and greedily associate each ground-truth rock region with the detection that overlaps by the highest percentage of its area. We match each ground-truth rock with at most one detection and label unmatched rocks as false positives.

To compute a precision/recall score each valid match must satisfy a correspondence criterion to ensure similarity between the detection and its associated ground truth rock. Performance scores are highly dependent on this correspondence standard so we evaluate two possibilities: a *permissive criterion* and a *strict criterion*. The permissive criterion requires that more than 50% of the detected region contain the matched ground-truth rock. This guarantees that a majority of pixels in the detected region lie on the rock. However, it is forgiving of situations where the detector finds only part of the target. We also investigate a strict criterion that requires a 50% overlap between the regions’ intersection and their union. The strict criterion demands a close correspondence between

the area of the true rock and the detected region, so it better indicates the accuracy of automatically-computed rock attributes like size or shape. Examples of each appear at right in Figures 2.14 and 2.15.

Performance for Navcam and Pancam images is similar for all algorithms and tests with one exception — a range data registration error prevented us from evaluating the stereo algorithm on the Pancam dataset. Therefore the explanation that follows favors Navcam imagery except in cases where the Pancam dataset’s additional data is particularly informative.

Figure 2.16 shows target selection performance for navigation imagery. The Rockfinder algorithm has the best overall target selection performance; it reliably achieves accuracies over 90%. Failure modes differ for each detection algorithm. The Rockfinder algorithm occasionally mistakes long cast shadows for rocks. The MVJ algorithm performs poorly in lighting conditions that differ from its training set (i.e. illumination from directly behind or in front of the camera). The stereo algorithm under-performs for small rocks or uneven terrain. Nevertheless, all rock detectors routinely score accuracies of 60–70% on the target selection task. In fact, these values may understate the performance of algorithms like Rockster and MVJ whose profuse detections could be filtered using size or shape heuristics.

Note that these precision-recall scores do not reflect different algorithms’ abilities to recover the exact contour boundaries of detected rocks. In fact, the MVJ method returns only bounding boxes and is not capable of recovering these contours in the first place. However, as previously noted, precise contour boundaries could be useful for extracting precise shape information. These analyses would require a different performance measure such as the Chamfer distance approach of Dunlop *et al* that measures the similarities between detected and ground-truth contours [38]. In our case, we are concerned with accurate spectrometer targeting for which the centroid and correspondence matching tests better indicate the expected system performance.

Figure 2.18 shows average performance on Navcam images for rocks in four size categories. The permissive detection criterion was used. Note that the MVJ algorithm’s performance increases with rock size, while the reverse is true for Rockster. Rockfinder and the stereo algorithm both perform poorly on rocks less than 10cm in length. This is expected in the case of stereo that relies on protrusion above the ground plane. However, there is no obvious mechanism that would cause a size bias in an edge-based method like Rockfinder. This algorithm regularly achieves 90% precision on medium to large rocks. Figure 2.19 shows performance on Navcam images plotted by range categories. Most detectors show improved detection precision in the areas near to the rover. Not only do distant rocks subtend fewer pixels, but they also cluster more closely due to camera foreshortening.

Figure 2.14 shows the same detection results for Navcam imagery on an image-by-image basis. This provides some intuition about the variance in performance from one image to the next. Again, both Stereo and Rockfinder algorithms exhibit high precision and low recall while the Rockfinder algorithm is the overall winner in precision. The Rockster algorithm achieves the highest overall recall for this test set at the cost of considerable inter-image variance. Finally, Figure 2.15 shows Navcam detection results using the strict correspondence criterion. A comparison with figure 2.14 suggests that few detections actually capture the target rock’s true contour. The MVJ algorithm

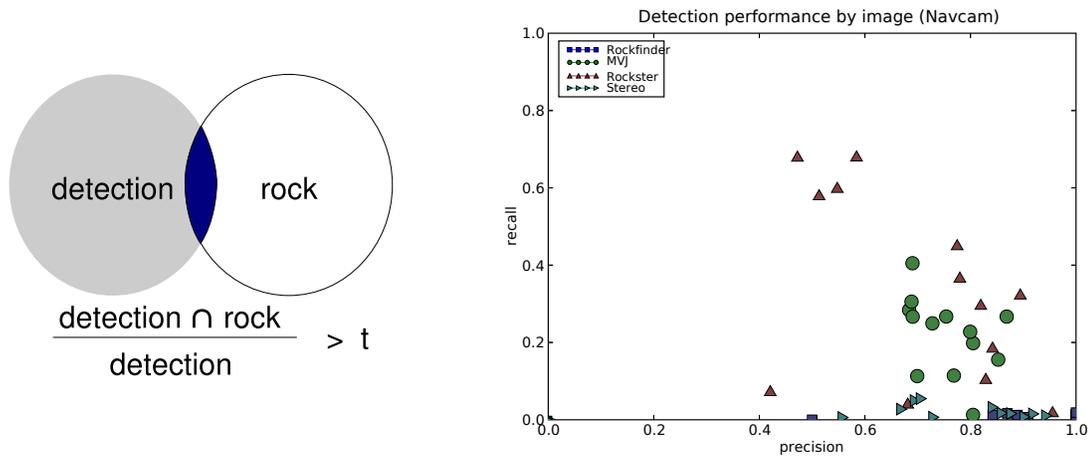


Figure 2.14: Left: Permissive criterion for rock detection. The ratio defined by the areas of intersection between detected contours and ground-truth rocks must be greater than a fixed threshold  $t$ . The correspondence criterion permits at most one detection per rock. Right: Rock detection performance for Navcam imagery using the permissive criterion. Points represent the precision/recall result for each individual image.

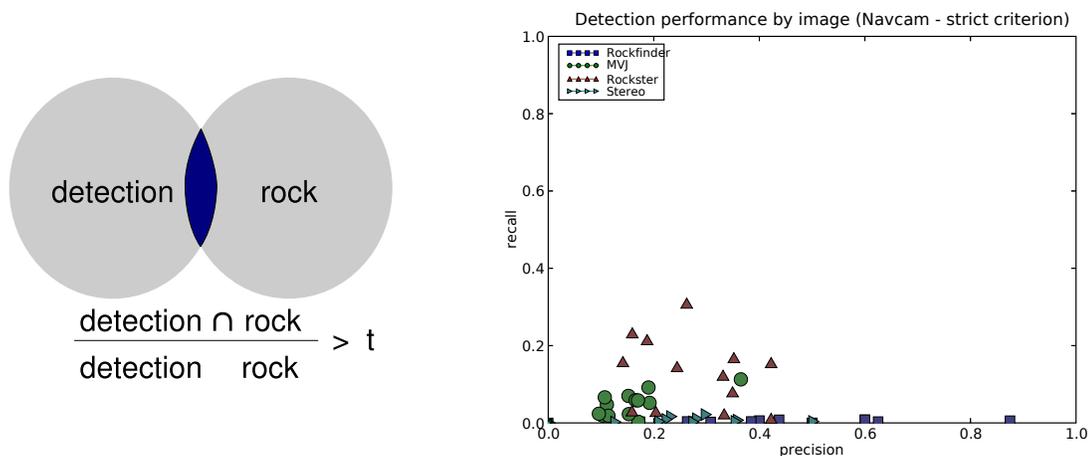


Figure 2.15: Left: Strict criterion for rock detection. The ratio defined by the areas of intersection between detected contours and ground-truth rocks must be greater than the union. The correspondence criterion permits at most one detection per rock. Right: Rock detection performance for Navcam imagery using the strict criterion. A close match between detected and ground-truth boundaries is difficult so strict correspondence results in worse (apparent) performance.

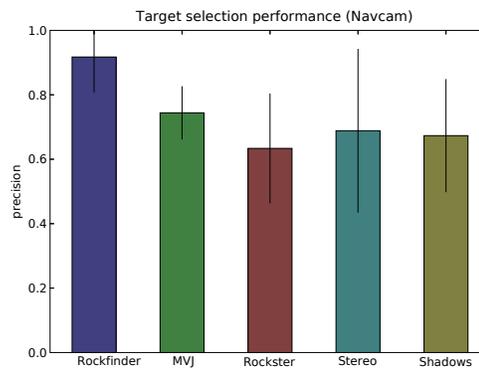


Figure 2.16: Target selection performance for Navcam imagery. The vertical axis represents the fraction of detected centroids that lie on a rock pixel.

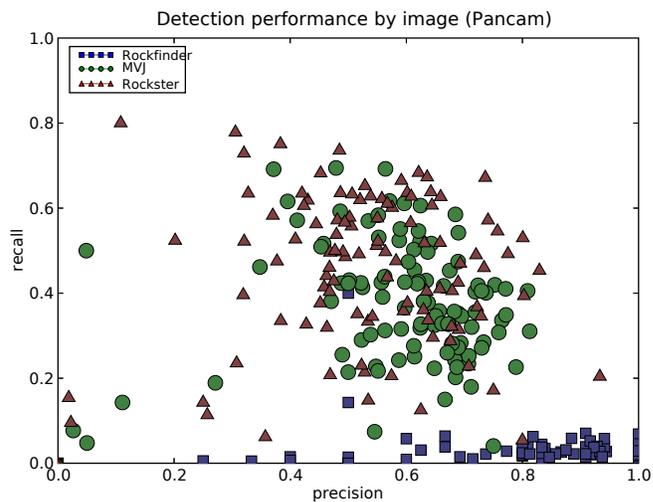


Figure 2.17: Rock detection performance for Pancam imagery using the permissive criterion. Performance is generally comparable to that for Navcam images, with MVJ and Rockster trading precision for recall.

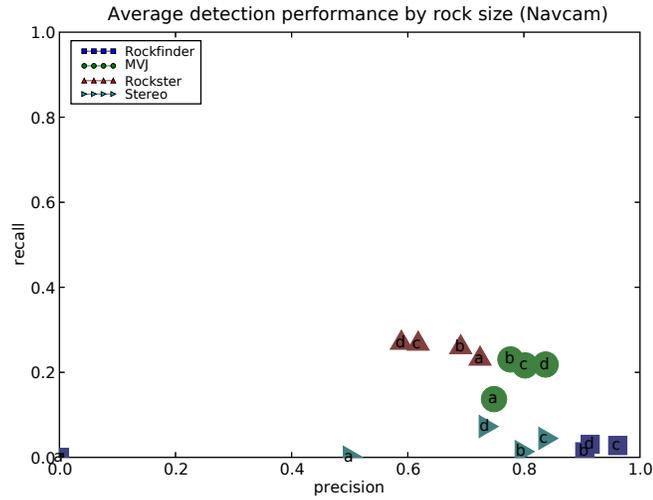


Figure 2.18: Performance by size for Navigation imagery. Points represent the average detection performance over the entire dataset for rocks in each size bin: a) < 10cm, b) 10 – 20cm, c) 20 – 30cm, d) > 30cm.

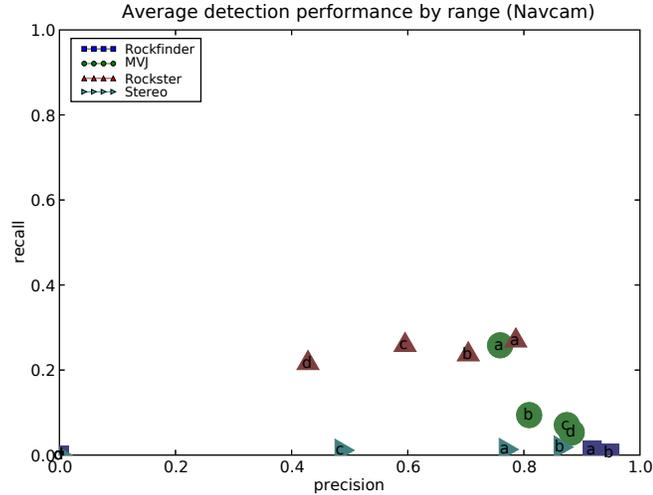


Figure 2.19: Performance by range for Navigation imagery. Points represent the average detection performance over the entire dataset for rocks in each range bin: a) < 2m, b) 2 – 4m, c) 4 – 6m, d) > 6m.



Figure 2.20: The “rock garden” at Amboy Crater is the site of the traverse experiments described in this chapter. Here the rover begins a spectrometer profile.

only detects bounding boxes and not contours, so its performance score suffers most from the stricter correspondence standard.

In summary, the new MVJ detector offers a good balance of features for autonomous spectroscopy. It is fast enough for real-time use. Its recall rate is near parity with the best available, while its precision exceeds 80% on the Mars Exploration Rover datasets we considered. Finally, it performs well for the large, close rocks that are the best targets for autonomous spectrum acquisition. In short, it has a high precision but also finds sufficient targets to “keep the spectrometer busy” for tasks involving multiple-spectrum acquisition.

## 2.3 Feature Tracking and Instrument Deployment

Here we present a new operational mode that leverages the MVJ rock detection to improve data collection during extended site survey operations. A new data product — a *spectrometer profile* — is a sampling of spectra collected from detected features at intervals along a rover transect. Field tests use a VIS/NIR Reflectance Spectrometer to acquire data from detected features at distance under solar illumination. This offers a spectroscopy acquisition rate significantly higher than any previously reported autonomous method. The system’s main innovation is the integration of rock detection into a mapping system capable of tracking and recognizing dozens of unique features. This demonstrates the feasibility of feature detection and classification in the intelligent mapping approach. We address the correspondence and association problems first broached at the beginning of the chapter using a target matching scheme based on scene geometry and motion-invariant keypoint locations of the image.

In contrast, previous research has examined automatic spectroscopy in the context of Single-Cycle Instrument Placement (SCIP) [47, 48, 49]. In the SCIP approach a human operator specifies each target in a panoramic image. The rover advances toward the feature with visual tracking, and

then performs a “hand-off” to instrument cameras for fine tracking and arm placement [47, 48]. Pedersen et al. extend this model to multiple targets [14]; their system demonstrated high-accuracy instrument placement on 3 targets in under 3 hours.

Our work differs from previous efforts in several respects. First, we focus on reflectance spectroscopy under solar illumination which obviates the need for slow approach and arm placement maneuvers. Additionally, instead of reacting to a single prespecified target our algorithm autonomously detects and localizes multiple target rocks simultaneously. The autonomous spectrum collection has a relatively high failure rate compared to supervised SCIP methods. However, since there is no significant recovery penalty these errors are not problematic. Thanks to the integrated rock detection, the rover can always move on to a new target and initiate another collection sequence. Over time it can acquire spectra from dozens of rocks, resulting in a spectrum map that provides a VIS/NIR profile of the transect.

Our tracking solution uses stable image points known as Scale-Invariant Feature Transform (SIFT) descriptors [50] to characterize previously-appearing rocks. These descriptors represent key locations that correspond to particular physical locations and whose appearance is preserved across minor viewpoint changes. This permits correspondence matching, yielding a database with a unique position estimate for each science target. The robot periodically pauses its forward motion, performs visual servoing to aim the spectrometer, and collects measurements from the most promising targets. Our approach allows a robot to collect dozens or potentially hundreds of spectrometer measurements in a single command cycle with no direct human intervention.

### 2.3.1 Rover Hardware

Here we demonstrate the spectrometer profile algorithm on an exploration robot “Zoë” (Figure 2.20). Zoë is used at Carnegie Mellon University to test mission scenarios involving long-range navigation [51, 7]. The rover platform is 2 meters in length and can travel up to 1 meter per second under solar power. A dual passive axle design steers the vehicle by driving the wheels independently at different velocities. This produces variable-radius drive arcs while at the same time accommodating obstacles up to a 20 centimeters in height.

Mast-mounted navigation cameras provide stereo vision for obstacle avoidance and terrain evaluation based on local slope and roughness. Onboard navigation software uses these terrain evaluations along with a D\* path planning algorithm [52] to identify the best paths between prespecified waypoints. Our software architecture treats terrain-based navigation separately from science-driven path planning. Autonomous science software gives each path segment in the form of a single end waypoint, and the terrain-based navigation algorithm attempts to reach a goal area within 5 meters of this location. This does not guarantee the rover will follow a straight-line path to the waypoint, as it must detour to avoid any obstacles or rough terrain that appears.

Zoë’s localization strategy emulates planetary exploration scenarios. We initialize rover position through differential GPS to simulate manual landmark-based techniques for finding accurate start-of-day positions. During the traverse the rover updates its position estimate with dead reckoning. Planetary rovers can derive absolute heading data from celestial and solar navigation systems



Figure 2.21: The spectrometer's foreoptic mounts rigidly to the stereo rig atop Zoë's pan tilt unit.

utilizing ephemeris calculations; an onboard magnetic compass simulates this capability with periodic heading corrections. Some localization error persists due to wheel slip; depending on the terrain this error ranges from 0 to 5% of distance traveled.

Zoë's science instruments include a 60cm-baseline stereo rig and a VIS/NIR 350 – 2500nm reflectance spectrometer. The spectrometer objective lens, or *foreoptic*, is mounted with the cameras on a pan-tilt base that provides full 360° coverage of the environment. The foreoptic's field of view forms a 1° sensing cone in space. Placing a target into this cone allows the spectrometer to collect reflectance spectra under solar illumination.

### 2.3.2 Feature Tracking Approach

The operator initiates the spectrometer profile mode by specifying a transect start location, a goal waypoint, and a time budget. During execution the rover drives from the start location to the goal, producing a map of detected rocks and spectra — the *spectrometer profile* of science features. It detects rocks in images, localizes them using groundplane and stereo methods, and if time remains, pauses to target the spectrometer using an iterative visual servoing method. Because the mapping procedure enforces a time constraint one could integrate it as an atomic command into a larger mission plan with other drive or sampling actions.

Figure 2.23 shows the software architecture developed for the spectrometer profiling task. Our architecture uses three independently-operating software modules for the tasks of data acquisition, image analysis, and navigation. In addition, a database of detected rocks records key information about each detected feature. This rock database is central to the system and acts as a medium of communication between the other components. We implement this structure with a relational database. It associates the following information with each rock:

- An integer used as a globally unique identifier.
- A global 3D position estimate given in Earth-Centered-Earth-Fixed (ECEF) coordinates.



Figure 2.22: Zoë at the Amboy crater site. The rover travels between waypoints toward the goal while collecting spectra of rocks. Actual field tests use a traverse length of 50m.

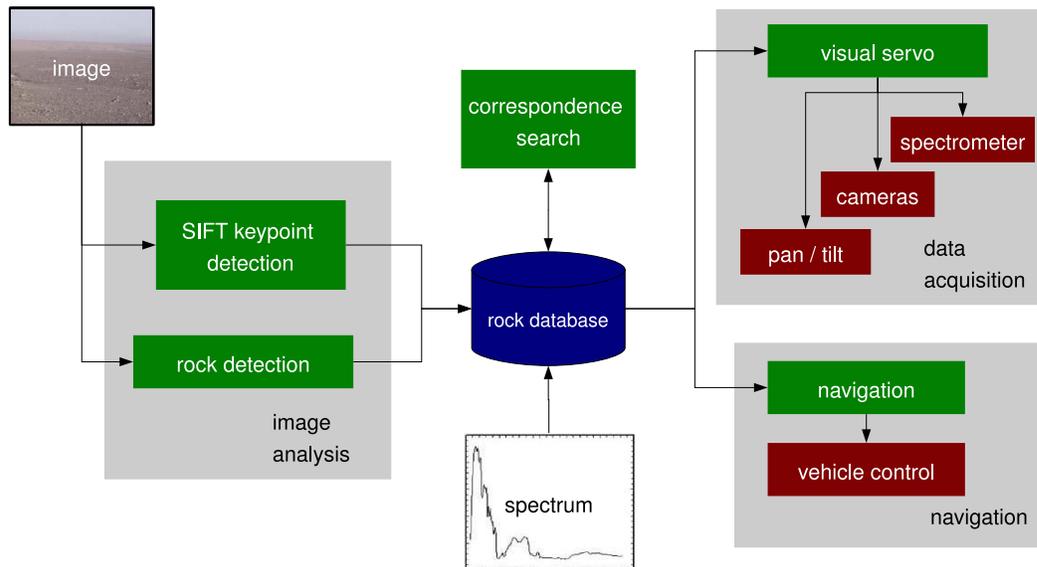


Figure 2.23: General architecture for feature tracking and spectrometer acquisition. The rock database contains a list of unique features and their associated attributes.

- Pixel bounding boxes for each image in which the rock appears.
- Scale-Invariant features for matching new instances of the rock against previous images.
- Visible attributes of the rock, such as its size.
- Spectrum attributes of the rock if available.

During operation the rover alternates between two states: a *travel* state for navigation and a *pause* state where it stops to collect spectra. In the travel state the rover navigates along a series of waypoints between the start location and the goal (Figure 2.22). The navigation module drives the robot from one waypoint to the next; it performs obstacle avoidance using stereo terrain analysis and onboard path planning. Meanwhile the data collection module commands forward-looking images, panning alternately to the left and right for maximum coverage of the terrain in front of the rover. The data analysis module examines these images and inserts newly detected rocks into a growing database of features.

A switch to the *pause* state requires two conditions: First, there must be at least one unmeasured rock within spectrometer range; and second, the current traverse segment must retain sufficient time for a spectrum collection sequence. Whenever these criteria are met the navigation module stops the rover, which signals the data collection module to begin spectrum collection. If the sequence succeeds it produces a new spectrometer measurement which enters into the database

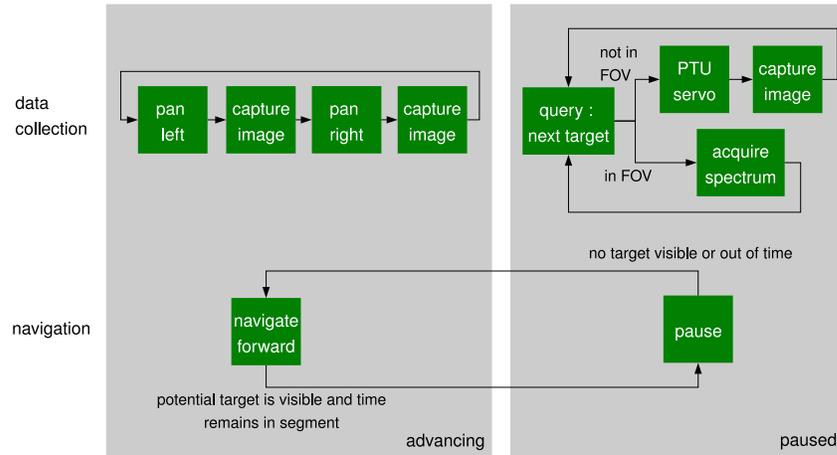


Figure 2.24: The system for generating spectrometer profiles has two states. The rover advances forward whenever it is out of time in its current segment or when no nearby potential spectrometer target is listed in the database.

to be associated with the correct rock. A flowchart showing actions taken by the independent data collection and navigation modules appears in Figure 2.24. Note that these two behaviors communicate implicitly through rover motion; the data collection procedure knows to begin the visual servo procedure whenever the rover comes to a stop.

The image analysis module detects rocks using the MVJ detector. We train the detector using typical images of rocks taken from the same cameras near the traverse site (Figure 2.9). The system discounts detected rocks less than  $10\text{cm}$  in size that are too small for accurate targeting. In addition, we eliminate rocks that are within  $2\text{m}$  of the rover center; this prevents spurious detections from the rover body or shadow from entering the database. The filtered result is a set of valid potential spectrometer targets. Figure 2.25 shows an example rock detection result from the Amboy Crater tests. The numbers associated with each bounding box show each rock's estimated position relative to the rover center.

We resolve the correspondence problem between images using Scale Invariant Feature Transform (SIFT) keypoints. SIFT keypoints are stable image locations that are invariant to scale and rotation and partially invariant to affine distortion, change in 3-D viewpoint, addition of noise and changes in illumination [50]. Scale-invariant SIFT keypoints are identified according to Lowe's method which searches for local extrema in the "scale space" of an image. The scale space  $L(x, y, \sigma)$  is defined by the convolution of a Gaussian kernel with size  $\sigma$  at the pixel  $(x, y)$ . Scale-invariance is achieved by searching for local maxima in the 3D space defined by the function  $D(x, y, \sigma) = L(x, y, \sigma) - L(x, y, k\sigma)$  where  $k$  is a constant factor. Even after rescaling the base image, the same extrema are likely to reappear at the same pixel locations for some value of  $\sigma$ .

The system pairs each rock with any SIFT keypoints that appear in its bounding box. The

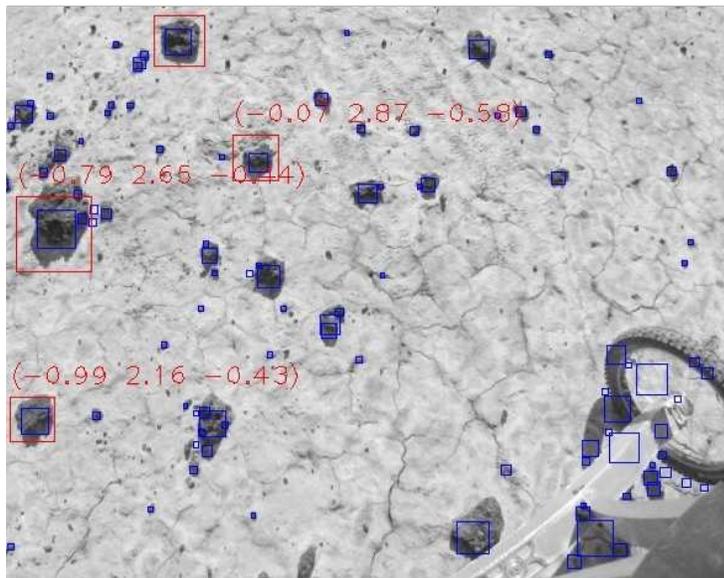


Figure 2.25: Detection result after filtering by rock size and distance. Large red bounding boxes are potential spectrometer targets, with numbers indicating rover-relative position estimates. Small blue squares indicate SIFT keypoints that are used for the correspondence search. The front axle is visible in the lower-right corner of the image.

rover establishes correspondence across images by matching the local pixel appearance of rocks' SIFT keypoints against the appearance of previous keypoints stored in the database. A single traverse might view thousands of unique keypoints; an unconstrained  $O(n^2)$  correspondence search is unsuitable for real-time visual servoing. Instead we constrain the search by exploiting geometric information. We will elaborate the stereo localization procedure shortly, but for now it is sufficient to know that the rover maintains a position estimate of each target. Given the known absolute pose of the rover and the relative orientation of the cameras the system determines an approximate region expected to contain the new SIFT keypoints. We use a search radius of 1 meter, and exclude any keypoint outside this boundary from the correspondence search.

The classical SIFT algorithm associates each keypoint with a 128-dimensional descriptor vector composed of a histogram of gradient magnitudes extracted from the local image region around the SIFT keypoint. The descriptors are defined relative to the principal gradient direction at the keypoint, and are invariant to slight changes in rotation, scaling, and — to some degree — affine distortions. A simple Euclidean distance metric applied to these descriptors can match physical points in the environment cross different views of the scene [50]. We evaluate correspondence between SIFT descriptors by calculating the Euclidean distance of the incoming descriptor against each of the remaining candidates in the database. We use the traditional SIFT match criterion that requires a nearest-neighbor distance at least  $n$  times smaller than the distance to the second closest neighbor for an empirically-defined threshold  $n$  [50].

The operating requirements for our automatic spectroscopy system differ from those of indoor environments for which the SIFT algorithm was designed. The original SIFT pose estimation procedure detects several SIFT matches for each object. It fits an affine homography to the altered pose using the RANSAC algorithm to identify multiple keypoints in a common configuration [50]. However, for spectrometer pointing we are only interested in a point estimate of the object's centroid (without orientation information), we only require a single SIFT match for each new image. This is a significant benefit because rocks themselves are feature-poor; they are often small, low-contrast objects that display many transient features.

In consequence our design favors the reliability over the quantity of SIFT matches. We alter the SIFT matching criterion with an additional “backwards” match to improve reliability. For each of the matches from the Euclidean search we do an equivalent search over SIFT descriptors in the new image. In other words, we apply a similar search for Euclidean nearest and second-nearest neighbors using the old SIFT descriptor as the query and the new image's SIFT descriptors as the search space (Figure 2.26). A match is valid when the thresholded Euclidean match is unique in both directions. In practice this excludes the vast majority of false matches.

A successful match implies that the SIFT descriptor lies on a previously-detected feature. We update the rock's visual and position information along with the SIFT descriptor's record to reflect the fact that illumination and aspect may have changed. We also merge all instances of the rock in the database to eliminate duplicate entries. If none of a new rock's SIFT descriptors match previous features it receives a new unique entry in the feature database. SIFT keypoints with no matches and that are not identified as part of any rock are assumed to belong to the background and discarded. Fortunately most rocks large enough to constitute spectrometer targets contain at least one feature

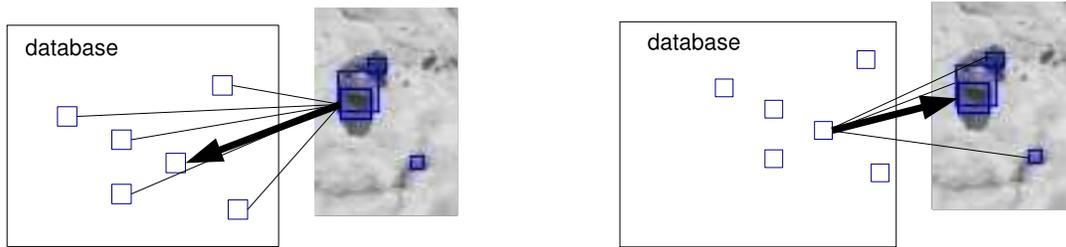


Figure 2.26: Correspondence with an existing SIFT descriptor in the database requires that the match be valid in both directions.

(Figure 2.25).

Note that we perform the correspondence search on all SIFT keypoints in the incoming image, not just those belonging to a rock detection. In general rock detection performance is less robust than SIFT keypoint extraction; the rock detector often fails to re-acquire a rock after changes in the viewpoint. Therefore all SIFT features in the incoming image are potential matching candidates. In order to avoid flooding the database with irrelevant “background” keypoints we only record those keypoints that match existing descriptors or belong to a newly detected rock. In other words, all SIFT descriptors stored in the database must have been identified as part of a rock at least in one image.

Two redundant methods estimate rock positions. An approach utilizing sparse stereopsis generates a 3D reconstruction based on the location of the rocks’ SIFT keypoints. The robust matching algorithm (detailed above in Section 2.3.1) identifies any keypoints that appear in both left and right cameras. Then a least-squares geometric reconstruction using a calibrated pinhole camera model and radial distortion [53] yields a sparse reconstruction. Some rocks contain more than one keypoint, yielding multiple position estimates. In these cases we use the 3D centroid of matched keypoints as the rock’s position estimate. We evaluate position estimation precision qualitatively by comparing the estimated position of rocks from the same rover position with different Pan/Tilt angles. This shows precision on the order of  $5cm$  for rocks within spectrometer range, which is more than sufficient to bootstrap our visual-servoing procedure.

Unfortunately, the stereo cameras are not time synchronized. This is not a problem during visual servoing while the rover is stationary, but it effectively prevents stereo while the rover is in motion. In these cases the rover position changes by an unknown amount between image acquisitions so computing a metric reconstruction based on stereo alone is infeasible. Instead we require a second approximate position estimation technique. When the robot finally stops to attempt a spectrum collection, this approximate estimate causes the first pan-tilt position to place the target rock in the cameras’ field of view. Then the robot re-acquires the rock from a stationary position using stereo, and completes the standard servoing procedure.

Our approximate pose estimation technique uses a single image and assumes the rock lies on a planar surface. This lets the rover calculate the rock position from a single image using known

kinematics. This also serves as a fallback position estimation strategy whenever the rover fails to find a SIFT match across stereo images for a target rock. Since the ground plane assumption introduces considerable inaccuracy we favor stereo whenever SIFT matches are available.

### 2.3.3 Visual Servo

The spectrometer pointing procedure attempts to aim the spectrometer foreoptic at a target based on the target's spatial position. This sequence has two stages. The first *kinematic* step leverages a kinematic model of the the pan/tilt mount and the most recent rock position estimates from the database. The second *feedback* step refines the initial pointing with an iterative visual servoing approach that places the target in a specific region in the image.

Both kinematic and feedback steps rely on generic models representing the location of the foreoptic cone. We expect to place the center of the cone on the target to maximize the probability of covering the sample. We fit the parameters of the physical model using calibration data obtained in laboratory tests. The main challenge of this calibration is to accurately determine where the spectrometer is pointing; there is no visual trace to determine the field of view. Instead, we infer the field of view using spectroscopic measurements of a known target: a commercial 532nm green laser pointer. For practical purposes the narrow-beam laser projects to a single point in space. The laser spot is easy to detect in collected spectra due to a high peak in its center frequency. Adjusting the point position to maximize the amplitude of this peak centers it in the spectrometer's field of view.

Calibration points produce a kinematic model of the foreoptic on the robot. Generic forward kinematic equations solve for the pan and tilt angles  $\theta$  and  $\alpha$ :

$$\begin{aligned}\theta &= f_{\theta}(x, y, z, h, d, \beta) \\ \alpha &= f_{\alpha}(x, y, z, h, d, \beta)\end{aligned}\tag{2.4}$$

Here  $(x, y, z)$  is the location of the rock in 3-D space,  $h$  is the height of the robot,  $d$  is the distance between the axis of rotation and the foreoptic, and  $\beta$  represents other less-relevant parameters such as rotations and translations of the foreoptic with respect to the base. We perform an unconstrained nonlinear optimization over the parameters trying to minimize the total square error of the predicted position. The result is an estimate for  $h$ ,  $d$  and  $\beta$ . Calderón provides additional details on the kinematic model [54].

It is impractical to maintain a sufficiently accurate model because many calibration parameters are unstable or vary with time. Examples are the changing tire pressure on the vehicle, inaccurate pan/tilt angle readings, and occasional replacement of parts. Therefore a spectrometer pointing system based solely on a feed-forward kinematic model will inevitably exhibit some pointing error. We incorporate a second closed-loop visual correction step to correct these errors and improve pointing accuracy.

Our visual correction method exploits the fixed transform between the camera and the foreoptic. Both devices are mounted to a single platform so their relative position is constant regardless of the

mobile base’s pan and tilt (Figure 2.21). Therefore the center line of the spectrometer sensing cone (*cc line*) is perceived as static by the camera and projects onto a line in the image. The intersection of the target with this line is fully determined by the distance from the foreoptic to the target.

A simple laboratory calibration procedure determines this geometric relationship. The *cc* line is not visible in imagery but we can infer its position using a laser reference aimed at a known location. We point the spectrometer at the laser reflection (maximizing the strength of the laser signature in the spectrometer signal) and capture an image of the scene. Since the laser point is centered in the foreoptic’s field of view, we can be sure that the visible laser spot in each calibration image corresponds to a point on the *cc* line. We can reconstruct the entire *cc* line from several laser points at different ranges. Perspective projection suggests a model relating the depth of the target and associated row/column pixels.

Pointing the spectrometer at a target is then equivalent to varying the position of the foreoptic until the target appears in a specific pixel of the image. This defines an “error” in visual space based on the current and desired position of the target. The system applies a simple visual servo [55] to correct the observed error. We determine the *Image Jacobian* that relates changes in the camera position to changes in the image. This permits a proportional visual control law:

$$\begin{bmatrix} \omega_{tilt} \\ \omega_{pan} \end{bmatrix} = K J_{\nu}^{-1} \begin{bmatrix} \Delta u \\ \Delta v \end{bmatrix} \quad (2.5)$$

In this equation  $\omega_{pan}$  and  $\omega_{tilt}$  are the estimated corrections in pan and tilt, respectively.  $K$  is a constant gain matrix of the control law,  $J_{\nu}$  is the image Jacobian, and  $(\Delta u, \Delta v)$  is the visual error in both image axes. The procedure iterates the visual correction until the error disappears. The result is a system that can accurately target features at ranges up to 5 meters.

The rover periodically renormalizes the spectrometer against a white reference target. This compensates for changes in illumination during the 40 minute traverse due to sun angle and atmospheric phenomena. We employ a standard VIS/NIR white reference target exhibiting high diffuse, Lambertian reflectance that is mounted to the rover deck. The rover collects a spectrum from the reference target every 7 minutes using a scripted sequence of pan-tilt actions. All other software modules pause for this collection procedure.

During the re-normalization the rover straightens its axles and servos the pan-tilt unit to target the white reference. Rough terrain sometimes results in poorly-oriented axles and — since the mast is fixed to the front axle — incorrect pointing. The rover senses this condition and aborts the attempt in favor of another 60 seconds later. Occasionally several attempts fail in succession; in these circumstances the rover continues trying every 60 seconds until a white reference spectrum is successfully captured. After a white reference the rover resets the standard 7-minute clock.

We chose the 7-minute recalibration time through trial and error by watching spectra change over time. Reflectivity values began to wander over longer timeframes; for example, recalibrations on the order of 15 minutes produced significant changes in reflectance and waveshape artifacts based on the growing miscalibration of separate internal sensors. Tests were performed under midday lighting conditions without significant changes in cloud cover. Future work under more dynamic lighting might require an even shorter calibration clock.

### 2.3.4 Feature Tracking Performance

Here we describe field rover tests performed at Amboy crater in the Mojave Desert, California (Figure 2.22). Here the rover traverses a field of rocks along the contact between an eroded basaltic lava flow and a sediment-covered plain. Accurate pointing proves necessary to produce clean spectra that contain the target signal. Tests show that rock detection combined with visual servoing significantly improves the diversity and quality of collected spectra over “blind” spectrometer pointing strategies.

We performed experiments at the lava fields of Amboy Crater in California over consecutive days in the fall of 2007. We identified a flat location near the border of the lava flow in which a basalt mound had partially eroded, leaving a scattering of rocks along a patch of ground approximately 50m on a side. The “rock garden” traverse site appears in Figure 2.20; clay predominates, but the scattered rocks up to 30cm in length are also visible. Basalt from lava flows appears as a dark material, while the underlying clay is lighter in color. This difference is also clear in VIS/NIR reflectance spectra.

Given that only these two materials are present, the amount of the target signal in each collected spectrum is a measure of the system’s targeting accuracy. In order to recover the proportion of basalt rock and clay sediment in each spectrum we presume that the spectra of each material mix linearly according to unknown weighting coefficients. This assumption holds experimentally for VIS/NIR spectra of macroscopic mixtures [56]. We establish reference spectra for a particular trial run by identifying clear examples of basalt and sediment in the spectra collected during that experiment. These reference samples comprise a basis for describing the rest of the collected spectra. We exclude water absorption bands; a least-squares linear deconvolution reveals the proportion of basalt in the remaining spectrum.

We performed 4 trials of the same traverse. We attempted to run each experiment under similar consistent lighting conditions, favoring trial times at midday when the sun angle changed slowly. Atmospheric conditions (clouds) were impossible to control, and varied somewhat across runs. We began each trial with the robot in the same starting location, and initialized its position using differential GPS. The robot traveled approximately the same 50m transect during each trial.

Figure 2.29 shows the result of each spectrum collection attempt for each of the four test runs. Table 2.1 provides performance results from each trial; the following sections explain each table column in detail. Figure 2.27 shows each of the runs independently, plotting feature detections against a georegistered visible-band overflight image. Small black dots show rock detections from the database, while large colored dots show collected spectra resulting from a completed visual servo procedure.

Individual detections can be hard to interpret or to compare across different runs. “Density plots” (Figure 2.28) facilitate visualization and comparisons across runs and demonstrate detection repeatability. These density plots presume the rocks are generated according to an inhomogeneous — that is, spatially varying — Poisson point process [57]. The point process density controls the mean number of rocks observed per unit area at any location in the map. The resulting local density, expressed in terms of rocks per square meter, permits comparison of detection results across runs

Table 2.1: Details of individual trials at Amboy crater.

Trial	Time	Rock Detect. Precision	Tracking Successes	Basalt Spectra
1	12:26 - 13:06	332/361 (92.0%)	26/50 (52.0%)	22/26 (84.6%)
2	13:15 - 13:55	249/280 (88.9%)	32/52 (61.5%)	27/33 (81.8%)
3	11:38 - 12:18	280/296 (94.6%)	29/43 (67.4%)	18/29 (62.1%)
4	12:27 - 13:07	299/340 (87.9%)	26/46 (56.5%)	17/26 (65.4%)

without requiring explicit correspondence between individual rocks.

We recover the density function with a simple a kernel smoothing strategy; Diggle supplies several practical methods for estimating local density while accounting for boundary effects [57]. Note that only rocks larger than the minimum size criterion contribute to the local density estimate.

The density plots reveal a high degree of clustering in the distribution of rocks at the site. Three prominent areas early in the traverse have a particularly high concentration of rocks. The high-density areas manifest in each of the trial runs, except for run 2 which appears to be missing one of the modes. It is unclear whether this was due to changes in lighting conditions or the starting position of the rover. Several other less-dense clusters appear late in the traverse.

We also evaluate the precision of the MVJ rock detector in the Amboy environment. We use the “strict” rock detection success criterion requiring that at least 50% of the detection contain a rock. The third column of Table 2.1 summarizes the resulting detection scores. The accuracy of our system is consistent with, or slightly better than, the MVJ detector’s performance on the MER benchmark datasets. We attribute any performance advantage to particularly favorable field conditions: the images all show dark, regularly-shaped rocks against an even light-colored background.

It is also instructive to consider system performance in targeting rocks based on their entries in the database. The main failure mode occurs when the system is unable to match a target’s previous SIFT descriptors to those in a new visual feedback image. For a successful correction sequence the system must match some new SIFT descriptor to previous descriptors for each of the 2 – 5 visual servoing iterations. Overall sequence success rates appear in the final column of Table 2.1. This measure does not consider the feature itself (i.e. whether the target was a true rock or a false detection), but is based simply on the system’s ability to track arbitrary SIFT descriptors. Thus, it eliminates the influence of the rock detector.

Many descriptors that are stored in the database never reappear during the visual correction process so it is not always possible to re-acquire a particular rock. Several causes may be at fault: changes of appearance of SIFT keypoints due to changes in illumination, rover pose, or shadows of the rock or the rover; a large number of similar SIFT descriptors in the database that prevents any confident match; or a SIFT keypoint that lies outside the camera’s field of view (unlikely in practice due to wide-angle cameras and robust rover position estimates).

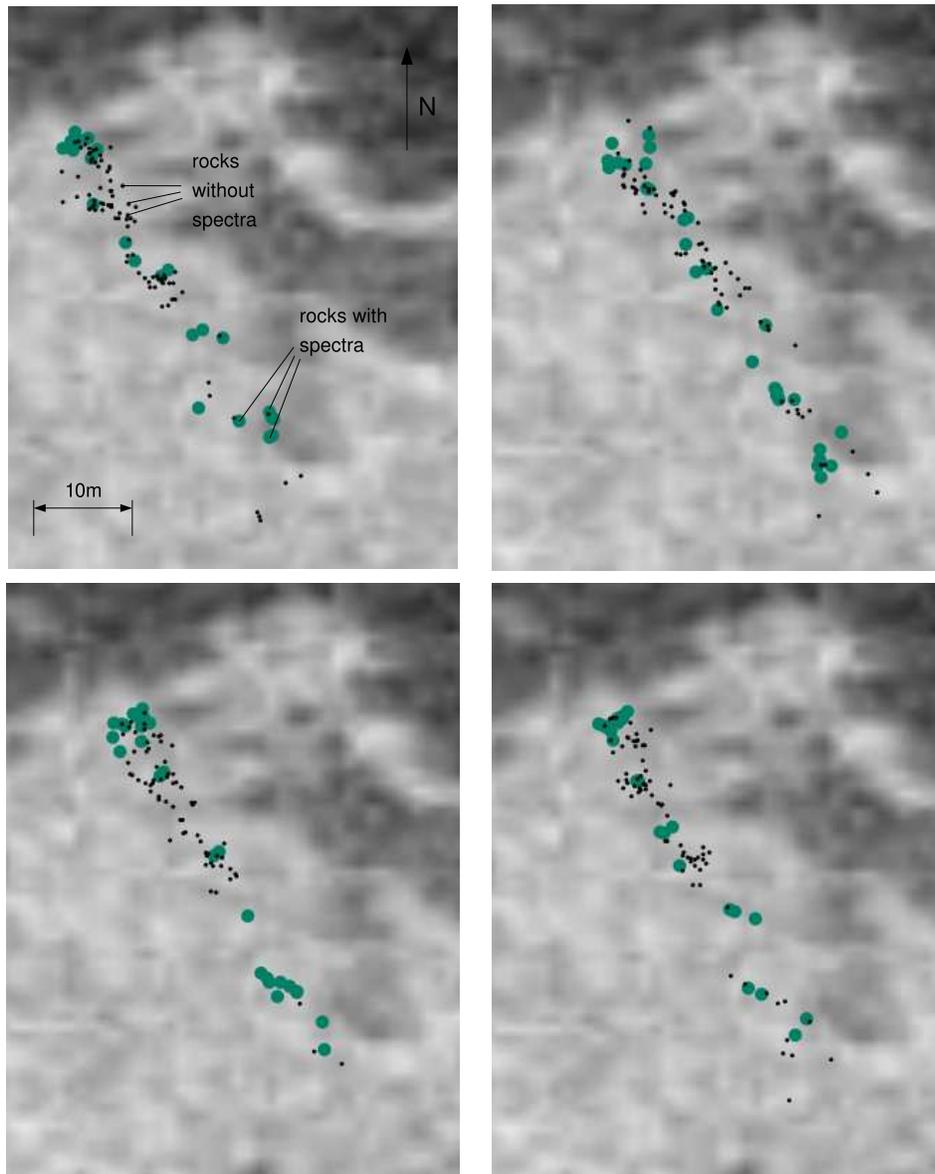


Figure 2.27: Maps resulting from the spectrometer profile trials. The four trials are ordered clockwise from the upper left. Small black dots show rock detections from the database, while large colored dots show collected spectra resulting from a completed visual servo procedure. Each run began in the upper-left corner, and ended at a goal waypoint located 50 meters to the Southeast.

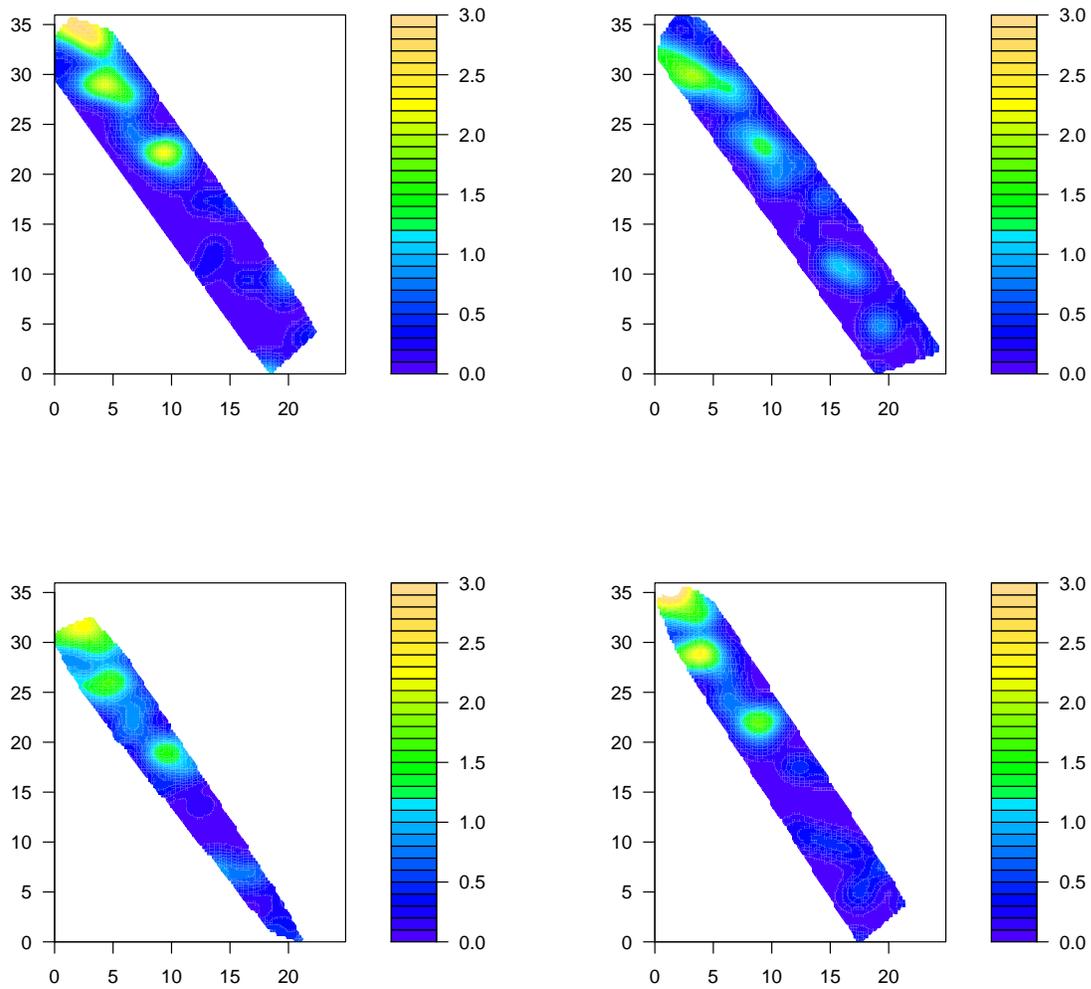


Figure 2.28: Density plots associated with each of the runs from Figure 2.27. Here the scale shows the number of detected rocks per square meter. We estimate the local density using a boundary-corrected kernel smoothing estimate [57].

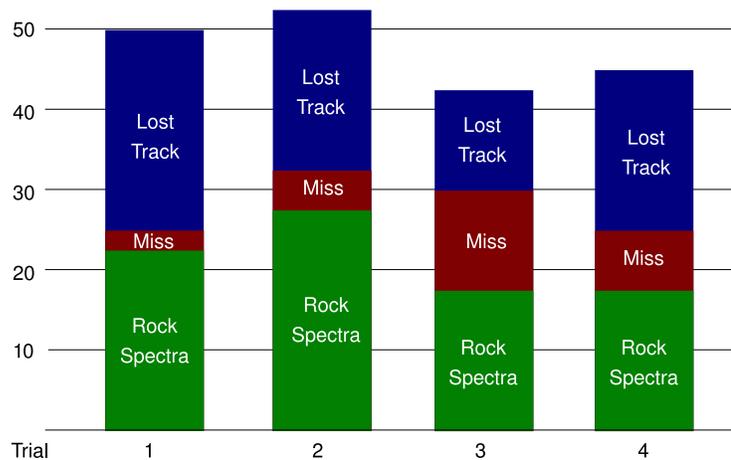


Figure 2.29: Result of spectrum collection by trial for each of four test runs.

Analyzing individual trials, we observe a lower performance in the first trial than the second (trials occurred consecutively on the same day). Many clouds were passing overhead during the execution of the first trial, causing significant changes in the light levels and complicating the SIFT descriptor matching. This is reflected in the captured data: failures cluster in adjacent tracking attempts whenever environmental conditions change abruptly. Trials 3 and 4 were also taken consecutively. Here we observe a slightly higher success tracking rate for trial 3. This may reflect a slower rate of illumination change when the zenith angle of the sun is higher.

In general we observe a tracking performance of over 55% when the light levels are favorable. An increased number of tracking failures will cause the system to spend more time looking for known rocks present in the database, but does not affect the quality of collected spectra.

Finally we quantify the accuracy of the spectrometer pointer to actually track rocks. We manually analyze each captured spectrum and classify it either as basalt or clay given that the spectra for both materials are known. The final column of Table 2.1 summarizes the result. Classification of signals as soil or basalt is based on a principal component analysis (PCA) noise-reduction strategy, followed by a simple linear deconvolution. In principal component space only one variable exhibits a large variance; this corresponds to the basalt/clay distinction. We use PCA as a noise-reduction strategy, projecting collected spectra onto the first component vector and reconstructing the signal to produce a filtered spectrum. In order to interpret the PCA coefficient we apply a linear deconvolution with representative samples of pure basalt and clay. The result is the percentage of basalt that is present in the mix. Spectra showing at least a 50% basalt are classified as a success.

The system presents a relatively high variance when hitting targets but consistently exhibits accuracies higher than 62%. Many factors influence the number of missed targets; some are intrinsic to the method while others relate to environmental conditions. This causes changes between trials and a large between-trial variance. Factors affecting system accuracy include: stereo estimate er-

Table 2.2: Tracking performance for periodic (blind) spectroscopy.

Trial	Basalt Spectra
1	0/68 (0.0%)
2	0/67 (0.0%)
3	0/64 (0.0%)

rors; errors introduced by the ground plane assumption when stereo is not available; spurious rock detections; tracking visual features lying out of the rock (e.g. background or shadow); tracking visual features lying in the border of the rock (SIFT features are not guaranteed to be in the center of the rock); and SIFT mismatches during the tracking procedure.

We compare system performance against *status quo* capabilities by performing three additional trials in the same testing area. In these trials the robot holds the pan/tilt unit at a fixed  $-30^\circ$  declination and blindly captures spectra at regular intervals during forward travel. Slight, inevitable perturbations of the start location (on the order of 10 – 20cm) ensure that the foreoptic’s field of view sweeps over different areas for each 50m traverse.

The results of the blind spectrometer mapping trials appear in Table 2.2; no measurement of basalt appears in any of the spectra from any of the trials. This demonstrates the difficulty of the testing area; rocks were generally small and scattered. Therefore it is highly unlikely that a target would fall by chance into the spectrometer field of view during the periodic spectrum acquisition. Our pointing strategy outperforms the blind method; each trial measures at least at 26 individual rocks over a 50m traverse lasting no longer than 40 minutes.

In summary, we have demonstrated that feature detection and tracking can produce science feature maps in real time with no supervisory input from human operators. This improves science data return in several respects. First, it enables spectrum sampling of targets encountered during long over-the-horizon traverses that scientists cannot see or specify in advance. This provides many more pure spectrum samples than can be gathered through current “blind” sampling strategies and improves the science return of the site survey.

Additionally, the resulting map constitutes a bandwidth-efficient data product to describe traverse images. The detection and tracking procedures are invariably imperfect so the resulting maps may still exhibit some biases (for example, the system may detect more rocks at midday when lighting conditions are ideal). Thus one should be cautious in using them as a quantitative measurement tool. However, the repeatability of the Amboy field trials suggests that these maps do capture objective physical properties of the environment. They evidence regional trends and transitions that can benefit interpretations and mission planning. In the following section we will consider methods for performing inference from these observations. Feature maps facilitate predictions of future observations that inform exploration decision making.



## Chapter 3

# Spatial Models

[These studies] provide explanation of presence or absence, of origins and limits. Distribution is the key to process. — Carl Sauer, *the Education of a Cartographer*” [58].

The spatial model is a generative map of the environment that extrapolates from previous observations and predicts new measurements at unvisited locations. Its input is a list of discrete classified science features. The model uses this data to infer the result of future data collection actions elsewhere, facilitating adaptive decisions about where to travel and what instruments to deploy (Figure 3.1). According to our intelligent mapping approach, the model describes spatial and cross-sensor relationships between data products. It must permit computationally efficient on-line learning so that the agent can modify model parameters on the fly in response to unanticipated trends.

This chapter details the specific case of Gaussian processes, which are a class of spatial models well-suited to these requirements. Gaussian processes are nonparametric Bayesian models commonly employed in geostatistics. We begin by reviewing previous relevant work in spatial statistics and robotic mapmaking. Then we present an overview of the Gaussian process regression models. Our work will treat the specific case where the science phenomena of interest can be expressed in terms of a scalar value at each observation location. The Gaussian process regression model is a good match for these tasks and readily incorporates input data from multiple sensing scales. We augment the traditional spatial model by introducing orbital data in the form of latent input dimensions. This improves performance in modeling surficial units and other geologic phenomena.

Our goal is to demonstrate accurate model predictions that extrapolate well to distant locations. This is important for any attempt to the resulting map for as a compressed representation of observations for transmission to Earth. It is also important for robots hoping to use these models to inform adaptive action selection and return decisions. We will deal with the question of optimal actions in Chapter 3, and focus for now on the interim goal of reliable inference.

Our evaluation method compares alternative models by reconstructing missing observations from partial datasets. After an entire dataset is collected, we can hide a portion of the observations and use the remainder as a ground-truth standard for evaluating model predictions. We will examine model performance in reconstructing surficial maps of Amboy Crater from rover data.

Traditional Gaussian processes based on samples' positions in the environment are a promising model for our purposes. However, some modifications can improve performance significantly. In particular, tests will show that the use of *latent input dimensions* — additional non-spatial inputs based on actual sensor data — can improve reconstructions over the basic formulation. In particular, the use of orbital data as a latent input promises significant benefits for autonomous geologic site survey.

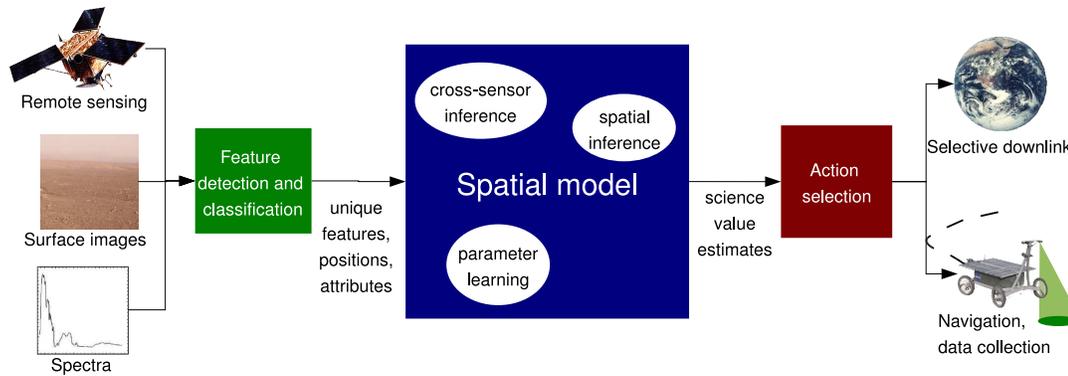


Figure 3.1: Spatial models

### 3.1 Previous work in Mapmaking for Robotic Exploration

The field of relevant work in automated robotic mapping is extensive; a complete review is beyond our scope. Instead we will consider a few representative examples of the classical approaches. One can bifurcate the broad field of robotic mapmaking according to the distinction between *metric maps* that reference features according to some common global coordinate system and *topological maps* that reference landmarks using local connectivity relationships [59].

*Occupancy grids* are a common metric representation [60]. Here the robot subdivides the environment into a grid and associates each grid square with a probability that the grid square contains an obstacle. These methods provide a natural means of combining range sensor data in cluttered environments; the map grows with the size of the area explored and not the number of features (that might otherwise be arbitrarily large). Researchers have recently used the occupancy grid approach to three dimensions for simultaneous localization and robotic mapping of submerged cave environments [61]. Dynamic data structures permit efficient caching for maintaining multiple hypotheses about the robot's position and the map.

Often grid-based approaches are appropriate for validating and extending orbital data products. For example, the Haughton crater mapping expedition deployed LIDAR and ground-penetrating

radar instruments that were later Coregistered to an orbital Digital Elevation Model, or DEM [6]. Alternatives to grid-based methods are simple point maps: lists of features referenced in a global coordinate frame. The Life in the Atacama expedition [7] treated rover surface measurements in this fashion, constructing an ECEF-referenced point map of microscopic image measurements over many square kilometers of the Atacama desert [62]. Scientists coregistered these images with point data from a variety of additional sensors, including humidity, temperature and spectrometer data; the result was a comprehensive metric survey of microbial habitats in the explored area.

In contrast with metric methods, topological maps express location in terms of positions relative to landmark features. If the landmark detection system is reliable, these methods can be more robust to localization drift. Examples of topological maps include room-and-corridor models used for indoor navigation, or the compressed Voronoi maps used in graph-search solutions to planning tasks. Kuipers expands on this idea with a “hierarchy of spatial relationships” [63] that infers metric geometry from Voronoi topologies. These topologies are in turn define in terms of low-level control strategies that move the agent from one location to the next. For example, a “wall-following” strategy could transport an agent between physical locations corresponding to the corners of a room.

Thrun et al. offer another joint treatment of topological and metric maps [59]. Their approach is statistical: the use the Expectation-Maximization algorithm [64], alternatively estimating topological and metric relationships that improve the likelihood of the observed sensor data. Note that in general researchers using topological techniques are concerned with structured indoor environments. Topological maps are less common in exploration domains where scientists may wish to correlate multiple sensors at different scales. Cross-sensor correlation, especially with respect to orbital data products, is straightforward when position estimates are expressed in a common geospatial coordinate frame.

These mapping techniques are *static* in the sense that they do not extrapolate beyond the agent’s sensing horizon to unvisited locations. Extrapolation requires statistical inference with a model that can describe conditional dependence relationships with unobserved locations. Indoor environments (in which much robotic mapping research occurs) are poorly-suited to inference. For example, it is difficult to extrapolate from observations of walls and corridors to predict what a particular, distant unobserved occupancy grid square might contain. Natural environments have simpler continuous structure more conducive to statistical modeling.

The Robotic Antarctic Meteorite Search study investigates the use of generative spatial models to improve *post hoc* classification performance on collected data [65]. Here Pedersen exploits the spatial continuity in rock distributions to compensate for classification noise. Neighboring observations share a spatially-smooth prior over rock classes. Pedersen’s model treats the class of a particular rock as a sample from a multinomial distribution with a Dirichlet prior. The Dirichlet parameters at any new location are a weighted sum of the previous observations, with each coefficient determined by an isotropic smoothing kernel.

Pedersen’s work differs from our proposed system in several ways. First, it stops short of using the spatial model to inform experimental design decisions. It only considers correlations in the spatial dimension, and provides no means of learning the degree of spatial correlation from the data

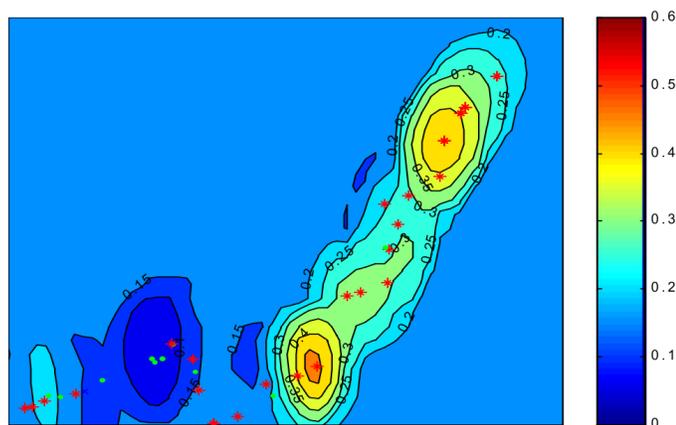


Figure 3.2: Generative models of spatial data improved rock classification decisions in the RAMS meteorite search. This image shows the learned probability of observing metamorphic rocks at various locations in an explored area. Image courtesy Pedersen [65]

directly. Nevertheless, Pedersen shows that these spatial smoothness assumptions improve classification scores (Figure 3.2). His work demonstrates that spatial geostatistics can be successfully applied to the rover exploration domain.

There is another notable gap in contemporary exploration robot mapping research. Studies have conducted autonomous science analysis of both ground-based and remote sensing instruments. However, no research has yet attempted to do both simultaneously on the same platform. This contrasts with site survey by humans in which geologists rely heavily on remote sensing data to corroborate trends and patterns observed from the surface [62, 51]. This analysis is important, not simply in post analysis, but also in mission planning — it provides important cues about scientifically interesting locations to visit next. Effective use of multi-scale sensing will become increasingly important as the daily traverse range continues to increase.

## 3.2 Gaussian Process Background

Chapter 1 outlined some important capabilities of spatial models in the geologic site survey domain. They should permit online learning and inference while representing both spatial and cross-sensor correlations in collected data. We will see that the specific the spatial dependence assumptions encoded in these models can exert a strong influence on the resulting exploration behavior. This is not necessarily a “side effect,” but rather an advantage of the intelligent mapping approach: it reduces questions about appropriate sampling strategy to model-selection issues to be resolved from intrinsic properties of the data.

Here we focus on observation sites characterized by a vector-valued independent variable  $x \in \mathbb{R}^d$ . The variable  $x$  represents known properties of the measurement location. For exam-

ple,  $x$  could include the site’s physical position, represented by latitude and longitude coordinates or as a position index along a transect. It could also include one or more “latent” input dimensions corresponding to preexisting measurements or remote sensing data. For example, Section 4 demonstrates a surficial mapping application where the salient dependent variables are the sample’s latitude, longitude, and the pixel intensities of the associated remote sensing data. We transform the latitude and longitude position  $[l_{i0}, l_{i1}]$  along with  $m$  remote image channels  $[c_{i0}, \dots, c_{im}]$  to a common scale, and append them to produce the input vector  $x_i$ :

$$x_i = [l_{i0}, l_{i1}, r_{i0}, \dots, r_{im}] \quad (3.1)$$

During the traverse the agent collects a set of *observations* — discrete features produced by detection and classification routines. Each observation associated with the vector  $x$  of independent variables yields a dependent scalar-valued geologic class  $y \in \mathbb{R}$ . This scalar could represent such varied phenomena as the presence or absence of a certain mineral, the composition of a mixture, or the altitude of an elevation model. In general observations might be vector-valued or discrete classifications. We will focus on scalar-valued classes, as they proved adequate for both the simulations and field trials we consider here.

Our site survey task aims to recover an underlying function  $f(x)$  which maps the input vector  $x$  onto the observed geologic class  $y$ . Formally,  $f(x)$  transforms the vector-valued independent variables  $X = \{x_1, x_2, \dots, x_n\}$ ,  $x_i \in \mathbb{R}^d$  onto the scalar geologic class. Introducing normally-distributed noise  $\epsilon$  yields the resulting dependent observations  $Y = \{y_1, y_2, \dots, y_n\}$ ,  $y_i \in \mathbb{R}$ .

$$y = f(x) + \epsilon \quad (3.2)$$

In other words, the observed value is the underlying function after perturbation by some small amount of noise. Given a set of these observations, we would like to predict the true value of  $f(x)$  at other locations. We can infer this probability distribution from a prior  $P(f(x))$  using Bayes’ rule.

$$P(f(x)|Y) = \frac{P(Y|f(x))P(f(x))}{P(Y)} \quad (3.3)$$

Computing this distribution is central to the intelligent mapping method. The probability distribution over possible  $f(x)$  provides a single surficial map of any desired extent and resolution; one need only evaluate *Maximum A Posteriori* (MAP) estimates  $\hat{f}(x)$  at unobserved sample sites to yield the most likely observations.

These MAP estimates are useful for visualization; they constitute a “point estimate” of the site geology. However, the complete posterior carries a much richer description of the environment. For example, the marginal prediction variance also describes our uncertainty about map contents at each sampling location. Moreover,  $P(f(x)|Y)$  expresses the probability of *any* combination of future observations. Given a set of observed data it can fully characterize the (potentially multimodal) space of possible maps.

A Bayesian approach raises the question of how to parameterize  $f(x)$  in order to represent these distributions and perform efficient inference. We utilize a Gaussian process approach. Gaussian processes [19, 20, 31] are known in the geostatistics community as “Kriging” models; they

are a powerful technique for nonparametric regression with several distinct advantages for the intelligent mapping domain. They are true Bayesian models with posterior distributions over map parameters, and germane to the large body of classical work in Bayes-optimal experimental design. Additionally, Gaussian process models are pervasive in terrestrial geostatistics so they are a natural fit for planetary geology and our site survey application. Rasmussen and Williams [66] offer a thorough introduction, but we will provide a brief review of the basic principles. Finally, the posterior distributions are all Gaussian, making efficient computation of entropies tractable for information-driven exploration strategies.

One can view Gaussian processes as a generalization of a traditional parametric function approximation method such as a neural network. Following MacKay [31] we express  $P(f(x))$  with a prior distribution over a set of parameters of  $f(x)$ . For example, we can use a series of basis functions  $\phi_j(x)$  where the free parameters are coefficient weights  $w_j$ . The evaluations of all  $m$  basis functions at a single data point form a length  $m$  vector; for the complete dataset of  $n$  input points we stack them to form an  $n \times m$  matrix  $R$ :

$$f(x_i) = \sum_j \phi_j(x_i)w_j = R_i w_j \quad (3.4)$$

Even simple basis functions can approximate complicated maps by using sufficiently many bases to express all of the “bumps and wiggles” in the underlying geology  $f(x)$ . This is easy to show in the extreme case of many basis functions consisting of infinitely-narrow convolution kernels. We can imagine that every possible observation location has a single nonzero basis function, with a weighted value corresponding to the true  $f(x)$  evaluated at that location so that the true  $f(x)$  is always equivalent to its representation as the linear combination of bases.

In the general case  $x$  is continuously valued and there are fewer bases than possible observation locations so that a linear combination of bases can only approximate  $f(x)$ . Regardless of the actual number of basis functions, the true values of  $w$  are unknown at the beginning of the mapping task. We can express a prior distribution over these values using a zero-mean Gaussian distribution:

$$P(w) = \mathcal{N}(0, \sigma^2 I) \quad (3.5)$$

We use the notation  $f(X)$  to express  $f(x)$  evaluated at the entire set of input points. The prior over  $f(X)$  is an affine transformation of the weights and is therefore a Gaussian with covariance matrix given by  $R\sigma^2 I R^T R^T$ . We can write the prior in terms of a covariance matrix  $K$ .

$$P(f(X)) = \mathcal{N}(0, R\sigma^2 I R^T) \quad (3.6)$$

$$= \mathcal{N}(0, \sigma^2 I R R^T) \quad (3.7)$$

$$= \mathcal{N}(0, K) \quad (3.8)$$

It follows that  $K$  is an  $n \times n$  matrix representing an inner product between basis functions:

$$K_{ij} = \sigma^2 \sum_h \phi_h(x_i) \phi_h(x_j) \quad (3.9)$$

We can make our function approximation more expressive by increasing the number of basis functions; this is tantamount to increasing the number of hidden units in a neural network or the wavelet

bases of a decomposed signal. However, as the number of basis functions increases the computational cost of learning and computing them becomes prohibitive. A Gaussian process approach exploits a *kernel trick* that supersedes these bases altogether, parameterizing  $K$  directly as a function of the input points. This amounts to placing a prior directly on the space of functions  $f(x)$ .

### 3.2.1 Inference with Covariance Functions

We define a *covariance function*  $\kappa(x_i, x_j)$  that represents expected correlations between points in the input space. Individual entries of  $K$  are given by the covariance function evaluated between each pair of datapoints:  $K_{i,j} = \kappa(x_i, x_j)$ . To define a valid Gaussian process prior this covariance function must itself be positive-definite, i.e. it must yield a positive-definite covariance matrix.

$$P(f(x)) = \mathcal{N}(0, K) \quad (3.10)$$

Recall that the geologic observations are perturbed by normally-distributed noise  $\epsilon \sim \mathcal{N}(0, \sigma^2)$ . The observations of geologic class  $Y$  are jointly Gaussian with the noise term contributing to the diagonal of the covariance matrix. The covariance matrix  $K + \sigma^2 I$  represents noise-perturbed observations:

$$P(Y) = \mathcal{N}(0, K + \sigma^2 I) = \mathcal{N}(0, K) = \frac{1}{Z} e^{-\frac{1}{2} Y^T (K + \sigma^2 I)^{-1} Y} \quad (3.11)$$

Recall that  $f(x)$  evaluated at a set of input points has a multivariate Gaussian distribution. It is fully specified by a mean vector (which we will take to be zero for simplicity) and the positive-definite covariance matrix.

A popular choice of covariance function is the *squared exponential* [31]. Excluding the noise term, we write the complete kernel function using the hyperparameters  $\psi_1, \psi_2$ , and additional “length scale” hyperparameters  $w_k$  for each dimension  $k$  of the input space:

$$\kappa(x_i, x_j) = \psi_1 + \psi_2 \exp \left\{ -\frac{1}{2} \sum_{k=1}^d \frac{(x_{ki} - x_{kj})^2}{w_k^2} \right\} \quad (3.12)$$

We can use the resulting prior to make inferences about future geologic observations from a set of collected data. For a set of observations  $X$ , given here in matrix form, we wish to predict values at new locations  $X^*$ . The posterior  $P(f(X^*)|X, Y)$  is normally-distributed according to  $\mathcal{N}(\hat{\mu}, \hat{\Sigma})$ . Our actual observations  $X, Y$  and predictions  $f(X^*)$  are jointly Gaussian so this distribution is simply a conditional Gaussian. We introduce matrix subscript notation  $K_{XX^*}$  to represent the covariance function evaluated between the set of training points  $X$  and the novel locations  $X^*$ :

$$K_{XX^*} = \begin{bmatrix} \kappa(X_1, X_1^*) & \kappa(X_1, X_2^*) & \dots & \kappa(X_1, X_m^*) \\ \kappa(X_2, X_1^*) & & & \\ \dots & & & \\ \kappa(X_n, X_1^*) & & \dots & \kappa(X_n, X_m^*) \end{bmatrix} \quad (3.13)$$

Following MacKay [31], the posterior mean  $\hat{\mu}$  of the predicted function at the new locations is given by the Gaussian distribution after conditioning on the vector of training outputs  $Y$ .

$$\hat{\mu} = K_{X^*X} (K_{XX} + \sigma^2 I)^{-1} Y \quad (3.14)$$

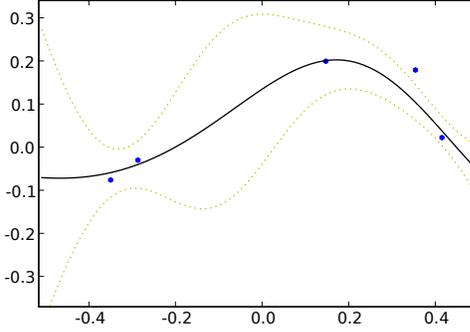


Figure 3.3: Gaussian process regression with a dataset of 6 training points. This plot shows the MAP prediction  $\hat{\mu}$ . The dotted envelope shows the variance of the marginal prediction associated with each location.

The posterior covariance matrix  $\hat{\Sigma}$  is also derived from the conditional.

$$\hat{\Sigma} = K_{X^*X^*} + K_{X^*X}(K_{XX} + \sigma^2 I)^{-1}K_{XX^*} \quad (3.15)$$

In practice, the inversion is usually performed through Cholesky decomposition to ensure numerical stability [67]. Decomposing  $(K_{XX} + \sigma^2 I) = LL^T$  yields factors that can be back-substituted to yield the solution.

Figure 3.3 shows a Gaussian process regression model applied to a small dataset of 6 collected observations. The dotted envelope illustrates the marginal prediction's variance at with each location in the input space. This is equivalent to the diagonal of  $\hat{\Sigma}$ , or alternatively, prediction variance  $\hat{\Sigma}$  for each of many separate inferences over the independent locations  $X^*$  along the horizontal axis.

We can rewrite the MAP prediction mean as a linear combination of kernel evaluations weighted by coefficients.

$$\alpha = (K_{X,X} + \sigma^2 I)^{-1}Y \quad (3.16)$$

$$\hat{\mu} = K_{X^*X} \alpha \quad (3.17)$$

Equation 3.17 demonstrates that Gaussian process regression is an example of a more general class of nonparametric regression methods known as *linear smoothers* [64]. Linear smoothers predict new observations using a linear combination of kernel functions evaluated at the training observations. We might worry that we have lost something in moving from a parametric form (with arbitrary degrees of freedom) to a nonparametric model determined by a covariance function containing only a handful of user-defined parameters. However, a powerful result, the *representer theorem* [68], shows that for any given kernel the optimal regression prediction must have the form of equation 3.17 for a large class of regularized cost functions.

For a Reproducing Kernel Hilbert Space  $\mathcal{H}$  and a loss function  $V(x, y)$ , we aim to choose a function  $\hat{f}$  that minimizes the regularized loss:

$$\hat{f} = \operatorname{argmin}_{f \in \mathcal{F}} \frac{1}{n} \sum_i V(\hat{f}(x_i), y_i) + \lambda \|f\|_{\mathcal{H}}^2 \quad (3.18)$$

Here  $\lambda$  is a coefficient that penalizes complexity of  $\hat{f}$  by weighting the function's norm within the Reproducing Kernel Hilbert Space. The optimal evaluation at a new point  $\hat{f}(x)$  is given by a linear combination of kernel evaluations:

$$\hat{f}(x) = \sum_i \alpha_i \kappa(x_i, x) \quad (3.19)$$

Once one has identified a covariance function the unstructured problem of choosing  $f$  from an infinite-dimensional space becomes a simple matter of computing the linear coefficient vector  $\alpha$ . This yields the form of the predictive mean used in equation 3.14.

### 3.2.2 Learning Hyperparameters

Our discussion so far has ignored the issue of how to identify hyperparameters of the covariance function. We can fix covariance function hyperparameters in advance, but in practice it is almost always preferable to learn them directly from the data. Convention favors two approaches to hyperparameter learning in Gaussian process models. The first, a pure Bayesian approach, begins with a prior distribution over hyperparameters and employs simulation strategies such as Markov Chain Monte Carlo to characterize the distribution over hyperparameters. This distribution is inevitably non-Gaussian.

Alternatively one can identify point estimates of hyperparameters using an evidence maximization (maximum likelihood) approach. The general approach is one of gradient descent. In theory this is susceptible to local minima, but we find that these are rare in practice for the Gaussian covariance function. Moreover, gradient descent methods offer a substantial speed advantage which makes them suitable for on-line learning by a remote agent during exploration.

For simplicity of notation we introduce  $Q$  representing the noise-perturbed covariance matrix of training observations.

$$Q = K_{XX} + \sigma^2 I \quad (3.20)$$

The data likelihood  $L(X, \kappa)$  for the squared exponential covariance function has the following general form [20]:

$$L(X, \kappa) = -\frac{1}{2} \log \det Q - \frac{1}{2} Y^T Q^{-1} Y - \frac{n}{2} \log 2\pi \quad (3.21)$$

The likelihood gradient with respect to any hyperparameter  $\psi$  can be written:

$$\frac{\partial L(X, \kappa)}{\partial \psi} = -\frac{1}{2} \operatorname{Trace}(Q^{-1} \frac{\partial Q}{\partial \psi}) + \frac{1}{2} y^T Q^{-1} \frac{\partial Q}{\partial \psi} Q^{-1} y \quad (3.22)$$

This requires the partial derivatives of the noise-perturbed covariance matrix  $Q$  with respect to each parameter of equation 3.2.1. The partials of an element  $Q_{ij}$  with respect to the constant and proportional terms  $\psi_0$  and  $\psi_1$  are the following:

$$\frac{\partial Q_{ij}}{\partial \psi_0} = 1 \quad (3.23)$$

$$\frac{\partial Q_{ij}}{\partial \psi_1} = \exp \left\{ -\frac{1}{2} \sum_{k=1}^d \frac{(x_{ki} - x_{kj})^2}{w_k^2} \right\} \quad (3.24)$$

The requirements of the geologic mapping task complicate the learning of length scales  $w$ . Often the subsets of input dimensions corresponding to spatial directions are *mutually isotropic*. In other words, smoothing is equivalent in all spatial directions. If there is no reason to believe that our chosen spatial coordinate frame is particularly special, or reflects some structure in the environment, then we should use a single length scale parameter to force equal smoothing along all spatial dimensions. A length-scale parameter responsible for a set of input dimensions  $D$  has the following partial derivatives:

$$\frac{\partial Q_{ij}}{\partial w} = \psi_1 \exp \left\{ -\frac{1}{2} \sum_{k=1}^d \frac{(x_{ki} - x_{kj})^2}{w_k^2} \right\} \left[ \sum_{k \in D} (x_{ki} - x_{kj})^2 \right] \frac{1}{w^3} \quad (3.25)$$

These permit us to calculate the likelihood gradient of equation 3.22.

After finding the likelihood gradient with respect to each hyperparameter we perform an iterative ascent of the likelihood function. Simple gradient ascent can take many iterations to converge; it is generally preferred to use conjugate gradient strategies that preserve gains from previous iterations [69]. Accordingly we use an iterative approximation to find conjugate gradient directions. The Polak-Ribière method perturbs the gradient at each time step after the first:

$$\text{dir}_t = \nabla L(X, \kappa_t) - \beta_t \nabla L(X, \kappa_{t-1}) \quad \text{for} \quad (3.26)$$

The value for  $\beta_t$  is given by the following expression:

$$\beta_t = \frac{\nabla L(X, \kappa_t) (\nabla L(X, \kappa_t) - \nabla L(X, \kappa_{t-1}))}{\|\nabla L(X, \kappa_{t-1})\|^2} \quad (3.27)$$

We perform a line search along the conjugate gradient direction to maximize  $L(X, \kappa)$  and recompute the conjugate gradient from the new hyperparameter values. This method offers a substantial speed improvement over simple gradient descent strategies; it converges in a handful of iterations. We avoid local maxima by restarting the optimization at random locations in the hyperparameter space.

### 3.3 Stationary and Nonstationary Environments

The squared exponential covariance function is *stationary*, meaning that the relationship between two locations' posterior predictions depends only on the locations' relative positions. Covariances are independent of the absolute location in the input space. This is apparent from Equation 3.2.1 in

which  $\kappa(x_i, x_j)$  changes based on the relative distance along each input dimension. In other words, the covariance function presumes a constant rate of change throughout the environment.

Thanks to their simplicity and generality, these stationary covariance functions pervade spatial modeling research. Nevertheless their assumptions do not necessarily apply to the remote site survey domain. Geology often violates the constant change requirement. Typical terrain classifications concern contiguous masses or units of material with common physical properties. These units often weather and transport *en masse*, resulting in homogeneous regions with sharp discontinuities. The geologic map of image 1.1 evidences these surface units separated by clean borders.

Second, stationary covariance functions carry strong implications for adaptive sampling. The discussion of Chapter 4 will show several useful adaptive observation strategies that depend exclusively on the covariance matrix of the posterior  $P(f(X^*)|X, Y)$ . In particular, the covariance matrix fully determines information-gain approaches to observation selection, so experimental-design strategies will require considerable attention to the second-order properties of the posterior distribution. Surprisingly, stationary covariance implies a periodic sampling strategy which removes any benefit from adaptive data collection. This implies that the standard covariance assumptions may be inadequate for the site survey task.

Here we discuss modifications to the stationary models to address these shortcomings. We will consider alternative nonstationary covariance functions as well as *latent input* methods that use additional input dimensions to describe nonstationary geologic phenomena.

### 3.3.1 Nonstationary Covariance Functions

A simple case with synthetic data elucidates the difference between nonstationary and stationary covariance. Consider the function portrayed in Figure 3.4. Here the underlying  $f(x)$  is a simple step discontinuity. We perturb it using additive Gaussian-distributed observation noise. The result is a highly nonstationary environment composed of two homogeneous legs and a boundary with a high rate of change. The ideal model should reflect the fact that predictions at two locations on either side of the boundary are loosely related, while two similarly-spaced points on the same side of the boundary have predictions that covary strongly.

The image at upper left models the step using a standard stationary squared exponential covariance function. We train covariance function hyperparameters using maximum likelihood. The resulting hyperparameters must compromise between estimating the smooth and discontinuous regions. This undersmooths in the homogeneous areas and oversmooths at the discontinuity. A dotted line shows prediction variance (rescaled for easy visualization). Prediction variance is determined completely by the density of the data, with high uncertainty near the periphery but not at the boundary.

The image at upper right shows an alternative model utilizing a nonstationary covariance function. Here our method is inspired by the work of Paciorek [70], who enables nonstationary covariance by allowing the length scale parameters to vary over the input space. The idea is to parameterize the length scales of the Gaussian kernel using a function of that point's input location. The covariance function for two points  $x_i$  and  $x_j$  in the input space is formed by the average of the

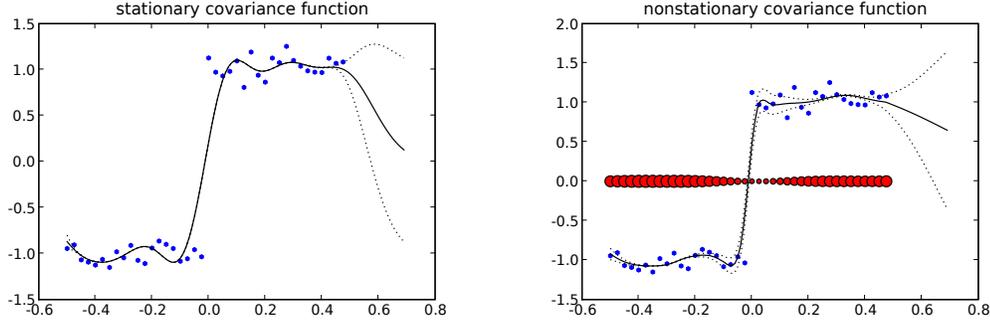


Figure 3.4: Left: The noisy step function modeled with a stationary squared exponential covariance function. The prediction variance, normalized to improve clarity, is shown by the dotted line; it indicates high uncertainty near the periphery of the data. Right: Nonstationary covariance functions using locally-varying length scales. The red circles illustrate the maximum-likelihood relative length scales for different locations in the input space.

variance parameters  $w_i$  and  $w_j$  at each of the two locations. Limiting the kernel to a single independent dimension (as in a time series model) the resulting covariance function, with a normalizing constant  $Z$ , is:

$$\kappa(x_i, x_j) = \psi_1 + \psi_2 \frac{1}{Z} \exp \left\{ -\frac{1}{2} \sum_{k=1}^d \frac{(x_i - x_j)^2}{\left(\frac{w_j + w_i}{2}\right)^2} \right\} \quad (3.28)$$

The covariance function for the training data is fully defined by a length scale parameter Paciorek shows that the result is a positive definite kernel function.

In order to evaluate the nonstationary function at a novel location one must identify length scale parameters at the new location. Paciorek uses a latent Gaussian process to model smoothing parameters, and learns the resulting parameters through a Markov Chain Monte Carlo procedure. This interpolates local length scales and permits extrapolation of local length scales outside collected data. In contrast, we take the alternative approach of parameterizing the data likelihood in terms of each independent length scale and identifying bandwidths using conjugate gradient descent. This results in an additional covariance function parameter for each training datapoint.

However, this nonstationary model employing local length scales is somewhat overparameterized; it fits each local length-scale independently without regard to local continuity. In principle we would like length-scale parameters to vary smoothly across the region, but evidence maximization provides no *a priori* support for this intuition. Thus the local length scales tend to be noise-sensitive, differing greatly even between neighboring input points. A second issue is that in order to make a prediction we must interpolate between datapoints to find an appropriate length scale for novel inputs.

A shortcoming of the local length-scale method related to the problem of overparameterization is computational tractability. The number of free parameters increases in proportion to the number of datapoints. The likelihood is closed-form for one-dimensional inputs, but higher-dimensions require computationally-expensive MCMC estimation procedures. This makes them unsuitable for any exploration tasks that require regular re-learning of multiple input dimensions.

We solve both problems simultaneously by applying a spline regression to the result of the evidence-maximization procedure, treating length scales as a smooth function of the input. The B-spline [64] is another linear smoothing technique in which smoothed predictions at a new point  $x$  are a linear combination of  $n$  polynomial basis functions  $\beta(x)$ , each weighted by coefficients  $\alpha_j$ :

$$\hat{w}(x) = \sum_{j=1}^n \alpha_j \beta_j(x) \quad (3.29)$$

Here each  $\beta_j$  is a polynomial interpolating *knot points* placed at the location of the training data. Inversion of an  $n \times n$  matrix supplies the appropriate coefficients for predicting  $\hat{w}(x)$  at a new location. A smoothing parameter  $\lambda$  controls the degree of regularization. We determine  $\lambda$  using bounded  $1D$  function optimization to minimize leave-one-out cross-validation error. The result is a smoothly-varying local length scale that one can substitute into Equation 3.28.

The local length-scale method outperforms the stationary covariance function for the noisy step function. In the image of Figure 3.4, red circles illustrate the smoothed local length scales which shrink near the boundary to reflect the high rate of change. This lets the smoothing parameters compensate between modeling the discontinuous boundary and the smooth outer periphery. It results in a sharper step boundary and less “ringing” beyond the step border.

Other methods for modeling nonstationary covariance involve warping the input space [71]. These share many of the advantages and disadvantages of locally-varying length scales. In particular, the requirement of expensive simulation methods for parameter inference motivates us to seek alternatives.

### 3.3.2 Augmented Input Spaces

Additional input dimensions can create spatial nonstationarity without changing the form of the covariance function. This method, pioneered by Pflingsten *et al.*, augments a  $d$ -dimensional input space with an additional value that separates training data along the new dimension [72]. Pflingsten synthesizes a new input using the prediction of an auxiliary regression model. The result is a Gaussian Process in a  $d + 1$  dimensional input space that better reflects varying rates of change in the original data.

Figure 3.5 illustrates this effect visually. Here observations are spaced at some distance  $d_1$  in the original 1-dimensional input space; data appears at regular intervals along the horizontal axis representing the input  $x_i$ . This corresponds to a situation where the agent performs periodic sampling. Due to stationarity the *a priori* covariance between any two adjacent points is equivalent. Augmenting the input with a stationary prediction  $\hat{f}(x_i)$  forms a multivariate 2-dimensional input space where interpoint distances (such as  $d_2$  illustrated) vary between neighbors. This is not the

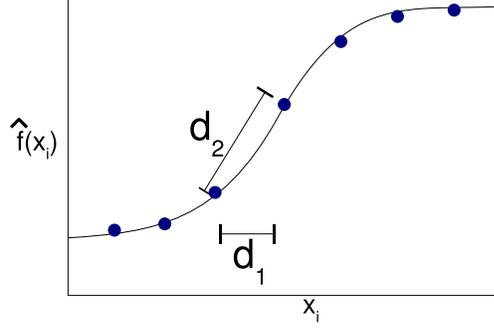


Figure 3.5: Expanding the input space with a latent prediction of the output variable permits non-stationary covariance by extending the distance between sample points along the new axis. Here  $d_1$  is the constant distance between datapoints in the original 1 dimensional input space.  $d_2$  is the alternative distance computed in the 1 + 1 dimensional input space augmented by the stationary prediction  $\hat{f}(x_i)$ .

same as crafting a nonstationary positive-definite Mercer kernel. Instead, we are literally embedding the data in a new input space where it is better-suited to the assumption of stationarity. The result is a Gaussian process with respect to the augmented data, but not necessarily with respect to the original raw input.

The raw prediction value itself suffices as a latent input, but it has the undesirable property that the input space becomes sensitive to the scaling of the output variable. For this reason a logit transform of the latent prediction improves stability [72]. We map the latent prediction to the  $[-1, 1]$  interval using the logit function so that its scale is commensurate with the (rescaled) input data, and append the result to the input space. The augmentation transforms the original  $d$ -dimensional independent variable  $x_i$  into an augmented version  $\bar{x}_i$ :

$$x_i = [x_{i1}, x_{i2}, \dots, x_{id}] \quad (3.30)$$

$$\bar{x}_i = \left[ x_{i1}, x_{i2}, \dots, x_{id}, 1 - \frac{2e^{\hat{f}(x_i)}}{1 + e^{\hat{f}(x_i)}} \right] \quad (3.31)$$

Finally, we train the hyperparameters of the main Gaussian Process model's squared exponential covariance function using the augmented input vectors. Here again we use the aforementioned gradient descent procedure.

In practice nearly any regression or smoothing technique to estimate  $\hat{f}(x_i)$  offers some advantage. Figure 3.6 shows three alternative regression models. The first is another Gaussian process using a squared exponential covariance function. We train the latent model using gradient descent and insert its MAP prediction  $\hat{\mu}(x_i)$  into the input space of the main Gaussian process.

The second latent input technique appears at lower right. Here we use a cubic B-spline [64] to smooth the input data and estimate  $\hat{f}(x_i)$  for all datapoints. This produces a smoothed prediction that we re-scale and append to the input space of the main model.

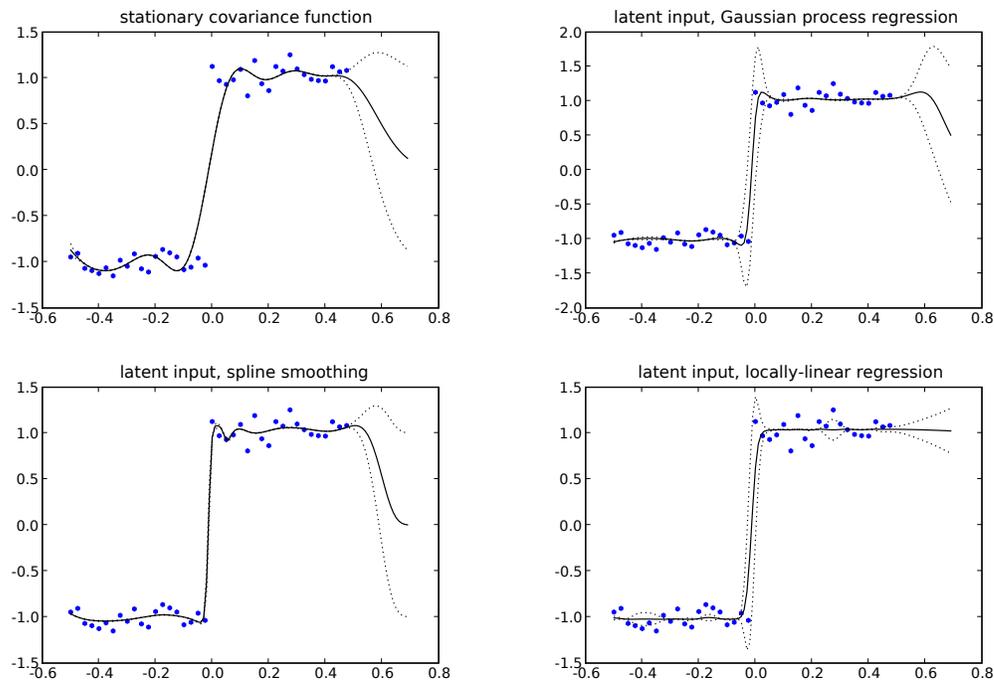


Figure 3.6: Upper Left: The noisy step function modeled with a stationary squared exponential covariance function. Upper Right: the noisy step function modeled with the nonstationary method of Pflingsten *et al.* Lower Left: latent input space generated spline smoothing regression. Lower Right: latent input generated by local linear regression.

Finally, we consider a *local linear regression* approach that predicts  $\hat{f}(x_i)$  by fitting a local one-degree polynomial. In order to evaluate the function at a new prediction point we perform a linear regression with training data weighted according to a local kernel function. The advantage to this procedure is that, unlike Gaussian process and spline regression methods, it tends to perform well at the extremes of the input space [64]. Local linear regression reduces boundary bias which is important for exploration applications that habitually predict observations near the periphery of collected data.

Our local linear regression uses a Gaussian kernel that presumes a single length scale for all input dimensions. This length-scale is initialized using the normal reference rule [64] to identify an approximate range. Then, we tune it more precisely using bounded  $1D$  approximation, driven again by squared cross-validation error. This produces an estimator  $\hat{f}(x_i)$  that we can use to reconstruct the original dataset.

Regression predictions are not the only source of latent inputs for the site survey task. Explicit predictions of  $\hat{f}(x_i)$  can be superfluous if additional *a priori* measurements such as remote sensing data is available. These spread the input points in the joint space of positions and *a priori* observations, resulting in models that are spatially nonstationary. Choosing the number and type of latent inputs is akin to a feature selection problem. Later chapters will investigate the use of orbital data in this context; for now we will limit ourself to a single latent input  $\hat{f}(x_i)$  produced by the three regression strategies described above.

### 3.3.3 Evaluation on Synthetic Data

A series of tests on synthetic data evaluate these stationary and latent-space models. These tests provide an idealized case for comparing the behavior of the different covariance functions in stationary and nonstationary environments. Our first case considers the noisy step function portrayed in figure 3.4. We generate data by assigning a random location to the boundary discontinuity, and produce latent inputs using various regression models to compute  $\hat{f}(x_i)$ . We produce performance scores by evaluating the MAP estimate  $\hat{\mu}(x_i)$  of the main Gaussian process prediction for each datapoint and compute mean-squared reconstruction error using leave-one-out cross validation.

Table 3.1 shows the reconstruction error for 100 trials. Several trends are apparent. First, all nonstationary methods outperform the stationary covariance model at reconstructing the underlying step function. The ability to exhibit variable rates of change permits them to better model both the step discontinuity and the areas of homogeneity.

As observation noise increases, so does the performance gap. Figure 3.7 shows reconstruction performance as a function of noise, with error bars indicating  $\alpha = 0.05$  confidence intervals. Much of the reconstruction error is concentrated around the border. This is evidenced in Figure 3.8, which shows reconstruction performance for the quintile of the input space centered on the discontinuity itself. The average error of the stationary model is much higher here, corroborating our visual intuition that the stationary covariance function is ill-suited to modeling the step.

It is also useful to examine the behavior of the different nonstationary models at the extremes. Exploration tasks routinely involve inference at the sparse periphery of the data, so good extrapo-

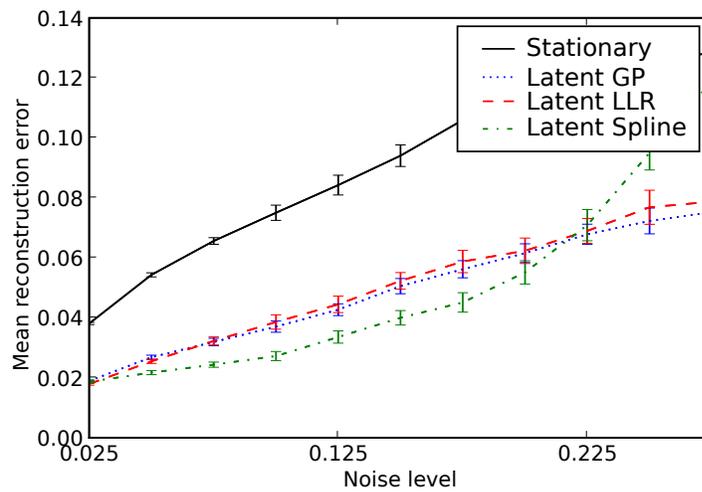


Figure 3.7: Reconstruction performance for the noisy step function using stationary and nonstationary methods for different noise values. 100 Trials. Error bars show  $\alpha = 0.05$  confidence intervals for the mean.

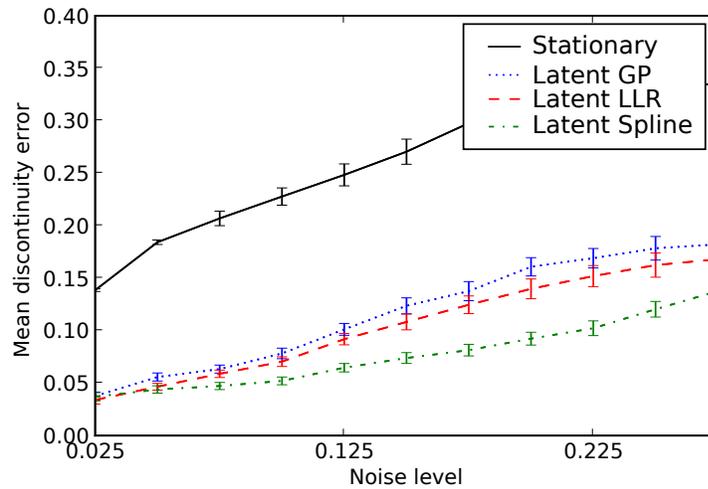


Figure 3.8: Reconstruction performance for the noisy step function for the quintile of the input space centered on the step discontinuity.

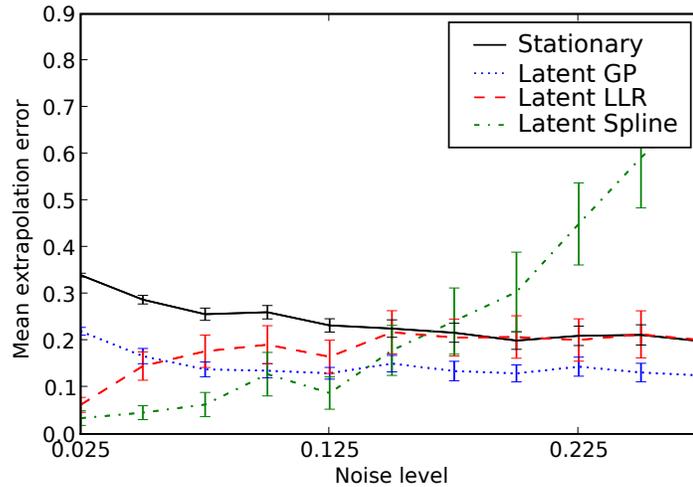


Figure 3.9: Extrapolation performance for the noisy step function for the quintile of the input space centered on the right edge of the training data.

Covariance	reconstruction	discontinuity	extrapolation
Stationary	$0.09 \pm 0.00$	$0.23 \pm 0.01$	$0.27 \pm 0.01$
Latent GP	$0.05 \pm 0.00$	$0.11 \pm 0.01$	$0.12 \pm 0.01$
Latent LLR	$0.05 \pm 0.00$	$0.19 \pm 0.03$	$0.10 \pm 0.00$
Latent Spline	$0.04 \pm 0.00$	$0.16 \pm 0.04$	$0.07 \pm 0.00$

Table 3.1: Mean absolute error in reconstruction for observations of the noisy step function (100 trials, noise = 0.150)

lation performance is important. Figure 3.9 shows reconstruction performance for the quintile of the input space centered on the rightmost edge of the training points.

The locally-linear regression model satisfies our expectations about good boundary properties - it properly extends the latent prediction according to the local polynomial fit. However, the latent Gaussian process model performs surprisingly well here as well, especially at higher noise levels where its length-scales grow longer. The B-spline extrapolates well when noise is low, but exhibits bizarre and unpredictable behavior at high noise values.

A second series of trials evaluated the different covariance functions on a continuously-varying environment where the assumption of stationarity holds better. Here we generate data sets by sampling from a sine curve, and evaluate performance at reconstructing this function from noise-perturbed observations. Unsurprisingly, the stationary covariance function has no significant disadvantage in this environment. The underlying function varies smoothly, and the stationary model with fewer “degrees of freedom” actually offers better generalization performance than the latent-

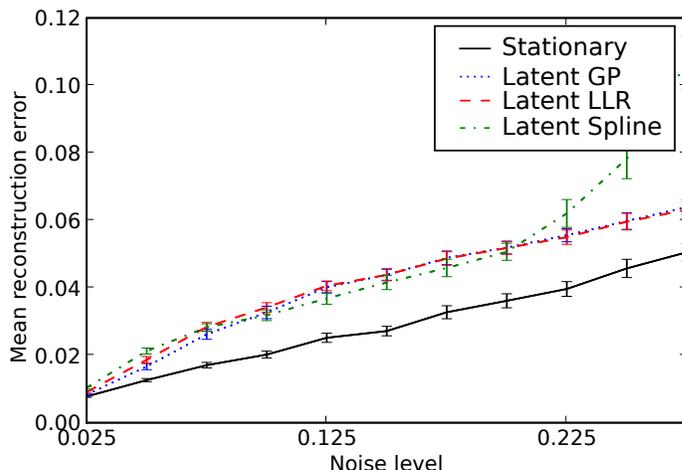


Figure 3.10: Reconstruction performance for the sine function using stationary and nonstationary methods for different noise values. 100 Trials.

space alternatives. Figure 3.10 shows the reconstruction performance as a function of noise.

In summary, the synthetic data suggests latent inputs are a useful tool for extending basic stationary covariance functions to nonstationary environments. The latent predictions can come from any source, such as a smoothing algorithm like local-linear regression or a separate Gaussian process. Here, the modeling power comes at a cost — the additional flexibility can produce additional error if  $f(x)$  is actually stationary. Regardless of the latent input’s source, the designer must be attentive to the extrapolation and generalization properties of the model. The appropriate input space for a particular task resembles a model selection question that is best addressed by standard cross validation methods.

In the next section we will examine a true intelligent mapping scenario in which the latent input is taken from remote sensing data that correlates with the desired observations.

### 3.4 Field Tests

Field data is ultimately necessary to evaluate the suitability of Gaussian process models for site survey tasks. To this end we tested reconstruction error on more realistic data from field sources. These tests demonstrate Gaussian process models that capture the spatial and cross-sensor correlations of a rover traverse dataset.

The following section describes field data collection at Amboy Crater, a lava flow in California’s Mojave Desert. The rover employs its Visible Near-Infrared Reflectance Spectrometer and learns to correlate the spectrometer signal with remote sensing data. The rover system begins with virtually no prior information about the correlations between these data products; instead, it learns

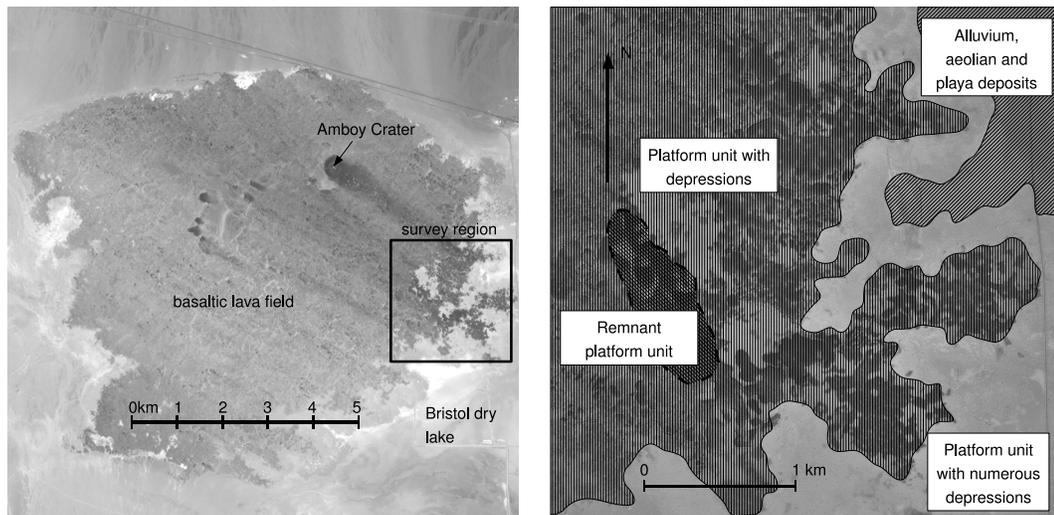


Figure 3.11: Left: An ASTER VIS/NIR image (band 1) shows the Amboy Crater lava field. The field itself measures approximately 2km in length. Right: Geologic map from Hatheway showing several distinct geologic units within the survey area on the East side of the lava flow. The background image is a USGS Digital Orthophoto Quadrangle product.

the appropriate model parameters on the fly.

We begin with a brief overview of the geology of the Amboy Crater lava field. We describe the specific exploration objectives and the experimental procedure for data collection. Analysis of the rover data suggests that latent inputs in the form of remote sensing data offer a significant benefit to the fidelity of automatically-generated surficial geologic maps. They also demonstrate that autonomous spectrum interpretation can recover the principal distinctions in surface material that are apparent from visual inspection and remote sensing imagery.

### 3.4.1 The Amboy Crater Lava Field

Amboy Crater is a cinder cone crater in the Mojave Desert of California. It is surrounded by basaltic lava flow of approximately  $70km^2$  in area that consists of layered vesicular pahoehoe lavas deposited over several distinct volcanic events. Figure 3.11 shows an orbital view of the crater area product from the ASTER instrument. The principal volcanic crater is clearly visible at the North end of the flow; note also the characteristic “wind streak” extending East from the crater.

Amboy is especially interesting to planetary geologists because the flows may be analogous to surfaces on the Moon and Mars [73]. Multiple studies over the past several decades have investigated the flow for various purposes related to planetary science and exploration. Greeley *et al.* consider the wind streak as a possible analog of aeolian features visible in Mars remote sensing data [73].



Figure 3.12: Zoë at the Amboy Crater lava field, approaching a series of basalt lava mounds.

The site has also been used for engineering and technology development experiments. The Marsokhod rover investigated Mars-relevant operational modes for remote science that incorporated techniques like regular command cycles and visual servoing [74]. Drop tests by Greeley and Iversen tested penetrator sensors for a potential Mars exploration mission. Recently, the Hyperion rover demonstrated mobility on the surface of the lava flow to show feasibility of the design for Lunar and Mars environments [75].

The flows themselves are generally composed of layered “platform units” with relatively flat surfaces; it is thought that these are the result of stagnant degassed lavas that cooled and solidified with very little sideways motion. The surface of these flows bears a strong resemblance to the planetary surfaces imaged by landers during the Viking mission.

Aeolian sediment deposits fill the low-lying areas throughout the flow. These sediment patches become more prominent as one travels East. The basalt units near this Eastern edge of the flow become progressively older [73]. Aeolian weathering processes are prominent, and sand-blasting features are common. Figure 3.12 shows the east edge of the lava field where highly eroded basalt mounds protrude from the sediment basin. Sparse vegetation is present throughout the area.

In contrast with the majority of the flow, which exhibits a fairly undifferentiated “hummocky” surface, the lava platforms near the East edge are characterized by regular “collapse depressions” with a distinctive circular shapes. Despite the name, it is possible that these depressions may have been caused by inflation of a plastic crust around a gap in the flow. These Eastern flows halt abruptly in a contact with the Bristol Lake Playa, with broken edges characterized by distinctive, heterogeneous platform shapes (Figure 3.11).

In this area the flow's surface comprises several distinct geologic units; we concentrate our studies in a particularly diverse area within the East end. Figure 3.11 (Left) shows its general position within the flow. Figure 3.11 (Right) shows a closer view, with geologic unit boundaries as identified by Hatheway and reprinted in Greeley *et al* [76, 73]. Erosion is prominent along the edges of the platform units; this, together with aeolian deposits into the depressed areas of the flow, creates a distinctive terrain characterized by isolated basalt mound features.

This latter figure employs a USGS Digital Orthophoto Quadrangle (DOQ) data product [77]. DOQ images are produced by mosaicing visible-spectrum photography from aerial overflight that has been registered to manually-placed ground control points and other known features. This results in a high registration accuracy. During our tests, we regularly observed registration accuracy within the margin of error of our commercial GPS units.

The precise location of the geologic boundaries in figure 3.11 should be considered uncertain due to the imprecise correlation between Hatheway's original map and the high-resolution DOQ data product. More significantly, it is unclear whether the platforms in this area all resulted from the same volcanic event or whether they originate from multiple different flows.

Our studies at Amboy focus on the task of surficial mapping, i.e. charting the different units of surface material present at various locations in the Eastern Amboy lava field. Relevant questions for such an investigation might include:

- What is the nature and extent of the basalt "mound" features? Of which flows are they a part?
- What is the nature of the erosion processes at work in the Eastern end of the lava flow? Are the basalt platforms themselves homogeneous, or is there variation in weathering and alteration processes across their surfaces?
- What is the relationship between orbital imagery and the morphology of the basalt platforms as visible from surface sensors? In particular, how do the orbital images relate to the density of the basalt, and the relative proportions in the mixture of basalt and sediment material that is present at a given location?

While we cannot conclusively resolve these issues during the scope of a single field season, they serve as important motivating questions suggesting several technology objectives which our rover maps should achieve. The maps should encompass enough terrain to contain meaningful data about a large region that could potentially contain one or more unit boundaries. The on-board instruments should be sensitive to the various compositions of surface material, and the map should tie these readings to appropriate remote sensing data.

The following experiments aim address these questions through a surficial mapping procedure involving Infrared Reflectance Spectrometer deployed dynamically by a mobile rover platform. The Gaussian process model learns the appropriate correlations between data products, extrapolating an MAP map estimate beyond the collected observations.

### 3.4.2 Rover Instruments

The field tests at Amboy crater utilize the Zoë rover described previously in Chapter 2. As before Zoë’s onboard science package includes a Visible Near-Infrared (VIS/NIR) reflectance spectrometer that can collect data from distant targets under solar illumination. The spectrometer’s objective lens, or *foreoptic*, mounts to a pan-tilt servo unit fixed to the rover mast. A calibration target mounted to the deck provides a white reference to determine targets’ absolute reflectivities despite lighting changes.

Basalt patches in this exploration area are dense enough to obviate the need for rock detection or targeted spectroscopy. During normal operation the reflectance spectrometer aims at a fixed  $-30^\circ$  angle of declination at the ground directly in front of the rover. At this range its field of view projects to an ellipse measuring approximately 14 and 6 centimeters along the long and short axes. One can easily distinguish the two surface materials with their VISNIR spectra (Figure 3.13). The Basalt is less reflective than the sediment, and exhibits a “bluer” signature with a larger proportion of reflected light in the short visible wavelengths. The clay signature is more reflective with a higher proportion of red components.

We perform several transformations on the raw spectrum to produce a geologic classification for our predictive map. In practice the absolute reflectance is a poor indicator because it drifts quite quickly after calibration. Instead we extract an attribute vector based on the *relative* amplitudes of five band windows chosen from the visible spectrum. The raw spectrometer data gives wavelengths down to  $2500nm$ . Our windows span the range from 350 – 1000nm, avoiding water absorption bands while capturing the visible-spectrum diversity of surface materials. We average the signal within each window, and re-scale the results to the  $[0, 1]$  interval. The result is a  $5D$  attribute vector describing the shape of the spectrum in the visible and near-infrared.

To summarize, a spectral signal at location  $x$  is given by  $s(x) = [s_{350}(x), \dots, s_{2500}(x)]$ . We produce the rescaled attribute vector  $s(x)_{\text{rescaled}}$ :

$$s_{\text{reduced}}(x) = \left[ \frac{\sum_{i=350}^{450} s_i(x)}{100}, \frac{\sum_{i=450}^{550} s_i(x)}{100}, \frac{\sum_{i=550}^{700} s_i(x)}{150}, \frac{\sum_{i=700}^{850} s_i(x)}{150}, \frac{\sum_{i=850}^{1000} s_i(x)}{150} \right] \quad (3.32)$$

$$s_{\text{rescaled}}(x) = \frac{s_{\text{reduced}}(x) - \min(s_{\text{reduced}}(x))}{\max s_{\text{reduced}}(x) - \min s_{\text{reduced}}(x)} \quad (3.33)$$

Reflectance spectra of macroscopic mixtures are generally linear combinations of the components [56]. In our case spectra describe a mixture of two physical materials (basalt and sediment) and their intrinsic dimensionality is basically one dimensional. This corresponds to the first principal component of the  $5D$  vector. Linear dimensionality recovers the principal distinction to yield the desired scalar observation: the basis derived from a linear projection that permits the best possible reconstruction of collected spectra. For the sample mean  $m(X)$  of all observed examples  $s_{\text{rescaled}}(X)$ , and the normalized first principal component  $b(X)$  given by the principal eigenvector of the covariance matrix of the mean-zero examples of  $s_{\text{rescaled}}(X)$ , we compute:

$$y = (s_{\text{rescaled}}(x) - m(X))b(X) \quad (3.34)$$

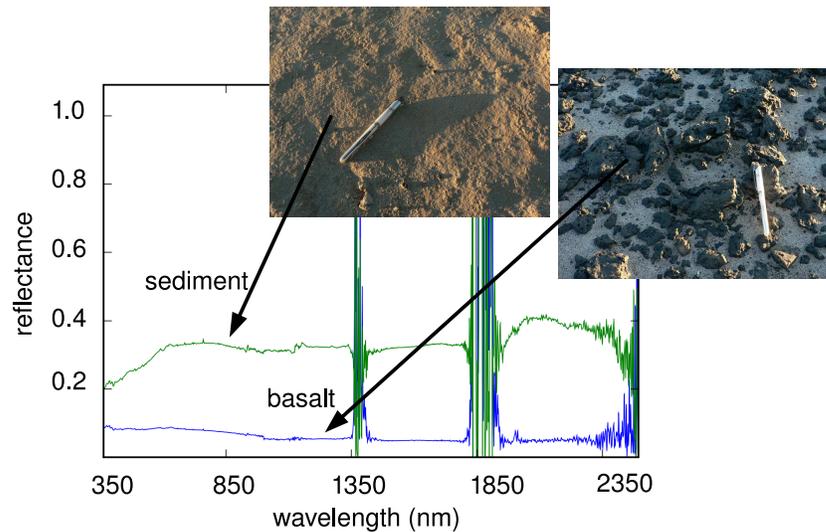


Figure 3.13: Spectra of basalt (blue - bottom line) and clay (green - top line) at Amboy Crater under solar illumination. High-noise ranges are caused by water absorption bands in the atmosphere.

A single-output Gaussian process — which predicts only scalar observations — can now model these spectra effectively. This method does not require labeled samples or other *a priori* knowledge about the spectra the robot will encounter. Instead, the rover identifies the principal distinctions on the fly.

We will consider remote sensing data in addition to the *in situ* spectroscopy. Our experiments use the remote sensing data product based on the USGS Digital Orthophoto Quadrangle (DOQ). This single-channel overflight image is sensitive in the visible spectrum [77]. The resolution of DOQ data is 1 meter per pixel, making it analogous to high-resolution instruments such as the HiRISE camera [78]. Registration accuracy is negligible, but we apply a 1-pixel Gaussian blur in order to compensate for localization drift. This amounts to a low-pass filter that preserves coarse information while excluding fine boundary detail that is likely to be inaccurate.

### 3.4.3 Procedure

During data collection the rover navigates between prescribed waypoints. It acquires spectra at 7 second intervals during forward motion, resulting in periodic measurements of surface material from the rover path. An automatic servoing procedure re-optimized spectrometer calibrations against white and dark references every 7 minutes using the scripted sequence described in Chapter 2.3.2. This procedure, along with a concerted effort to perform runs at midday in good weather conditions, helped minimize the effects of varying light levels on the spectroscopy data. Despite these efforts there were occasional modulations of the spectrometer signal from clouds and other atmospheric phenomena. Other outlier signals resulted from the rare instances where the spectrom-

eter field of view fell on a salt deposit or a plant.

We draw our data from a series of traverses covering the “Bonnie and Clyde” site at the east end of the Amboy lava field. The complete set of experiments included 20 traverse attempts over the course of three days. From this set we will consider four representative traverses from three separate days. Our selection is based on several criteria. First, we will only consider traverses where the rover finished the run in a timely fashion without significant operator intervention or navigation errors. We will also exclude traverses that contain obvious mis-calibrations or for which abrupt clouds lighting changes disrupt the spectrometer data. These four traverses constitute a representative mix of datasets that represented a diverse selection of navigation strategies, while representing performance possibilities of a mature rover system operating under nominal conditions.

The traverses appear in Figure 3.14 - 3.25. They span an array of different navigation strategies, including fixed coverage pattern and adaptive navigation strategies to be discussed in chapter 4. The traverses are approximately 500m in length, and cover a terrain area approximately  $100m \times 300m$  in size. Trials were performed over the course of three days. The rover was fully autonomous during the trials, each of which lasted approximately 20-30 minutes.

Our analysis evaluates model predictions using the partial datasets that are available to the rover during the traverses. Each prediction trial corresponds to a particular rover location. For each trial we construct a training set consisting of the previous 70 observations and train Gaussian process hyperparameters using maximum likelihood. Finally, the model extrapolates from the training set to predict a future observation at a location several samples into the future. The only free parameter is the observation noise, which we set manually to a value of 0.1 for all tests.

We varied several experimental parameters during the trials. First, we tried various combinations of input data. One trial’s input space consists only of the samples’ temporal positions in the sequence, creating a time series problem similar to those of the previous section. A more sophisticated model combines temporal position and remote sensing data.

In addition we tested each model on various prediction horizons. The far-horizon case inferred distant observations 20 samples into the future, corresponding to distance up to 100 meters depending on the rover’s future trajectory. On the opposite extreme a near-term prediction estimated the very next observation in the sequence. These tasks correspond roughly to different inference tasks that the robot might be expected to perform during remote exploration: the distant prediction is relevant to large-scale navigation issues while the proximal case is germane to issues of instrument deployment for a nearby target.

#### 3.4.4 Results

The first traverse followed an adaptive sampling strategy that resulted in a fairly even coverage of the traverse area. Figure 3.14 (Left) plots the remote sensing data and the surface measurements at each location. The color of each dot (from green to yellow) shows the scalar geologic class derived from spectroscopy. Qualitatively, green corresponds to a high percentage of sediment while yellow corresponds to high basalt signal. The unsupervised dimensionality reduction procedure automatically retrieves this distinction.

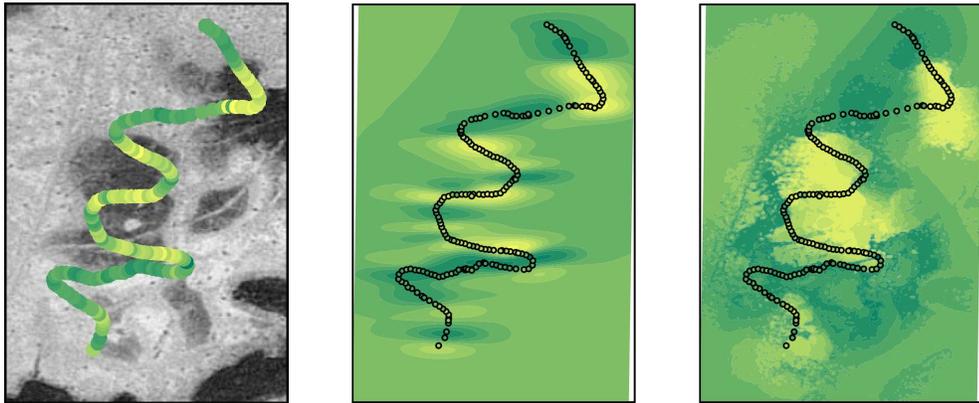


Figure 3.14: First traverse from the “Bonnie and Clyde” site. Left: Spectrometer measurements are overlaid on the DOQ orthophoto product that comprises the remote sensing data from the experiments. The color of each dot (from green to yellow) shows scalar geologic class derived from spectroscopy. Center: Map reconstruction using position information. Right: Map reconstruction using position and DOQ data.

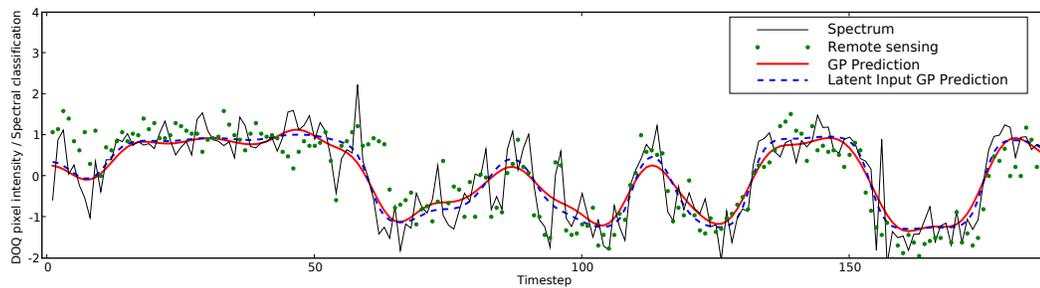


Figure 3.15: Time series plot of spectrometer classifications and remote sensing data. The first traverse exhibits a strong ( $r = 79$ ) correlation.

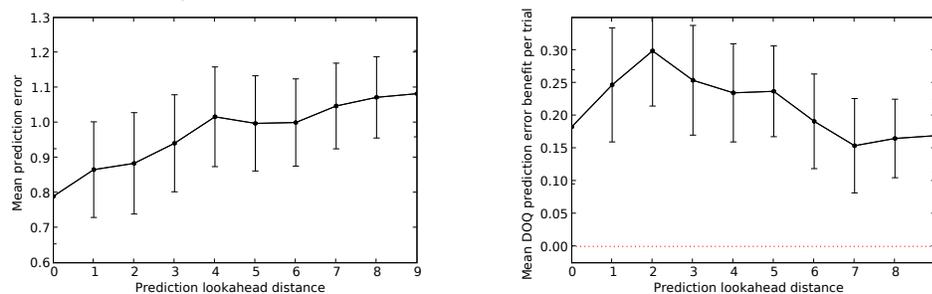


Figure 3.16: Left: Prediction performance plotted by lookahead distance. Right: Mean improvement per trial in prediction accuracy using remote sensing data.

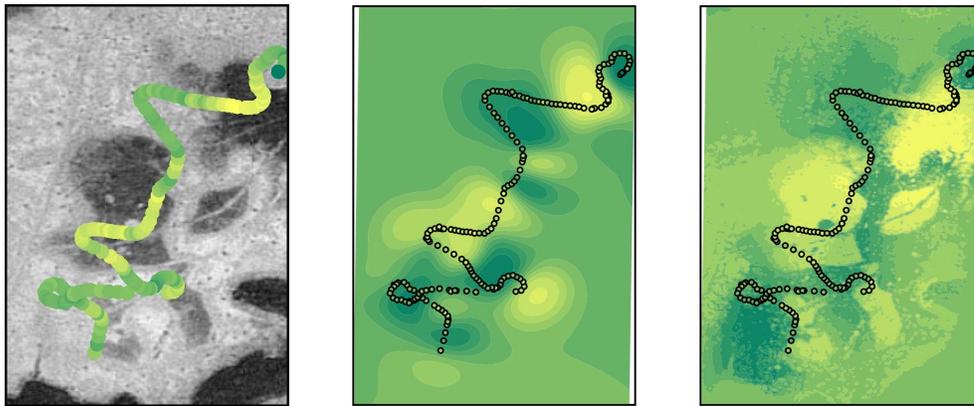


Figure 3.17: Second traverse from the “Bonnie and Clyde” site. Left: Spectrometer measurements overlaid on the DOQ orthophoto product. Center: Map reconstruction using position measurements only. Right: Map reconstruction utilizing position and remote sensing data.

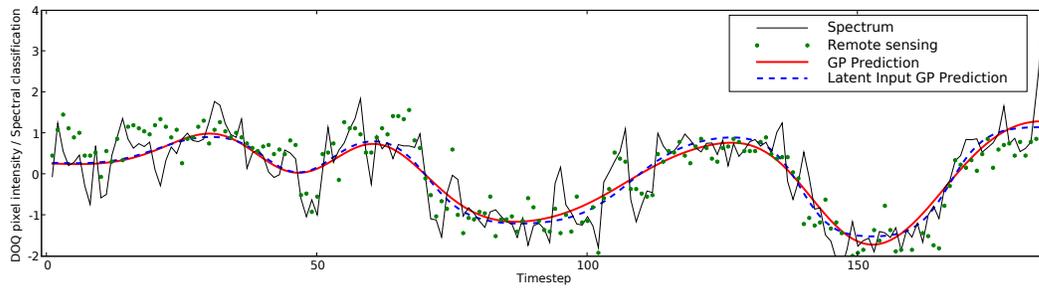


Figure 3.18: Time series plot of spectrometer classifications and remote sensing data

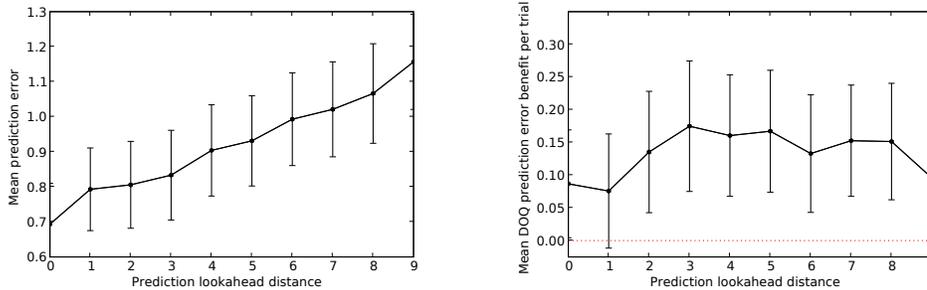


Figure 3.19: Left: Prediction performance plotted by lookahead distance. Right: Mean improvement per trial in prediction accuracy using remote sensing data.

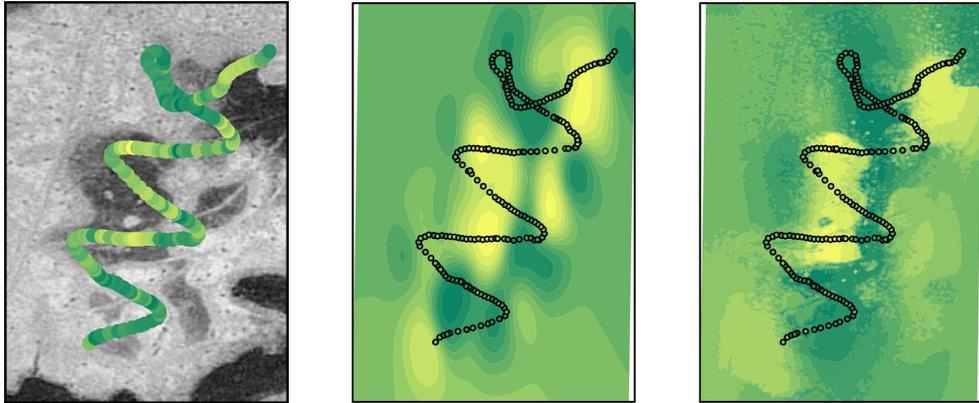


Figure 3.20: Third traverse from the “Bonnie and Clyde” site. Left: Spectrometer measurements overlaid on the DOQ orthophoto product. Center: Map reconstruction using position only. Right: Map reconstruction utilizing position and remote sensing data.

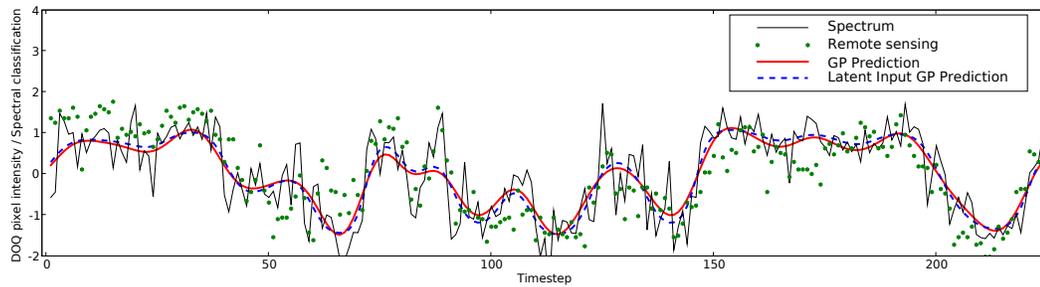


Figure 3.21: Time series plot of spectrometer classifications and remote sensing data

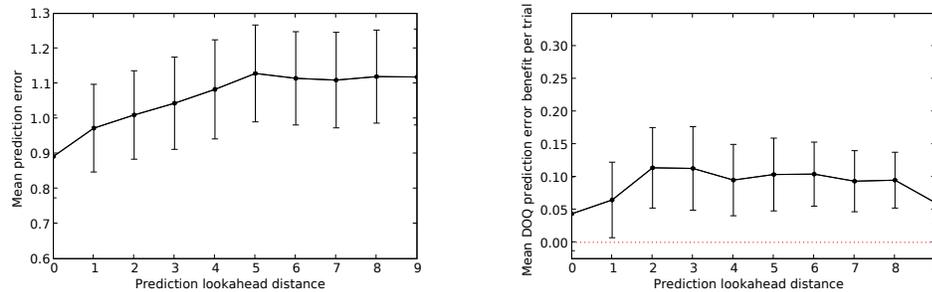


Figure 3.22: Left: Prediction performance plotted by lookahead distance. Right: Mean improvement per trial in prediction accuracy using remote sensing data.

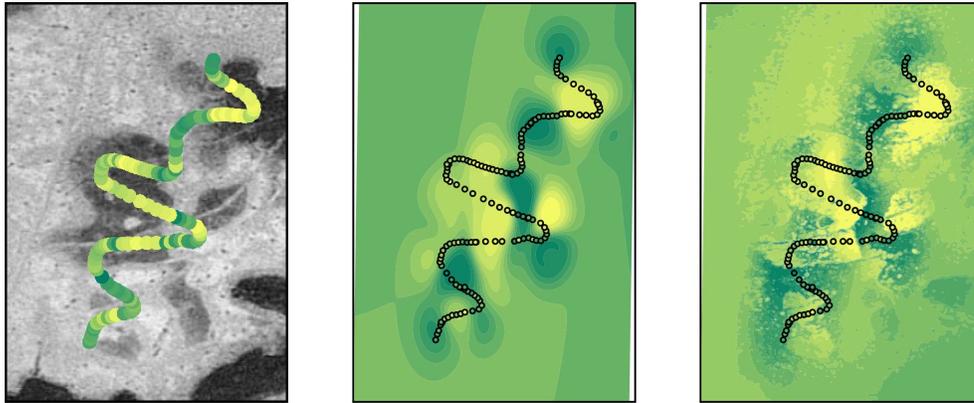


Figure 3.23: Final traverse from the “Bonnie and Clyde” site. Left: Spectrometer measurements overlaid on the DOQ orthophoto product. Center: Map reconstruction using position only. Right: Map reconstruction utilizing position and remote sensing data.

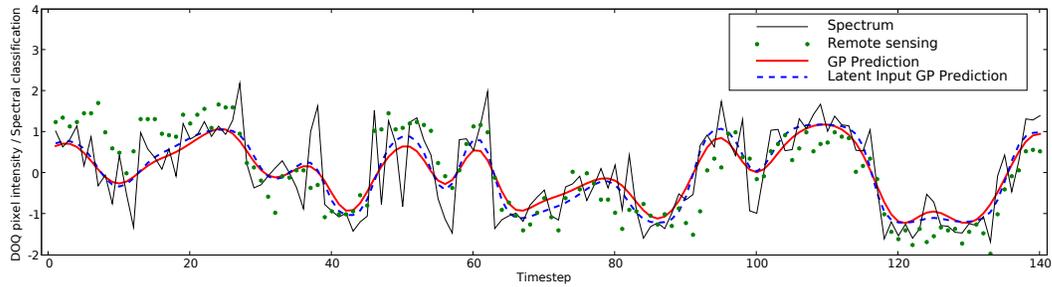


Figure 3.24: Time series plot of spectrometer classifications and remote sensing data

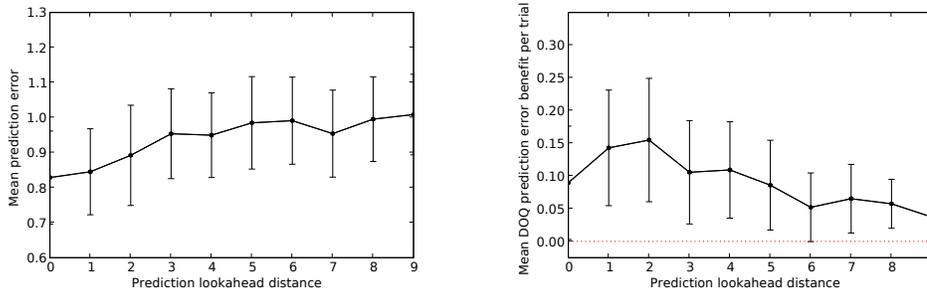


Figure 3.25: Left: Prediction performance plotted by lookahead distance. Right: Mean improvement per trial in prediction accuracy using remote sensing data.

Figure 3.14 (Center) shows the MAP map estimate using a Gaussian process model trained to predict the spectral classification based only on the physical position of each sample. We extrapolate observed values to unobserved locations with the mean prediction estimate of the Gaussian process at that location in the map given the entire set of spectrometer observations and their positions. The Gaussian process identifies a spatial correlation and extrapolates uniformly from the sampling sites.

In contrast, the illustration at right shows the improved reconstruction gleaned by including the DOQ remote sensing data product as a latent input. The model leverages the strong correlation between its learned spectroscopy classification and the associated pixel intensity of the DOQ image. This results in sharper boundaries and more faithful extrapolation. Note that the MAP prediction falls to the mean value near the edges of the map. This occurs because spatial and remote sensing correlations are both significant for this dataset. These locations are spatially removed from any other samples, the model is “less certain” about its predictions in this area.

Figure 3.15 shows the same problem represented as a time series, with pixel intensity values in the DOQ image and the spectrum classifications derived from spectroscopy. We normalize both input dimensions so that the training data has mean zero and unit standard deviation. The distribution of terrain types is approximately bimodal, reflecting the two discrete surface unit types. Time series plots show the sharp, discrete transitions from one unit to the next as the rover travels forward.

The center time series plot also shows interpolated values of  $\hat{f}(x)$  from two Gaussian process models using the temporal position of the sample as the input space. The red solid line is the estimate generated by a stationary covariance model, and the blue dashed line indicates the prediction resulting from an additional latent input from a stationary Gaussian process prediction. This demonstrates similar behavior to that of the noisy step function; the latent input model does a better job at following discrete state transitions. This benefit is not significant for reconstruction error. Moreover, any visible difference between the predictions is further mitigated by the addition of the DOQ remote sensing data. Since the DOQ data correlates strongly with the spectrometer signal it functions as a latent input and renders the additional latent Gaussian process redundant for capturing discrete boundaries.

There is a strong correlation ( $r = 0.75$ ) between the rescaled orbital intensity values and spectrum classifications. This suggests that the autonomous spectrum classification procedure recovers the environment structure reflected in qualitative observations and remote sensing data. The system retrieves this distinction without any access to spectral libraries, reference samples, or explicit hand-labeled supervised training data: the only human-selected component is the set of bandwidth windows that comprise the uncompressed attribute vector.

Figure 3.16 (Right) shows the prediction performance for various lookahead distances. We calculate mean absolute reconstruction error across all locations; error increases more or less monotonically with extrapolation distance.

The Gaussian process gleaned a significant improvement from the addition of remote sensing data. Figure 3.16 (Left) evidences this benefit; it shows the mean per-trial improvement achieved by augmenting the input space with remote sensing data. Remote sensing offers considerable

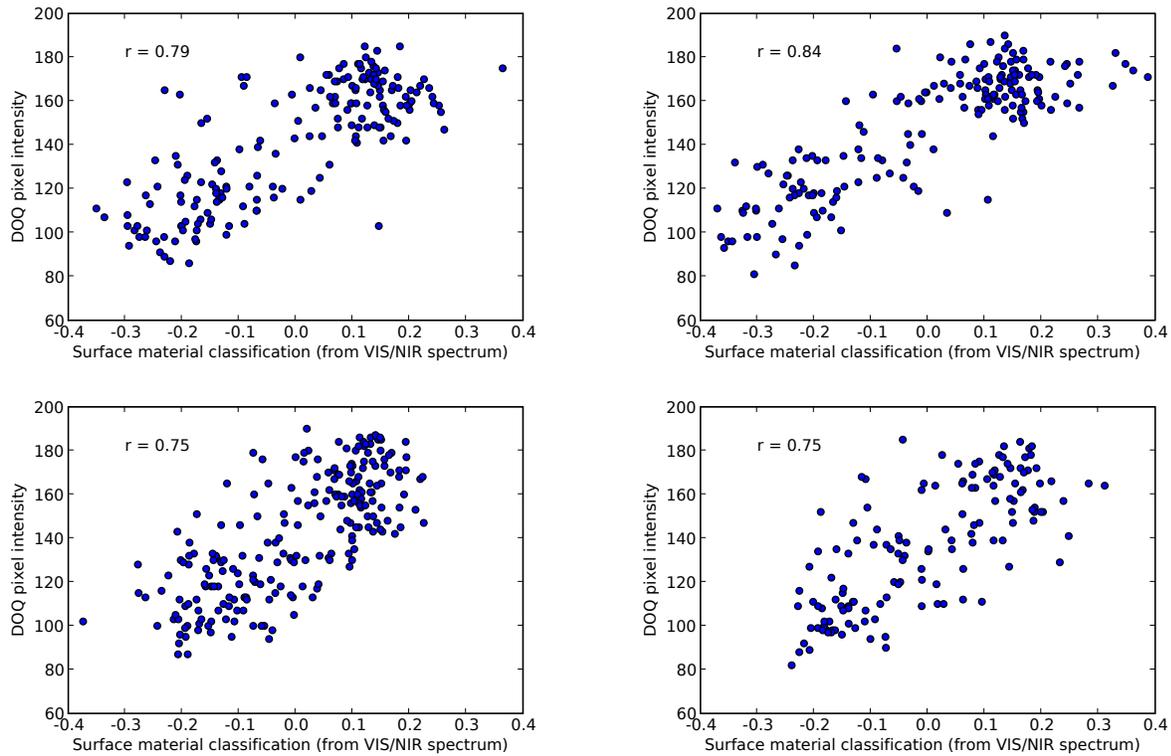


Figure 3.26: Spectrometer readings correlate strongly with remote sensing data in each of the four runs. Upper Left: First traverse. Upper Right: Second traverse. Lower Left: Third traverse. Lower Right: Fourth traverse.

Traverse	Number of Samples	Correlation Coefficient	p-value
1	186	0.79	$p < 0.01$
2	184	0.84	$p < 0.01$
3	224	0.75	$p < 0.01$
4	130	0.75	$p < 0.01$

Table 3.2: Autonomous interpretation of spectra produces a classification value that correlates strongly with the normalized albedo of the DOQ image in each of the four runs.

improvement near the rover, with benefits peaking for a prediction 2 samples into the future. The latent input continues to offer significant improvements up to 6 samples forward, after which the sample is presumably too far away in the joint position/remote sensing space to benefit from the accumulated data.

Figures 3.4.3 through 3.25 show similar results for the remaining three traverses. The different paths result from coverage patterns and adaptive sampling routines that were occasionally perturbed by navigation errors. However, all traverses cover the same approximate area in the lava field. The sequence lengths range from 140 to 220 images. In each case the addition of remote sensing data significantly improves the map reconstruction result (assuming normally-distributed error, and  $\alpha = 0.05$ ) for model predictions beyond 2 steps ahead. We conclude that spatial and remote sensing data can both contribute to the fidelity of model predictions.

Each of the traverses shows a strong correlation between the autonomously-derived spectrum classes and the remote sensing pixel intensity (Figure 3.26). Plotting these two values reveals strong linear correlation (Table 3.2). It also provides visual evidence of bimodality, reflecting the two distinct terrain types present in the traverse area.

Navigation images from the traverse reaffirm that the principal component analysis approach to spectrometer classification accurately matches the qualitative interpretation of terrain types. Figure 3.27 shows mean and extreme observations from the final traverse, along with context images illustrating the terrain at each location. Here the spectrometer view touches the area in the right/center of each image. The mean image has a lower reflectivity; it is unclear whether this is due to the rover shadow or the sparse basalt visible to the right. The extreme values correspond to archetypal examples of thick sediment and thick basalt.

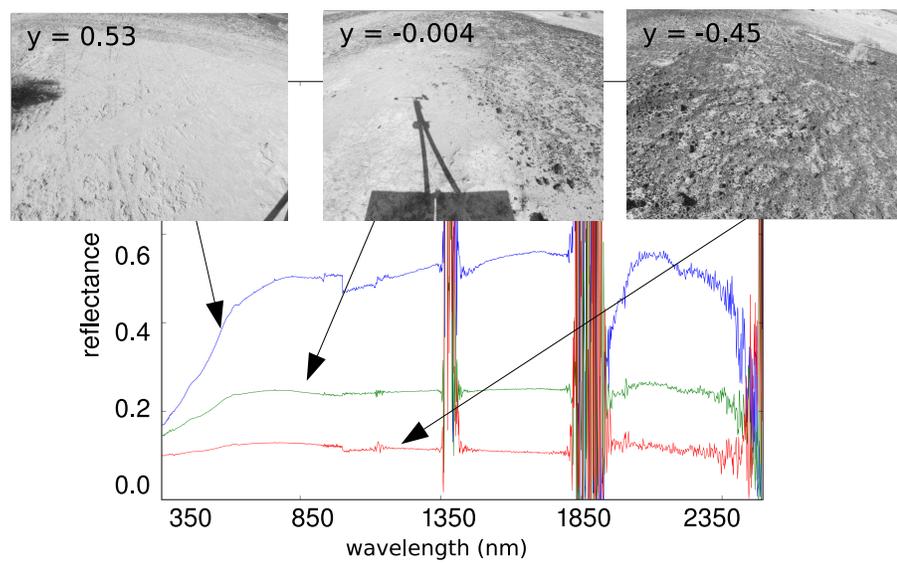


Figure 3.27: Extreme and mean observations from the traverse of Figure with associated images and spectra. We exclude noisy  $H_2O$  absorption bands near 1350, 1850, and 2350nm.



## Chapter 4

# Adaptive Data Collection

The field excursion and field class need not be concerned with a predetermined organization of observation, such as is contained in a synoptic map legend. Leads aplenty ... will turn up in the course of walking, seeing, and exchange of observation. — C. O. Sauer, *The Education of a Geographer* [58]

Science autonomy is motivated by resource bottlenecks — limits on time, energy, or bandwidth that preclude exhaustive sampling. These limits may take the form of hard constraints or incremental costs that score against the value of detected data. The *action selection* component chooses actions to improve an expected utility function related to the science value of the down-linked data and the cost of acquisition and transmission (Figure 4.1). This chapter addresses action selection strategies for adaptive data collection.

We begin by formalizing the data collection and transmission task. We present a *decoupled* solution which treats data collection and transmission as two separate, sequential optimization problems. Then we address the question of a suitable reward function for quantifying science value and survey previous science value metrics. A survey of the literature reveals a variety of options, including classic approaches such as “target signatures” and “representative sampling.” These have met with considerable success in their own domains, but are less than ideal for geologic site survey.

Instead we advocate an approach grounded in information theory. Our science value function treats remote exploration as an experimental design or active learning problem. Here the science value of a new observation is determined by the information it provides over parameters of the spatial model. Specifically we use a *spatial design* strategy that seeks subsets of observation sites that maximizes information gain about unobserved locations subject to resource constraints or incremental costs.

Simulated and field scenarios demonstrate active decision making that leverages map predictions to improve exploration efficiency. We use our science value functions for constrained sampling along a linear transect. We then introduce a *corridor exploration* scenario in which a rover plans paths to investigate a 2-dimensional exploration area. A recursive path planning strategy identifies informative paths that maximize science objectives based on the spatial modes we have

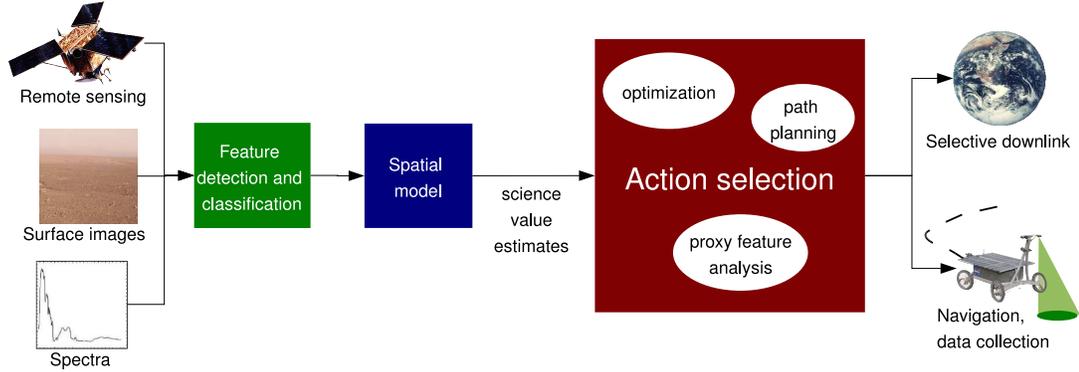


Figure 4.1: “Action selection” computes favorable collection and transmission actions.

considered in the previous section.

Finally, we demonstrate the use of adaptive science-driven navigation in real time on a field rover. Our approach permits adaptive kilometer-scale traverse at the Amboy Crater site. The rover is capable of learning to interpret spectrometer data products on the fly. This constitutes a complete closed-loop science autonomy system that discovers the principal variations in surface material and their relationship to remote sensing data.

## 4.1 Objective Functions for Remote Exploration

Meaningful resource allocation decisions require that the agent evaluate the science value of candidate datasets. We define  $\mathcal{A}$  to be the space of all possible collected data products. An element from  $\mathcal{A}$  includes the raw data itself and any associated metadata — such as the instrument status, and the time and location of acquisition — required to interpret the observation. For instance, data products could include images, spectra or other science measurements. Action selection evaluates data sets  $A$  consisting of  $n$  items from  $\mathcal{A}$ .

$$A = [a_1, a_2, \dots, a_n], \quad a_i \in \mathcal{A} \quad (4.1)$$

The complete exploration procedure consists of two sequential stages (Figure 4.2). Here  $\Theta$  represents unknown parameters of the explored environment that generate the observed features. In the initial *collection* phase the explorer agent collects a series of observations represented by the set  $A$ . In the second *transmission* phase the agent transmits a subset  $A'$  of these observations to scientists.

This two-phase description simplifies several aspects of remote exploration. By characterizing exploration as a piecewise process with a single discrete communications opportunity we preclude continuous or opportunistic communications. We must also disregard future activities beyond the next transmission phase. However, these simplifications are reasonable for our problem domain.

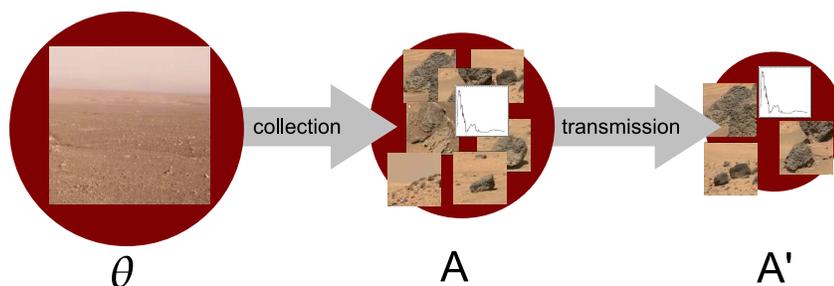


Figure 4.2: Illustration of the remote exploration task consisting of sequential data collection and transmission phases.

Each communications opportunity is a chance for scientists to change the agent’s behavior so plans over longer intervals are subject to obsolescence from new mission goals. Moreover, our scenario still captures the essential resource constraints that characterize remote exploration in general. The examples referenced in the introduction suggest that the use of discrete, sequential exploration and transmission stages pervades robot exploration in terrestrial and planetary science. Planetary exploration in particular favors discrete windows due to the periodic orbits of relay spacecraft and prescheduled timesharing on the Deep Space Network.

Our treatment makes another useful simplification, disregarding the general encoding problem of how to compress the data products themselves. Instead we will presume that the individual data products are the atomic elements of the transmission stage, and that each is associated with a fixed bandwidth cost. We should note that researchers have considered the more general problem of science-driven data product compression. For example, they have produced images that have higher resolution in areas of science interest [79]. This is beyond the scope of our work, as data product compression methods are highly dependent on the form of the data itself. Moreover, flight systems generally compress each data product independently using domain-specific algorithms. This benefits the designers that can address different data products independently, and scientists that can presume a required level of fidelity for each data product. Our method can easily accommodate the variable costs associated with differently-sized data products, and such formulations are common in the experimental design literature [80]. We consider the simpler case here for clarity.

#### 4.1.1 Coupled and Decoupled Solutions

The overall objective concerns a value function  $U$  that one can decompose into reward and cost terms [60]. A reward function  $R(A, A')$  maps the sets of collected and transmitted data to a scalar score,  $R(A, A') : \mathcal{A}^n \rightarrow \mathbb{R}$ . We include  $A$  to reflect that the value of returned data may depend on its relationship to the entire set of collected data. The use of a scalar reward score subsumes a wide range of other relevant science value classifications; scalar scores permit rank orderings of datasets, as well as binary “good / bad” classifications (provided a suitable threshold).

An observation cost  $L_O(A)$  is a scalar value accounting for resource requirements (like en-

ergy and time) required to collect the dataset. Similarly, transmission cost  $L_T(A')$  accounts for bandwidth limits. The remote explorer aims to improve a utility function  $U(A, A')$  given here as a function of the transmitted data  $A'$ .

$$U(A, A') = R(A, A') - L_O(A) - L_T(A') \text{ for } A' \subseteq A \quad (4.2)$$

The subset constraint creates a significant challenge. The explorer chooses a downlink  $A'$  using the *entire* set  $A$  that results after completing the data collection stage. This makes the contribution of each individual data product difficult to determine until all measurements are finished. During data collection, the agent must solve the entire joint collection and transmission optimization in order to evaluate any single observation. It must consider all the potential returned subsets of data that it will later collect. For example, the agent may decide to forgo collection of informative data products because they will not complement future data that would be included in the downlink.

We denote the optimal solution to this *coupled* observation/transmission problem with a subscript  $c$ . The solution defines a set of collected observations  $A_c$  and transmission subsets  $A'_c$  that maximize the joint utility function  $U(A, A')$ . During data collection the agent optimizes the expected utility of collected observations, considering all of the potential downlinks.

$$A_c = \operatorname{argmax}_A E[R(A, A') - L_O(A) - L_T(A')] \text{ for } A' \subseteq A \quad (4.3)$$

The joint optimization over  $A$  and  $A'$  requires integrating over all possible values of data that will be collected before initiating the transmission phase. We desire more tractable approximate solutions to the joint collection and transmission problem. Consider the case where we substitute a new reward function: a sum of sequential observation and transmission terms, where the reward ascribed to the transmitted data is contingent only on that data that was collected.

$$R(A, A') \equiv R_O(A) + R_T(A') \quad (4.4)$$

The agent can approximate an optimal solution by optimizing observation and transmission reward sequentially. We will use the term *decoupled* to describe this class of exploration tasks, and denote the decoupled solution with the subscript  $d$ :

$$A_d = \operatorname{argmax}_A E[R_O(A) - L_O(A)] \quad (4.5)$$

$$A'_d = \operatorname{argmax}_{A'} E[R_T(A') - L_T(A')], \quad A'_d \subseteq A_d \quad (4.6)$$

The agent pursuing this value function computes data collection actions while disregarding the transmission budget. This results in a solution set of observations  $A_d$ . Only later, during the transmission phase, does the agent assemble a transmission subset  $A'_d$ . We summarize the basic procedure in Algorithm 1.

The form of Equation 1 suggests that there are important parallels between data collection and return objectives. Nevertheless, this chapter will focus exclusively on the problem of data collection and reserve the selective data return task for Chapter 5. Our basic decoupled data collection strategy is the following. Based on the current Gaussian process hyperparameters, the agent chooses observations  $A$  to maximize a utility function where reward is based on the observations' information gain. Then during data collection, it periodically reestimates hyperparameters by training on the set of collected data and regenerates the experimental plan for the remaining resource budget.

```

Objective:  $\max_{A, A' \subseteq A} R(A, A') - L_O(A) - L_T(A')$ ;
for data collection do
   $A_d \leftarrow \operatorname{argmax}_A E [R_O(A) - L_O(A)];$ 
for data return do
   $A'_d \leftarrow \operatorname{argmax}_{A'} E [R_T(A') - L_T(A')];$ 

```

Figure 4.3: Decoupled approach to data collection and transmission.

Before describing the particulars of the algorithm we will survey previous efforts in assessing science value, formulating suitable reward functions, and resolving the optimization problem that results.

## 4.2 Previous Work in Adaptive Data Collection

Practical issues of quantifying science value make the reward function  $R_O(A)$  difficult to specify. Combinations of data products could yield arbitrary co-dependencies, redundancies, and synergistic effects. For example, a reflectance spectrum might be complementary with visual data or redundant with respect to other data products. The space of potential collected datasets,  $\mathcal{A}^n$ , is invariably too large to enumerate all these interactions. The principle challenge of representing science value is to approximate  $R_O(A)$  using factorizations that are tractable yet accurately capture the basic structure of the scientists' preferences.

Another difficulty is the inherent subjectivity of science value judgments. Preference elicitation surveys suggest that scientists can disagree significantly over the science value of collected data. Their independent mission goals may not map cleanly to a scalar reward score. Smith *et al.* surveyed a group of scientists during simulations of remote rover operations in the Atacama Desert [18]. The survey queried scientists about their preferences for different sets of surface images, and found considerable discrepancies between individuals. While scientists often professed similar preference orderings, they usually provided different rationales. Some preferred wide-field images that would provide the best information about where to travel next. Others looked for distinctive novel, unexpected features such as shadows or rock piles.

Other blind studies in planetary geologists' preferences for image sets have uncovered qualitatively similar results [81, 29]. Human scientists' value assessments may rely on current hypotheses of interest, short term objectives, or personal familiarity with the instrument. We desire reward functions that make these assumptions as explicit as possible.

On the other hand, we do not want the reward function to be so finely-tuned to the anticipated observations that it precludes serendipitous discoveries. In surface exploration to date, many of the most significant revelations involved unanticipated phenomena: consider the "blueberries" of the Opportunity landing site [82], Martian "dust devils" [83], and the foreground motion phenomenon recently evidenced in Mars Exploration Rover images [83]. These discoveries imply that the re-

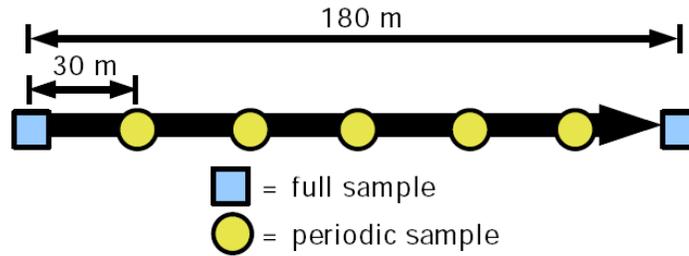


Figure 4.4: Right: Periodic sampling strategy used in Atacama Desert rover astrobiology tests. Image from [84].

ward function should be *adaptive*; in other words, it should reflect trends and anomalies in the environment and not simply the *a priori* judgment of scientists. We desire formal methods for balancing this goal against the fidelity of our value function to scientists’ expressed preferences.

### 4.2.1 Periodic Sampling

*Periodic sampling* is a common data collection strategy in which the agent makes observations that evenly cover the explored space. This is not a science reward function *per se*, and does not require any onboard data analysis. For this reason it is used as the default procedure in most remote exploration tasks and pervades previous site survey work. Periodic sampling strategies are appropriate when the task requires a dense feature map, or to produce an unbiased sample for statistical comparisons between sites. We will use periodic sampling as a baseline strategy for the comparative experiments that follow.

Often remote site surveys use periodic sampling along a “transect,” or standardized straight line path across a site. Rover expeditions to the Atacama desert, Chile provide a typical example [8]. These tests simulated site survey operations for planetary exploration with the Zoë rover platform [7, 33, 84]. The science goals were to characterize the distribution of life in the harsh desert environment, paralleling future astrobiology missions to the Mars surface. Scientists used remote sensing data to choose over-the-horizon navigation waypoints, and uploaded each day’s traverse in a daily command cycle [51].

The rover carried a microscopic imaging instrument designed to detect chlorophyll fluorescence and identify photosynthetic organisms. Remote scientists deployed the instrument according to a Standard Periodic Sampling Unit, or SPSU [62]. The SPSU consisted of 7 measurements taken at 30m intervals along a transect. The first and final sampling locations in the traverse employed a “full sample” that included more informative photosynthetic dyes with a long incubation period. Periodic sampling provided a consistent set of measurements that facilitated comparisons between locales.

The *farming pattern* is a two-dimensional generalization of the transect that attempts to cover an area of interest with evenly-spaced measurements. This strategy has been used for construct-

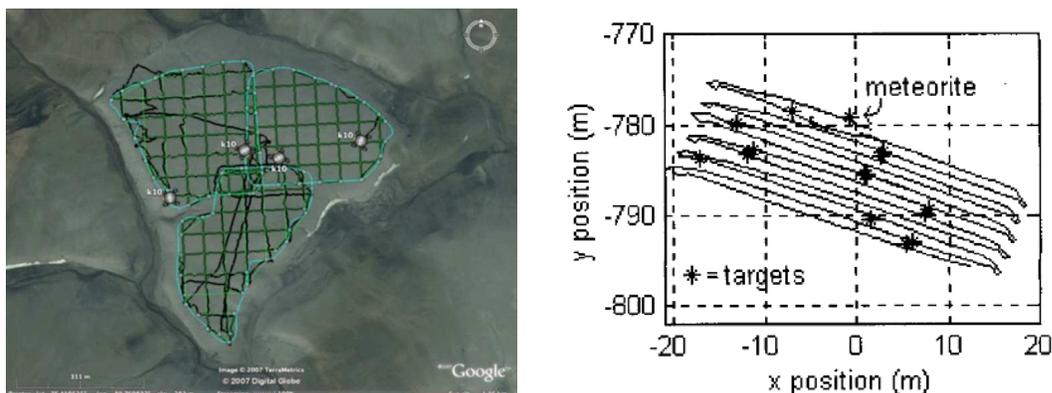


Figure 4.5: Farming patterns for rover site survey. left: Site survey of Haughton Crater; image from [6]. Right: Farming pattern for meteorite search in the search for Antarctic meteorites; image from [85].

ing dense maps with ground-penetrating radar during rover expeditions at Haughton Crater [6]. Here the application was not science-driven but rather a characterization of soil structure for base construction and engineering. This required dense, uniform coverage of the site (Figure 4.5 Left). Farming patterns were also used during the Robotic Antarctic Expedition [25]. The robot searched for rocks by traveling in parallel tracks that covered a fixed region of the Antarctic plateau [85]. In this manner operators could be sure that the robot had completely searched the exploration area (Figure 4.5 Right).

### 4.2.2 Target Signatures

The *target signatures* approach responds to observations through on-board data analysis. It pursues specific, predefined phenomena of interest [10]. Formally the target signature is a predefined subset of data products,  $T \subset \mathcal{A}$ , where the value function favors elements from  $T$ . Often  $T$  is simply a the set of target features with a specific classification. This reduces the science value function to a series of independent binary-valued questions: “is a target feature present, or not?” The value of the entire dataset is directly proportional to the number of target data products that are present.

The target signature reward function  $R_{\text{TARGET}}(A)$  totals the scores of each independent data product. We use the indicator function  $I_T(a)$  to represent membership of the data product  $a$  to the target signature set  $T$ . The reward function decomposes as follows:

$$R_{\text{TARGET}}(A) = \sum_{a \in A} I_T(a) \quad \text{where} \quad I_T(a) = \begin{cases} 1 & a \in T \\ 0 & a \notin T \end{cases} \quad (4.7)$$

A common variant [10, 86] uses several target signature classes  $\{T_1, T_2, \dots, T_m\}$  where each class  $T_i$  is associated with a different reward coefficient  $w_i \in \mathbb{R}$ . The total score for the data packet is

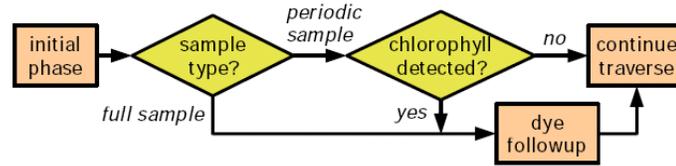


Figure 4.6: Flowchart of the target signature strategy for chlorophyll followup in Atacama rover tests [18].

equal to the sum of weighted value scores for each data product:

$$R_{\text{TARGET}}(A) = \sum_{i=1}^m w_i \sum_{a \in A} I_{T_i}(a) \quad (4.8)$$

The target signature approach has several advantages. Note that the value of the dataset is independent of the actual explored environment. Explorer agents might only see a small portion at a time, but they can determine the science value of a dataset with no penalty for this partial knowledge. Optimizing datasets involving multiple weighted target signatures is straightforward. The agent can evaluate the overall science value without considering interrelationships between data products, so the task of choosing data products to fit a resource budget is a “knapsack” problem with tractable solutions [87].

If feature detection and classification is robust, then the target signatures approach yields behavior that is conceptually simple for the user to predict. This facilitates both parameter elicitation and retrospective interpretation of system behavior. The target signature approach suggests a “treasure hunt” metaphor for the remote science, and it is a natural match when scientists are interested in a specific predefined phenomenon that is known to be rare or transient.

Field-deployed systems have demonstrated performance improvements from applying the target signatures approach to data collection. The aforementioned Atacama project compared periodic sampling to target signature that favored images containing desert organisms [84, 18]. Here the rover analyzed the output of each sample and triggered additional dye measurements only when the simpler tests showed evidence of life 4.6. Researchers noted a significant improvement in the number of dye-based life detection observations returned by adaptively-triggered measurements.

The Earth Observing 1 (EO-1) project deployed a target signature algorithm to an Earth-orbiting spacecraft as a part of the Autonomous Sciencecraft Experiment, or ASE [88, 41, 89]. In a typical ASE scenario, scientists specify a specific list of science targets which the spacecraft would monitor with regular imagery. The appearance of target phenomenon, such as a fresh lava flow from a target volcano, results in the spacecraft planning additional followup observations. An onboard planner, CASPER, schedules additional observations while respecting time, memory and power constraints. As the first comprehensive science autonomy system in flight, EO-1 demonstrated the feasibility of the target signature approach and captured several of the desired events.

Target signatures are also germane to the downlink selection we will address in the following chapter. The Mars Exploration Rovers employ target signatures in detecting and returning images of atmospheric phenomena [90, 83]. A pixel-differencing operation identifies dynamic atmospheric features like clouds and dust devils in images. This lets the rover capture long sequences of images and return those that contain the desired features. The procedure maximizes the number of appearances of the target signature for a given fixed-size downlink budget.

While the target signatures approach has conceptual and practical advantages, some obvious limitations prevent its application in more complex science tasks such as site survey. Site survey involves characterizing a remote environment rather than simply seeking some predefined phenomenon. This implies that the reward function is contingent on the actual state of the environment. Field geologists performing site surveys regularly change their sampling to avoid redundant information; they must react to trends in spatial distribution, geologic class, and cross-sensor correlations. However, existing implementations of target signature rewards are inherently non-adaptive.

Similarly, the total value of a dataset is often quite different than the sum of its component features. Since contextual evidence — the scientist’s prior expectations for the site, recent data from neighboring locales and coincident sensors — is essential to interpret new data products, single data products seldom carry sufficient information to determine their own science value. Scientists may find it counterintuitive to assign numerical value scores to isolated, independent features. There is no “optimal answer” so scientists with the same goals and assumptions might disagree over parameter settings.

Revisions to the target signatures schemes posit additional reward discounts and thresholds to account for each of these shortcomings. Additions proposed by Smith *et al.* introduce a maximum quantity for any single feature in the downlink or a minimum distance between samplings of a particular feature [86]. The additions address some failure modes at the cost of additional complications in the parameter-setting process. For example, the distance threshold may encode implicit assumptions like the spatial continuity of the environment. If the actual spatial structure differs from the presumed one — more homogeneous or irregular — then the onboard science decisions may be poor. In contrast, our adaptive exploration strategy will base rewards on an explicit statistical model of the environment structure. Learning model parameters on the fly ensures that the robot’s behavior reflects the discovered structure of the environment.

### 4.2.3 Action Selection with Target Signatures

Target signatures lend themselves to straightforward action selection strategies. One natural solution treats science autonomy as consisting of occasional, opportunistic departures from a pre-planned command sequence, with a fixed reward for each opportunistic action. The OASIS project at NASA’s Jet Propulsion Laboratory demonstrated one such system which gathered data on rocks in a Mars-like test environment [10, 12]. Here onboard data analysis module identified rocks in rover navigation imagery and matched them against target profiles supplied by scientists. When a match was found, the data analysis module signaled the onboard planner with a new “science data request.” This token entered the rover’s daily plan as a new action.

The planner, based on the CASPER architecture, incorporated each science data request into the rover's daily plan while respecting mission constraints on time, energy, and the desired end-of-day location. A series of rigorous tests demonstrated rovers' abilities to identify interesting rocks and divert from their pre-scheduled path to approach these new targets. After completing the opportunistic data collection the rover returns to the original path. More recently, researchers demonstrated interrupt-driven plan alteration and repair in the context of geologic anomalies for the SARA science agent, a project underway with potential applications to the ESA Exomars rover [17]. Both cases were simplified somewhat by the fact that the rover did not generally have to weigh competing observation opportunities. There was no explicit treatment of detection noise. In other words, they assumed that each observed target signature match was correct.

An alternative to opportunistic autonomy is a branching command sequence that chooses from among scientist-specified alternatives. This restricts the scope of decision-making to discrete decision points. An example is the contingent planner/scheduler architecture (CPS) [91]. This planning system was developed along with a contingent rover planning language to guide science operations during Mars exploration. A ground-based planner would specify a contingent plan in terms of a series of branching possibilities. During execution an onboard executive follows the relevant branch of the plan.

Another recent development in planning for robot planetary science is a strategy for action selection and reordering in situations with oversubscribed resources [92]. A planner based on this strategy was recently demonstrated in autonomous instrument placement tasks for a Mars exploration scenario [15]. Here a group of humans and robots selected targets for followup analysis while another robot performed visual servoing and instrument deployment.

Recent work at the Carnegie Mellon Robotics Institute integrated a Partially-Observable Markov Decision Process (POMDP) representation [93, 18]. The POMDP planner considers observation uncertainty explicitly. Experiments demonstrated a simulated autonomous science environment where a rover attempted to verify surface targets from probabilistic orbital data. The rover navigated along an exploration corridor while moving only forward — thus, an ill-advised approach to one possible target might preclude following up on other nearby targets.

Several trial runs compared the performance of a rover equipped with a POMDP planner against purely reactive exploration strategies. By balancing the immediate benefits of pursuing a nearby target with observation uncertainty and long-term plans, the POMDP planner verified targets more efficiently than the reactive alternative. By anticipating future observation opportunities, the POMDP planner was able to keep the rover in an advantageous position to take advantage of anticipated targets. The work suggests that an explicit model of observation uncertainty can improve science autonomy performance. Moreover, it underscores the value of generative spatial models to predict collected data.

In general the experiments of Smith *et al.* complement the work we will present here. Smith considers a discrete state/action space with a fixed number of possible observations. In contrast, the Gaussian process formulation permits point features at arbitrary locations in a continuous state space of arbitrary dimensionality. Our approach will disregard issues of partial observability; we do not, for instance, take into account the ability of future measurements to influence model pa-

rameters. Instead, we focus on action selection in a spatial model so that the reward adapts over time to discovered spatial and cross-sensor trends.

#### 4.2.4 Representative Sampling

Unlike target signatures, *representative sampling* [10, 86] aims to produce *faithful* datasets. It assumes that scientists are interested in the diversity of data products; they attempt to include archetypal examples that best represent the full range of features in the environment.

Recall that each data product  $a_i$  is associated with a vector of independent attributes  $x_i \in \mathcal{X}$ . For a set  $V = \{v_1, v_2, \dots, v_n\}$ ,  $v_i \in \mathcal{X}$  of potential observations, the preferred subset  $A$  minimizes the distance of each  $v_i$  to its nearest observation.

$$R_{\text{REPRES}}(A) = \sum_{v_i \in V} \min_{x_i \in X} |x_i - v_i|^2 \quad (4.9)$$

Note that if we define  $x_i$  to be the physical location of the data product, the best solution is periodic sampling in the spatial domain. On the other hand, if  $x_i$  represents one or more prior observations then representative sampling favors an even distribution of observations from these attributes. For example, consider the decision of where to deploy a spectrometer, where  $x_i$  is a preliminary visual classification. Scientists wanting a representative sample could use this reward function to identify good spectrometer targets based on visual similarity.

The k-means algorithm [31] provides a local optimum to Equation 4.9. K-means seeks a clustering that minimizes the distance of each datapoint to its cluster centroid. A variational fitting procedure iteratively assigns each location  $v_i$  to its Euclidean-nearest class, and recalculates cluster centroids to be the means of their associated datapoints [69]. After several iterations, the algorithm converges to a local minimum of the error function: a mixture model that constitutes a compressed representation of the dataset. The algorithm’s only free parameter is the number of clusters. One can set this value either by model selection techniques (cross-validation or an appropriate substitute [64]). Alternatively, resource constraints may determine the number of collected data products and therefore the appropriate number of clusters.

In practice representative sampling methods can employ nearly any unsupervised data clustering technique. After clustering potential observations, those near cluster centroids constitute a representative subset to favor during data collection or return [10]. Representative sampling’s performance may generalize better than target signatures because the quantization need not be specified in advance; classification thresholds adapt to reflect the variance in data with new observations. In fact, we will show in the next chapter that representative sampling is a special case of our general formulation of information-optimal observation selection. Our Gaussian processes extend this idea to scenarios where data is related by other relationships beyond mixtures of clusters.

#### 4.2.5 Novelty Detection

*Novelty detection* is essentially the opposite of representative sampling. It presumes that scientists

prefer unusual or outstanding features, and triggers additional data collection when a data product’s attribute vector  $x_i$  exceeds a novelty threshold [10, 33].

Existing novelty metrics generally fall into one of a few basic categories [94]. Classical hypothesis testing flags datapoints that fail a hypothesis test, where the null hypothesis states that the data is derived from the distribution which produced a training set. Here designers set a novelty threshold using a suitable value for the confidence threshold  $\alpha$ . Other novelty metrics derive from probability density estimation: parametric or nonparametric density models characterize the entire set of collected data, and data points receive novelty scores according to their inverse likelihood. Finally, maximum-margin methods using a Support Vector Machine (SVM), use a novelty measure based on distances from a hypersphere enclosing the datapoints [95].

Smith et al. used informal surveys to evaluate these density estimation methods in a rover exploration application [18]. They subsampled sequences of rover navigation images by first detecting rocks as in [32] and then applied a hand-tuned mixture of representative sampling and novelty detection. A control group used periodic sampling that did not consider image content. Scientists involved in remote rover operations identify a preferred dataset. While the reported rationales for scientists’ preferences varied significantly, they expressed a significant overall preference for the intelligently-subsampled data. This suggests that the clustering methods captured useful properties of the environment.

Our work in this section doesn’t ascribe explicit value to novel observations, though the Maximum Entropy criterion tends to favor this data. Nevertheless one could easily introduce novelty as a parallel value function. The learned spatial model easily identifies outliers by their low likelihood.

#### 4.2.6 Learning Mixed Reward Criteria

Scientists may prefer a combination of specific targets and dataset diversity. DesJardins, Eaton and Wagstaff characterize these mixed reward functions with a language DD-PREF to describe user preferences over sets of objects [96]. A DD-PREF preference profile consists of a tuple  $P$ , where  $P = \langle q, d, w, \alpha \rangle$ . For each dimension  $x_{ij}$  of the attribute vector  $x_i$ , we associate a corresponding  $q_j$ , a density function mapping the real-valued  $X_{ij}$  to a scalar “depth” score:  $q_j(a_i) : \mathbb{R} \rightarrow [0, 1]$ . This value represents the attribute’s match to a specific value the scientists prefer. The depth value of an entire set  $A$  is the sum of average depth values for each feature, weighted by the feature coefficients  $w_j \in [0, 1]$ .

$$R_{\text{DEP}}(A) = \sum_j w_j \frac{1}{|A|} \sum_{a \in A} q_j(a_j) \quad (4.10)$$

Conversely  $d_j \in [0, 1]$  is a desired “spread” within the set for the feature’s values.

$$R_{\text{DIV}}(A) = \sum_j w_j (1 - (d_j - R_{\text{DIV}_j}(A))^2) \quad (4.11)$$

A final global parameter  $\alpha$  controls the importance of  $w_f$  factors versus  $q_f$  across all features. The final score  $R_{\text{SETPREF}}(A)$  for the dataset is computed by a linear combination of diversity and

spread scores, weighted by  $\alpha$ :

$$R_{\text{SETPRF}}(A) = (1 - \alpha)R_{\text{DEP}}(A) + \alpha R_{\text{DIV}}(A) \quad (4.12)$$

In principle setting these parameters is no more intuitive than identifying target signature coefficients. Fortunately, scientists need not specify them explicitly. A major contribution of the DesJardins work is to *learn* these parameters by extrapolation from training subsets provided by scientists. In the case of the DD-PREF language designers can leverage nonlinear optimization techniques to fit the preference profile to training sets.

Learning preferences from data is advantageous if scientists find it easier to assemble example subsets than to fix quantitative feature scores. However, the resulting science value score is only appropriate insofar as the original preference language is valid for the task at hand. DD-PREF lacks some characteristics that we want for our site survey task. In particular, scores must be specified in advance as a linear function of deviations from predefined target values and diversities. These scores are independent of the actual environment, so the preference profile cannot represent a dataset’s fidelity to the site. Attributes are evaluated independently, DD-PREF cannot capture between-feature correlations such as redundancy across different sensing modes.

#### 4.2.7 Information-Driven Reward Functions

Information-theoretic reward functions rate datasets according to the information provided about parameters of an environment model. New data is valuable insofar as it reduces our posterior uncertainty about these model parameters. Figure 4.7 depicts the strategy. Here scientists begin with a broad prior distribution  $P(\Theta)$  over parameters of interest  $\Theta$ . After observing data  $A$ , Bayes’ rule yields a posterior distribution  $P(\Theta|A)$  that is at least as certain as the prior. We quantify this uncertainty with the Shannon entropy  $H(\Theta)$  of the distribution over environment parameters  $\Theta = \{\theta : \theta \in \Omega\}$  [97].

$$H(\Theta) = - \int_{\theta \in \Omega} P(\Theta = \theta) \log P(\Theta = \theta) d\theta \quad (4.13)$$

Information-theoretic value functions favor observations with the best expected *information gain* with respect to  $H(\Theta)$ . A common way to quantify this gain is the reduction in posterior entropy  $H(\Theta|A)$  that represents the uncertainty  $H(\Theta)$  after observing  $A$ . This expresses how far the probability density “bubble” of  $P(\Theta)$  shrinks for a Bayesian posterior that incorporates the new observation.

Chaloner and Verdinelli provide an excellent review of information-driven approaches to experimental design problems [98]. Lindley first proposed to treat experimental design decisions in terms of maximizing an expected utility function based on Shannon information [99, 100]. More recently Ryan et al. presented an Markov Chain Monte Carlo sampling scheme to determine expected information gain in experimental design decisions [101]. In each case the designer chooses experiments before any data is collected and executes the strategy that maximizes the expected reduction in entropy.

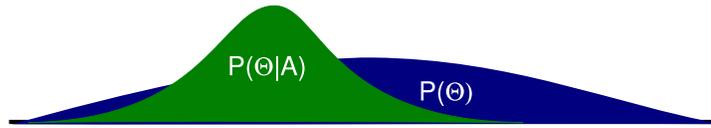


Figure 4.7: Information-theoretic value functions score each observation in proportion to the reduction in uncertainty over values of  $\Theta$ . Here our prior knowledge of  $\Theta$  is imprecise so the prior distribution  $P(\Theta)$  is broad. An observation  $A$  results in a more constrained distribution, reflected by the reduction in posterior entropy  $H(\Theta|A)$ .

Information gain was used in the Robotic Antarctic Meteorite Search (RAMS) to improve rock classification decisions [102, 26]. Here an onboard science planner makes decisions about which instrument to deploy on the basis of information gain with respect to the classification decision. By predicting the time cost of data collection with both spectroscopy and the magnetometer, and the information gain associated with each, the robot can make principled decisions about tradeoffs between classification accuracy and time resources spent on sensor deployment. The RAMS study is significant in that it combined multiple sensor sources to classify detected features. In addition, it integrated feature detection, classification, and adaptive behaviors to perform a significant autonomous science task.

Information gain is a popular principle for action selection in spatial mapping domains. Single- and multi-robot systems have used information gain as an objective for choosing paths to map occupancy grids [103, 104, 105]. These *frontier-driven* exploration strategies do not learn the distribution over future observations; instead, they assign a common information value to each unobserved grid square. This encourages robots to plan paths for maximum coverage of the environment (Figure 4.8).

Moorehead et al. expand the frontier exploration idea to outdoor exploration and planetary survey domains. They incorporate a model of sensor uncertainty where occupancy readings of distant cells are more noisy than nearby readings [60, 106]. In this work a planner seeks to collect information about multiple weighted maps that describe features such as cells' traversability or height. It uses a utility measure weighting data collection cost against the expected information gain.

Stachais et al expand the notion of information-driven mapping in occupancy grids to the full Simultaneous Localization and Mapping problem [107]. This results in mapping behavior that appropriately balances reductions in observation and localization uncertainty. It bears repeating that these exploration strategies differ from Loredó's adaptive learning [108]; they do not model the structure of the environment or predict future observations. Instead, they assign a constant information value to each unobserved location. For this model, the information gain metric reduces to a measure of sensor coverage.

Shewry and Wynn show that under some weak assumptions one can instead maximize the entropy of the observations themselves [109]. It constitutes a classic strategy known as Maximum Entropy Sampling (MES) studied extensively for both linear and nonlinear models [80, 110, 111,

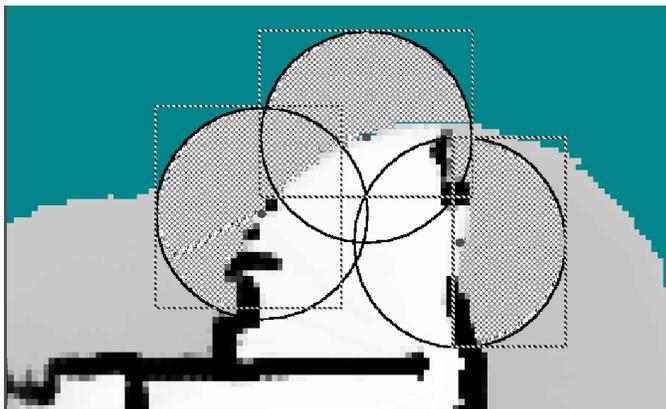


Figure 4.8: In frontier-driven exploration, the information gain of an observation is proportional to the unobserved area that will be revealed. This diagram compares three sensor readings (circles). White areas are known to be clear, black areas are occupied, and grey areas are unobserved. The information gain from each observation (represented by crosshatching) is the area of intersection of the sensor observation window with the unobserved areas. Image courtesy Simmons et al. [103]

112]. MES requires several assumptions about the experimental design problem.

- The system's entropy is bounded; we can eventually learn about the whole environment to within some required degree of precision by collecting many observations.
- The conditional entropy  $H(\Theta|A)$  is functionally independent of experimental design decisions. Observing the environment doesn't significantly change it or alter the character of future observations.
- Observation noise is constant throughout the explored environment.

These assumptions are valid in most practical cases and have been extended in practice to continuous-variable domains. They permit us to think of  $H(A, \Theta)$  as a fixed quantity and ignore this constant for the purposes of observation selection. The general experimental design problem decomposes:

$$R_{\text{MAXENT}}(A) = -H(\Theta|A) \quad (4.14)$$

$$= -H(\Theta, A) + H(A) \quad (4.15)$$

$$\equiv H(A) \quad (4.16)$$

Therefore minimizing  $E_A[H(\Theta|A)]$  is equivalent to maximizing  $H(A)$  [109]. The optimal experimental design maximizes the entropy of the set  $A$  of observations [112]. Shewry and Wynn's result reflects the intuitive notion that to learn the most about the system we should make observations where our results are least certain.

### 4.2.8 Spatial Design

Spatial design is a specific experimental design problem that aims to choose a subset  $A$  out of possible sampling locations  $V$  in the environment. It is also conducive to an information-theoretic objective function. If we take our variables of interest  $\Theta$  to be the observations at unmeasured locations, so that  $\Theta = V \setminus A$ , this results in decomposition parallel to that of equation 4.16:

$$R_{MAXENT}(A) = -H(V \setminus A \mid A) \quad (4.17)$$

$$= -H(V \setminus A \cup A) + H(A) \quad (4.18)$$

$$\equiv H(A) \quad (4.19)$$

Shewry and Wynn apply this result to spatial design problems with continuous observations where measurements at sample sites are normally-distributed. For  $V \in \mathbb{R}^n$ ,  $V \sim \mathcal{N}(\mu, \Sigma_{V,V})$ , we associate sample locations with dimensions of a multivariate Gaussian and reorder the covariance matrix to gather the observations into a block submatrix  $\Sigma_{A,A}$ .

$$\Sigma_{V,V} = \begin{bmatrix} \Sigma_{V \setminus A, V \setminus A} & \Sigma_{V \setminus A, A} \\ \Sigma_{A, V \setminus A} & \Sigma_{AA} \end{bmatrix} \quad (4.20)$$

Traditionally spatial design represents data collection cost  $L_O(A)$  as a hard constraint on the number of observations. These budgets can never be exceeded under any circumstances: cost is zero until the budget is exceeded at which point the cost of additional data collection becomes infinite. In these cases it is simpler to write the utility function in terms of a constrained optimization with maximum observation budget  $m_O$  and transmission budget  $m_T$ :

$$U(A, A') = R(A, A') \quad \text{for } A' \subseteq A \quad (4.21)$$

$$L_O(A) \leq m_O \quad (4.22)$$

$$L_T(A') \leq m_T \quad (4.23)$$

The decoupled objective function becomes:

$$U(A, A') = E[R_O(A)] + E[R_T(A')] \quad \text{for } A' \subseteq A \quad (4.24)$$

$$|A| \leq m_O \quad (4.25)$$

$$|A'| \leq m_T \quad (4.26)$$

The resulting MES optimization aims to find  $n$  sample locations to maximize the differential entropy of the marginal Gaussian  $\mathcal{N}(\mu, \Sigma_{A,A})$  associated with the sub-covariance-matrix  $\Sigma_{A,A}$ . This quantity is given by the following expression:

$$R_{MAXENT}(A) = \frac{1}{2} \log\{(2\pi e)^n |\Sigma_{A,A}|\} \quad (4.27)$$

This is proportional to a constant plus the submatrix determinant, so we can use this quantity instead. For this reason, the determinant maximization is also known as *d-optimal* experimental design [98]. We have:

$$R_{MAXENT}(A) \equiv |\Sigma_{A,A}| \quad (4.28)$$

These MES strategies have been widely applied to a range of spatial design problems [19] such as those we consider in the field experiments.

Recent work by Guestrin *et al.* investigates information-driven sampling in the context of spatial design and sensor placement for monitoring networks [113]. They introduce an alternative *mutual information* criterion that maximizes information gain with respect to several preselected observation points. The *mutual information*  $I(A; \Theta)$  measures the number of bits of information shared between variables. It is equivalent to the following reward function:

$$R_{\text{MUTINF}}(A) = H(V \setminus A) - H(V \setminus A|A) \quad (4.29)$$

This is similar to an MES approach [109], but with a second term that removes any reward from reduction in entropy at the observed locations themselves. In practice this mitigates MES's inherent bias toward high-entropy boundary locations, and can improve the predictive accuracy of the model over MES [113]. The submodularity of the mutual information function provides good theoretical bounds on optimality for greedy sensor placement algorithms.

For a subset  $A$  of size  $n$  from a set of sample points  $V$  with size  $m$ , where sample locations in  $V$  are jointly-Gaussian, the mutual information is given by the expression:

$$R_{\text{MUTINF}}(A) = H(A) + H(V \setminus A) - H(V \cup V \setminus A) \quad (4.30)$$

$$= H(A) + H(V \setminus A) - H(V) \quad (4.31)$$

$$= \frac{1}{2} \log\{(2\pi e)^n |\Sigma_{A,A}|\} + \quad (4.32)$$

$$\frac{1}{2} \log\{(2\pi e)^{m-n} |\Sigma_{V \setminus A, V \setminus A}|\} - \quad (4.33)$$

$$\frac{1}{2} \log\{(2\pi e)^m |\Sigma_{V,V}|\} \quad (4.34)$$

Optimal subset selection in the spatial sampling problem for both conditional entropy and mutual information objectives is an NP-complete problem [80]. However, good fast solutions are still available for most practical situations. Similar approaches have been developed separately in the communities of spatial statistics [19] and active learning [114].

A simple but effective approach is the greedy algorithm that maximizes the objective function with each new added observation. This results in a fast approximate solution for both MES [109] and mutual information [113] criteria. Indeed, in the latter case the simple greedy approach has provable performance bounds [113]. MacKay uses a greedy sequential algorithm to choose labeled datapoints from a set of training data, and shows how adaptive data selection improves performance in classification and regression tasks [114].

An alternative to the greedy algorithms (and a step up in both complexity and runtime) is an exchange algorithm of the type proposed by Shewry and Wynn [109]. Here we initialize the observation plan with the greedy solution and then refine the solution by testing pairwise swaps between observed and unobserved locations. This increases runtime from a linear to low-polynomial time algorithm. Alternatives involve stochastic exchange procedures like simulated annealing and other

standard combinatorial optimization strategies. More recently, Ko has introduced a branch-and-bound solution algorithm that offers performance guarantees at the cost of considerable additional computation [80].

### 4.2.9 Discussion

We believe information-driven objective functions, and spatial design in particular, offer promising candidates for remote exploration. They are especially well-suited to our geologic site survey task, offering several immediate benefits over heuristic approaches such as target signatures or set preferences.

First, they are *principled*. They draw a formal relationship between action selection, the underlying assumptions about the structure of the environment, and the model that relates observations to parameters of interest. These assumptions are quantitative in form and can be validated against datasets as they are collected. Action selection follows logically from these initial choices in a predictable fashion. Heuristic methods also make assumptions about the data’s structure (for example, periodic sampling also presumes some degree of spatial continuity), but these implicit claims are difficult to validate without recourse to a formal model.

Second, information-driven objectives are *flexible*. They can accommodate a wide range of models, including those that integrate remote sensing and surface measurements at different scales. Information gain is a natural framework for handling discovered trends and redundancies across these sensing modes. With on-line learning the agent can adapt the model as it gathers new data to refine its description of the environment. In this sense, the value functions constitute a methodical approach to introducing adaptive behavior into remote science systems. They are an excellent match for our intelligent mapping methodology.

Except where otherwise noted we employ a greedy selection strategy in the following experiments. The greedy solution produces a natural “best  $n$ ” rank ordering of data products; this ensures the most valuable collections occur first, hedging against execution errors or other disruptions.

Information-driven objective functions hold interesting implications for stationary Gaussian process models. For Gaussian process models with a static, stationary covariance function the information gain of a new sample depends only on *proximity to existing data* [113]. This is apparent visually from the stationary step function of Figure 3.4. The marginal variance increases at the periphery of the data and is smallest at the center of the training data. Thus, observation selection is invariant to the actual values of the observations, and there is *no benefit* to adaptive data selection in the decoupled reward function. The best experimental design inevitably spreads observations to evenly cover the sampling domain [113, 19].

While we will ultimately forgo these stationary models in favor of augmented input spaces and spatially-adaptive sampling, they have an interesting property: it is possible to optimize the joint data collection/return problem directly. Consider a spatial design task of allocating  $n$  measurements to a set of sampling locations that do not change during exploration. For the conditional entropy design objective, all that matters is the entropy of the downlink observations that are received by

scientists. This yields the following value function:

$$U_{\text{MAXENT}}(A, A') = E[R_{\text{MAXENT}}(V \setminus A', A') - L_O(A) - L_T(A') \quad A' \subseteq A \quad (4.35)$$

$$\equiv H(A') - L_O(A) - L_T(A') \quad A' \subseteq A \quad (4.36)$$

$$\equiv |\Sigma_{A'A'}| - L_O(A) - L_T(A') \quad A' \subseteq A \quad (4.37)$$

Assuming  $L_T(A)$  increases monotonically as new data products are added, observed and transmitted datasets can be made equivalent - there is no benefit to making observations beyond those that can be transmitted.

$$U_{\text{MAXENT}}(A, A') \equiv |\Sigma_{A'A'}| - L_O(A') - L_T(A') \quad (4.38)$$

This is a standard experimental design problem of maximizing a size  $n$  submatrix determinant subject to row and column costs. This simplification also works for the mutual information objective.

$$U_{\text{MUTINF}}(A, A') \equiv E[R_{\text{MUTINF}}(V \setminus A', A')] - L_O(A') - L_T(A') \quad (4.39)$$

$$\equiv \frac{|\Sigma_{A'A'}| |\Sigma_{V \setminus A'}|}{|\Sigma_V|} - L_O(A') - L_T(A') \quad (4.40)$$

The static covariance structure implies that the information value of an observation set is independent of the observed value itself. The stationary covariance models we have investigated exemplify this phenomenon; the observations' covariance relationships are the same as long as their positions relative to each other remain constant.

However, field scientists performing site survey seldom sample at regular intervals; instead, they target their exploration to the most informative areas and adapt their own spatial exploration on the fly. This fact suggests that the site survey task requires non-static covariance functions. We can identify several ways in which the covariance function might change so that spatial models benefit from adaptive exploration in practice:

- A Gaussian process covariance function is nonstationary in some data-dependent manner.
- A latent input space is data-dependent.
- The agent re-estimates unknown covariance hyperparameters using collected data.

In general, exploration rovers can improve their performance by modifying the experimental design plan on the fly. This idea was presented by Loredó et al. who expand the idea of Bayesian experimental design to permit this adaptive modification. Information gain functions as a criteria for choosing each new measurement in a series of experiments [108, 115, 116].

At each stage in the sequence, previously collected data produce predictions about future observations (Figure 4.9). The experimenters chose the observation that provides maximum information gain over model parameters given data already collected. Adapting one's experimental strategy with each new measurement is an important innovation; it can provide far more information per measurement than traditional hypothesis testing paradigms with a fixed experimental plan. This *sequential design* strategy has much in common with active learning approaches from economics and statistics [114, 117].

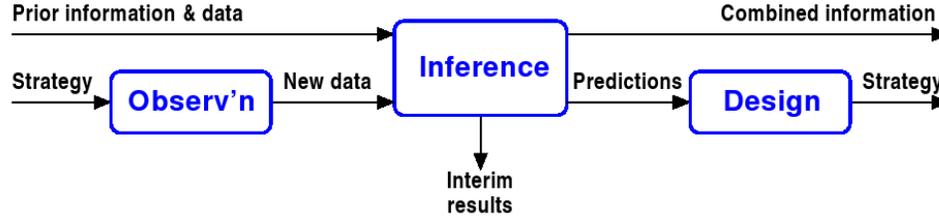


Figure 4.9: Loredó’s sequential design method exploits the learned structure of the process under investigation. Each new observation is incorporated into a generative model of the hidden process. By extrapolating to the future behavior of this process, one can improve predictions about the information gains that will result from new observations. Image courtesy Loredó et al. [115]

Krause investigates non-myopic strategies for adaptive exploration that consider the effect that the observations will have on the model parameters themselves [118]. His work examines an exploration/exploitation tradeoff that balances gathering information about model parameters with simultaneously refining the existing model. Krause concludes that one can bound the possible improvements in prediction accuracy from adaptive sampling using the entropy of the hyperparameters. In other words, high *a priori* certainty about covariance function parameters mitigates the advantages of adaptive observation selection. But as long as the agent can modify the covariance function on the fly, it can potentially benefit from collecting extra observations beyond those that can be transmitted.

Other models besides Gaussian processes can be more difficult to characterize. In the more general case the conditional entropy of untransmitted data may depend on the values of the transmitted observations. This is common, for example, in the case of the graphical models that we will examine in Section 5. If  $R(A)$  is data-dependent, the simplifications of equations 4.38 and 4.40 are not possible. These situations again call for a decoupled approach in which experimental design guides active data collection without considering the downlink.

To summarize, our data collection strategy reestimates Gaussian process hyperparameters using the entire set of collected data before each replanning opportunity. Then we choose observations  $A$  that satisfy the remaining resource budget while maximizing  $R_{\text{MUTINF}}(A)$  or  $R_{\text{MAXENT}}(A)$ . Selective data return employs a separate optimization routine, an issue that Section 5 will address at greater length.

### 4.3 Adaptive Sampling from a Transect

This section investigates adaptive sampling using information-driven utility functions with different Gaussian process models and environments. We will compare performance across several axes of variation:

- Sequential vs. static design: Do we reevaluate the observation set after every new observa-

tion, or always follow the original plan?

- Experimental design objective: How does MES sampling differ from that induced by the mutual information criterion?
- Model: Is data best described by stationary or spatially-varying covariance relationships? How do latent inputs affect sampling behavior?

We aim to explore the effect of these different choices on sampling and reconstruction on data collection.

We will focus on the spatial sampling problem where a set  $V$  of potential new measurement locations has Gaussian-distributed prediction values with a covariance matrix  $\Sigma$ . For MES, the best sampling strategy maximizes the joint entropy of the future observations; we must choose  $m_c$  collection points  $A \subseteq V$ ,  $|A| \leq m_c$  that maximize the entropy of the associated multivariate Gaussian.

In the more general case the agent must balance the time cost of data collection against that of moving to a new physical location, a topic we will address later. Here we will isolate the experimental design issues from those of path planning. We will represent resource constraints by an artificial limit on the number of observations, and treat data collection costs as fixed throughout the input space. This simplification renders the problem analogous to a classical spatial design [19] except that the location of observations is determined sequentially. In other words, the agent can jump to a new location after each observation with no penalty; measurements are “expensive” but physical motion is “cheap.”

Our synthetic environment consists of a noisy step function as in Figure 3.4. For each trial there is a single step discontinuity located at a random position in the input space. We seed the algorithms with four initial measurements at one extremum. This provides enough data to prevent numerical instabilities while revealing as little of the unexplored environment as possible. We characterize sampling performance by reconstructing the step function based on each method’s selected datapoints.

After each new observation the agent retrains its onboard model and chooses a new observation site from any unobserved location along the traverse. This corresponds roughly to MacKay’s active learning scenario [114] or Loredó’s adaptive experiment design [115] in which the agent has continuing access to the complete set of unobserved datapoints.

The algorithm for these experiments greedily observes the best new observation at each time step. It begins with a current set of observations  $X, Y$  at sites  $A$ , and a set of discrete candidate observation locations  $V \setminus A$  that are sampled from the continuous input space. It evaluates the expected reward  $R(A \cup a_i)$  of each new potential datapoint based on the trained model and the known input attributes at the sample site (See the algorithm of Figure 4.10).

Figure 4.11 shows the sampling behavior for a squared-exponential covariance function utilizing MES and mutual information criteria and relearning the hyperparameters after each observation. Here the red triangles show the selected observation locations and each digit corresponds to

```

Input: potential sample sites  $V$ 
         independent attributes  $y$ 
         budget  $m_c$ 
Result: observed sites  $A$ 
         independent attributes  $X$ 
         dependent attributes  $Y$ 
seed  $A, Y$  with initial observations;
while  $|A| < m_c$  do
   $a_{\text{new}} = \max_{a_i \in V \setminus A} E[R(A \cup a_i)];$ 
  observe  $a_{\text{new}}$  consisting of  $(x_{\text{new}}, y_{\text{new}});$ 
   $A \leftarrow A \cup a_{\text{new}}, \quad X \leftarrow X \cup x_{\text{new}}, \quad Y \leftarrow Y \cup y_{\text{new}};$ 
  relearn covariance function using  $X, Y;$ 

```

Figure 4.10: Greedy sequential algorithm for datapoint selection

the (greedy) selection order. Dotted lines show the prediction variance for the underlying noiseless  $f(x)$ . These variance values are normalized within each image to emphasize the differences between locations in the input space and highlight the uncertain areas according to each model.

The sampling decisions evidence several distinct behaviors. Both aim to fill the entire available space, choosing observations that are well-spaced along the input dimensions. However, the maximum entropy criterion exhibits a slight bias towards the extrema where prediction variance is high. This is apparent from the rightmost image, where the first observation consists of the datapoint farthest from any previous data. The second choice is an observation between the previous observations, favoring the right edge where the data is still sparser.

The image at right shows greedy sampling with a stationary model using the mutual information criterion. As in Guestrin *et al.* we notice this does not exhibit this affinity for boundaries [113]. However, note that both stationary models are indifferent to the step discontinuity.

Augmenting the model with latent inputs (as described in Chapter 3) produces pronounced qualitative differences in the sampling behavior. The latent input of Figure 4.12 stretches the input space around the step discontinuity, creating an area of sparse data — an artificial periphery that draws additional samples. This reflects the high rate of change near the step. The observation order approximates a binary search for the boundary, still favoring the rightmost extremum as before. Because it pursues areas of high prediction variance, the overall variance for the reconstruction is low.

Interestingly, the latent input has the opposite effect when used with the mutual information criterion (Figure 4.12). Because data is sparse in the regions of high change, there is less predictive overlap with the remainder of the data and mutual information selection actively avoids these areas. This has interesting implications for exploration in models with latent inputs; while both criteria aim to fill the available space (and do so more effectively than sequential sampling), the mutual

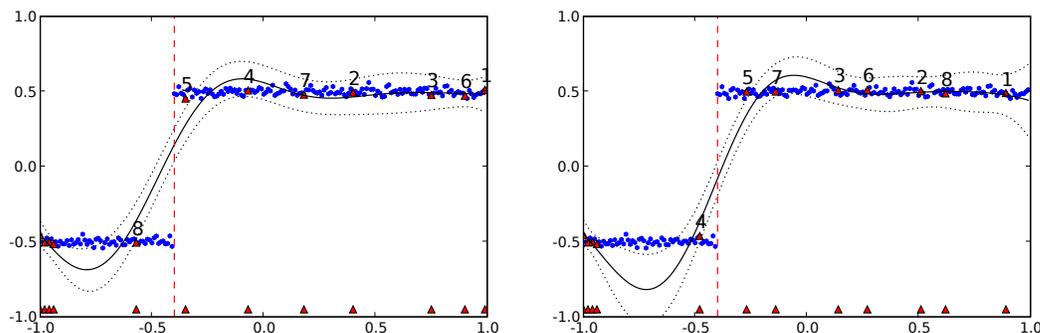


Figure 4.11: Red triangles indicate the location and of observations and digits the selection order. Left: A stationary model with active sampling according to the maximum entropy criterion. Selection favors an even spread of points throughout the input space, with a slight bias toward the boundary where entropy is high. Right: A stationary model with sampling according to the mutual information criterion, which favors points that are centrally-located with respect to the entire set of potential observations. Both stationary models are indifferent to the step discontinuity. Dotted lines show prediction standard deviation at each location for the stationary model.

information standard encourages “border avoidance” rather than “border affinity.”

The preferred choice will probably depend on science objectives: MES favors novel data products drawn from sparsely-sampled areas, while mutual information provides better reconstruction accuracy over the potential sites  $V$ . If the distribution of potential sites  $V$  parallels the science value of each location in the input space, then mutual information may be preferable. However, MES may be superior in situations with dynamic or augmented input spaces. The arrangement of sample locations in the augmented input space could change in response to the agent’s observations, and the density of locations in  $V$  might not reflect their value to scientists.

These examples also indicate that the covariance structure of stationary and nonstationary models has significant influence on sampling behavior. In the stationary case prediction entropy depends only on the prediction’s proximity to other data, so as long as covariance hyperparameters remain fixed adaptive data collection is not helpful to refine the model. With latent inputs and covariance re-learning, prediction entropy is data-dependent so adaptive sampling can be used to improve the utility of future observations. The model parameters determine the appropriate trade-off between boundary-following (for a nonstationary process), spatial coverage, and representative sampling along any additional input attributes.

We evaluate the sampling behaviors quantitatively by reconstructing the complete step function from each transmitted data set. Our performance score is the mean squared prediction error across the sample space. To ensure fairness we used the same model for all reconstructions: a Gaussian process using  $\hat{f}(x)$  a latent input. Table 4.1 and Figure 4.13 shows performance scores for active sampling for increasing numbers of observations. For this environment and number of samples there is no significant difference between the reconstruction error of the entropy-based and mutual

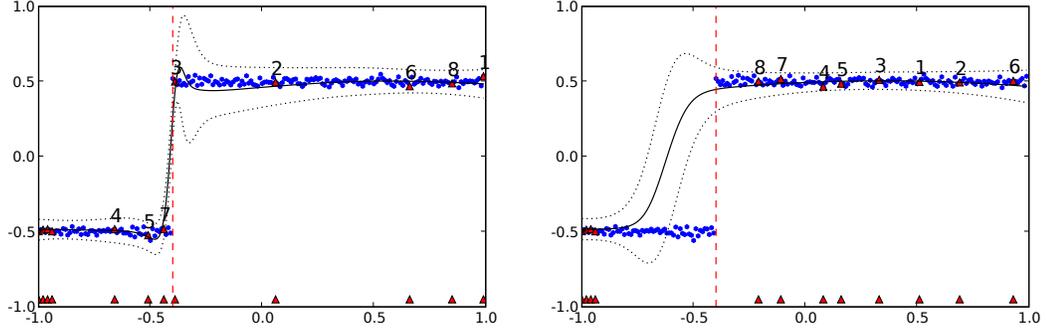


Figure 4.12: Left: A latent-input model, where Maximum Entropy Sampling adjusts to the step discontinuity. Right: The mutual information criterion actively avoids the high-change areas. In the augmented input space the data is sparser (and their distributions less correlated) in these areas; the observations provide less mutual information.

Method	accessible	constrained
Stationary, MaxEnt	$0.028 \pm 0.001$	$0.031 \pm 0.001$
Stationary, MutInf	$0.028 \pm 0.001$	$0.029 \pm 0.001$
LatentGP, MaxEnt	$0.021 \pm 0.002$	$0.031 \pm 0.001$
LatentGP, MutInf	$0.026 \pm 0.001$	$0.031 \pm 0.001$

Table 4.1: Mean absolute error in reconstruction with active sampling.

information criteria, with one significant exception: the latent input model coupled with MES scores significantly better than the alternatives. This reflects the qualitative behavior observed above in which MES quickly narrows down on the correct border area. The result is a superior overall estimate and a faster convergence.

We also consider a more constrained scenario analogous to a transect in which the agent cannot move backwards. Here, our algorithm allocates an entire sequence of observations to those locations after its most recent observation, and observes the closest site. In the cases with latent inputs we re-learn the covariance hyperparameters and compute a new observation plan for the remaining budget. The result, appearing in the third column of Table 4.1, is near-equivalent performance among all the different sampling strategies. This occurs because the forward motion constraint prevents the latent-space models from any further action to localize the boundary once it has been discovered. In the next section, we will consider forward motion in a 2-dimensional environment; there it becomes possible to re-cross any boundaries that are not to the rover’s forward progress.

It is also instructive to evaluate active sampling in the context of the Amboy Crater traverse data that we introduce in chapter 3. Here our goal is to characterize any improvement in the reconstruction gleaned by active sampling regimes using remote sensing data. The Gaussian process inputs take the form of the data products’ temporal position, and optionally, the associated DOQ remote

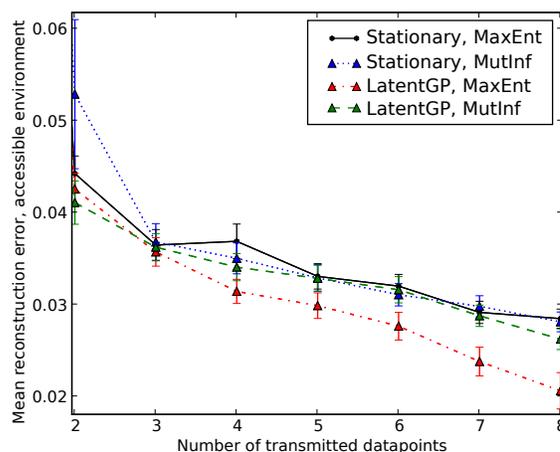


Figure 4.13: Mean squared error performance for 100 trials of active and static sampling. Error bars show 95% confidence intervals.

sensing image. In both cases the observation selection task is the same as before: to produce an accurate reconstruction of spectroscopy measurements at all of the unsampled locations.

We utilize the sequential selection algorithm (Algorithm 4.10), initializing the Gaussian process with three observations at the beginning of the traverse. Maximum entropy sampling chooses each new observation in the sequence, and we re-learn covariance function parameters at each step. Finally, we compute the reconstruction error by retraining a Gaussian process using the set of revealed observations, their positions, and the complete remote sensing data, which is presumably available to scientists in any case. The reconstruction score is the mean absolute difference between predicted and actual spectroscopy measurements.

Figure 4.14 shows the prediction accuracy for various numbers of observations using both position information and position information coupled with remote sensing data. The charts show error for each of the four traverses. In each case, the addition of remote sensing data dramatically improves the reconstruction error. Figure 4.15 shows the mean improvement in reconstruction facilitated by adding remote sensing data as an additional input to the active selection decision. The benefit is significant for every observation budget, but the greatest discrepancy occurs when the budget permits only a small number of observations. This reflects the active sampling leveraging orbital data to quickly converge on a good reconstruction.

## 4.4 The Corridor Exploration Task

This section presents a new operational mode for autonomous traverse that leverages orbital and surface measurements to efficiently map a prespecified corridor under time and position constraints. We present the basic task and describe path-planning algorithms for MES exploration. Finally, we

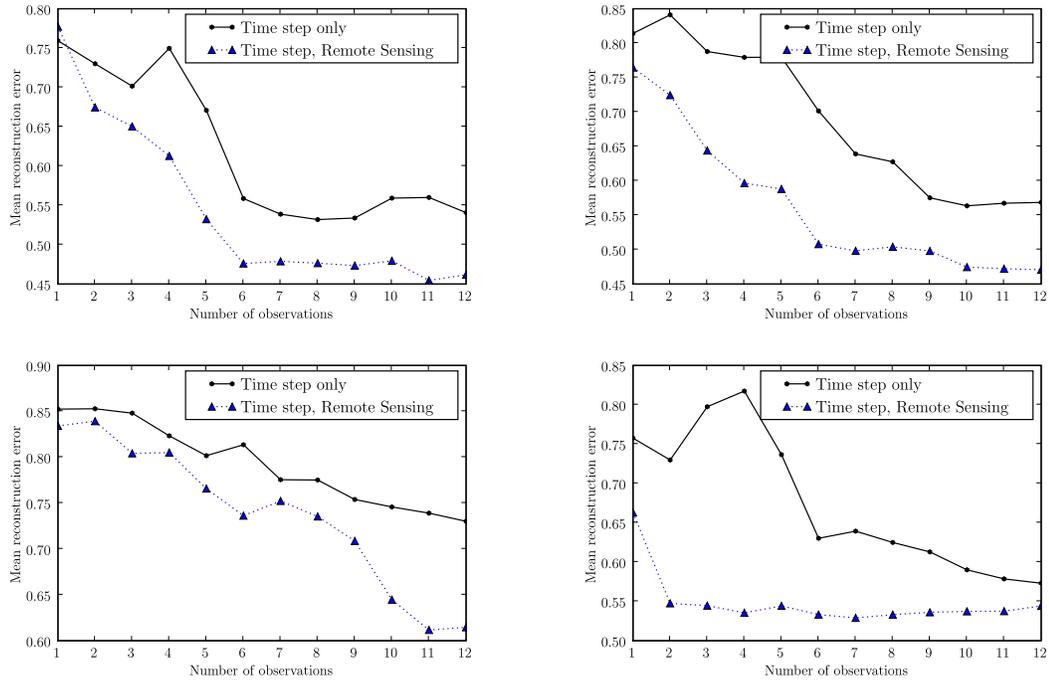


Figure 4.14: Reconstruction of field traverses from Camp C using both time sequence and remote sensing data. Plots show the ability of the active sampling strategy to reconstruct spectrometer measurements at unobserved locations. The graph shows reconstruction error for various observation budgets. Upper Left: Traverse 1. Upper Right: Traverse 2. Lower Left: Traverse 3. Lower Right: Traverse 4. Including remote sensing information in the observation selection decision dramatically improves performance, even though the reconstruction has complete access to remote sensing data in all cases. Maximum-likelihood hyperparameters are reestimated for each set of observations before choosing the next sampling location; since these observations are noisy the ground-truth reconstruction error does not decrease monotonically.

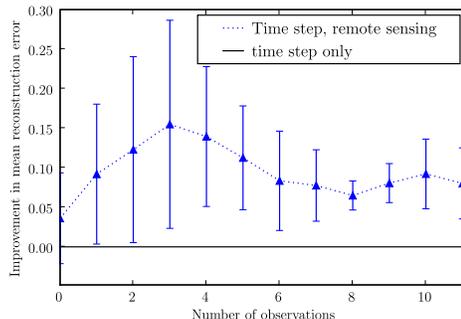


Figure 4.15: Active sampling regimes including latent remote sensing data outperform those that rely on position information. Here we provide examples of mean reconstruction error from four traverses along with 95% confidence intervals. Equivalent reconstruction methods were used for both sets of observations.

demonstrate experiments at Amboy crater in which a field rover generates automatic geologic maps of surface material in real time.

#### 4.4.1 Problem Formulation

The corridor exploration task is a 2-dimensional analog of the spectroscopy transects we considered in chapter 2. As before, operators specify a goal location and a time budget. In addition they provide a horizontal offset, defining a corridor within which the robot can travel freely. The agent moves forward inside this predefined exploration area toward the end-of-day location, constructing a surficial map based on observations made at regular time intervals. Figure 4.16 illustrates the scenario: a rover travels the path associated with maximum observation entropy that satisfies the available resource budget.

This corridor exploration task simplifies the general problem in several convenient ways. First, the agent incurs no explicit cost for data collection; instead, physical navigation is the only significant resource consumer. This is consistent with instruments like cameras or short-integration reflectance spectra that have trivial collection cost, or with any data collection that occurs automatically at a fixed rate. Incorporating data collection cost into movement cost reduces an unconstrained optimization problem to one of path planning.

The task also imposes movement constraints: the rover must move only forward and remain within the safe area defined by scientists at the outset. Our formulation facilitates integrating corridor exploration in modular fashion into mixed mission plans containing pre-scripted sampling and navigation waypoints. For example, scientists wishing to perform a long distance over-the-horizon traverse to a new site could use corridor exploration to permit limited science autonomy while conservatively subordinating the mapping task to mission navigation goals. Nevertheless, it provides sufficient autonomy to permit meaningful adaptive exploration; by choosing the appropriate lateral

offsets the rover can decide which locations within the corridor it must visit to produce the most faithful map reconstruction.

Formulation also ignores issues of navigability. Instead, we will focus on large-scale paths composed of waypoints. Our implementation architecture passes these waypoints to a lower-level local navigator that uses whatever terrain analysis and obstacle avoidance is necessary to reach these waypoints. This provides a clean separation between mobility-based navigation decisions, which require high-resolution plans, fast updates, and local information, and the more computationally-expensive science planning, which deals with longer plans and larger spatial scales.

Despite some significant constraints, corridor exploration is well-suited for testing the basic principles of autonomous site survey. It provides a meaningful test case for extrapolating predictions because the rover can recross any boundaries that it discovers that do not run perpendicular to the corridor direction. Corridor exploration involves significant resource constraints yet remains tractable for real-time implementations.

We represent each path  $\mathcal{P}$  using interpolated waypoints  $\mathcal{P} = \{\mathcal{P}_1, \mathcal{P}_2, \dots, \mathcal{P}_{\text{goal}}\}$  that extends to the goal. The entropy  $H(\mathcal{P})$  associated with this path is the joint entropy of  $n$  observation sites spaced equally along the interpolated path, conditioned on all existing observations. This is only an approximation since the lower-level navigation routine might not track the interpolated path exactly, and the turning radius of the vehicle is finite in any case. The number of sampling locations  $n$  controls the fidelity of our pointwise approximation of the total entropy of observations that will be collected by following the idealized path. It should be large enough to permit accurate approximations without imposing onerous computational burdens - the experiments that follow use  $n = 40$  for total (discretionary) path lengths that can grow as large as 200m.

Recall that general utility functions may have multiple terms balancing expected information gain against the cost of time and other resources [60]. Here movement cost is implicit in the requirement that the robot be able to reach the end-of-day goal location. This is a discontinuous utility function: cost is zero until the budget is exhausted, but failing to reach the goal incurs infinite cost. The path planning algorithm need only explore the space of paths capable of reaching the end-of-day goal.

We represent this as a constrained optimization with path cost  $L_O(\mathcal{P})$  equal to the travel time, the Euclidean path length multiplied by a conservative estimate of the average rover velocity. The maximum path budget  $m_O$  is prespecified by scientists at the start of the traverse. For a data set  $A$  of existing observations, the utility function of a forward path  $\mathcal{P}$  is given by the following expression based on the MES criterion:

$$U(\mathcal{P}) = H(\mathcal{P}|A) \text{ S.T. } L_O(\mathcal{P}) \leq m_O \quad (4.41)$$

Our utility function uses the constraint as a computational tool but it cannot ultimately guarantee that the rover will reach the goal. In particular it disregards execution uncertainty: if the robot is delayed unexpectedly due to unforeseen obstacles, the cost of the current plan could become infinite. We will account for this by replanning regularly during the traverse so that the current plan being executed reflects the actual time resources remaining.

A conservative estimate of the average rover velocity compensates for the inevitable delays

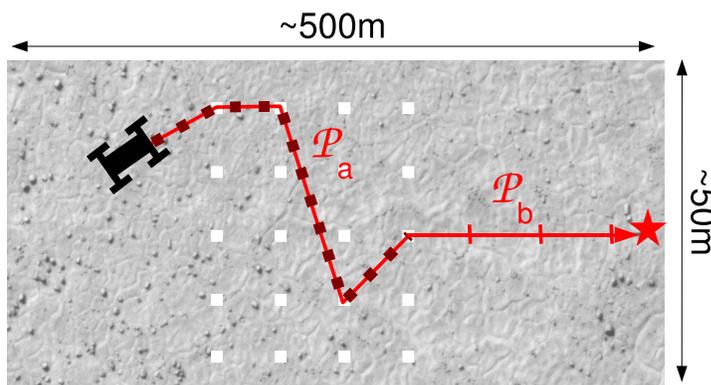


Figure 4.16: The corridor exploration problem. The rover travels forward within a predefined rectangular region toward an end-of-day location (represented here by a red star). The rover chooses an informative path that satisfies prescribed time constraints. Our solution plans waypoints within a local time horizon by choosing lateral offsets for each of several forward positions. We compare candidate paths by simulating sequences of observations. The remainder of the traverse  $\mathcal{P}_b$  over the local planning horizon is a straight drive towards the goal.

introduced by spectrometer calibration activity, pauses for replanning, and obstacle avoidance maneuvers initiated by low-level navigation routines. We use an expected velocity of  $0.33m/s$  for a commanded vehicle forward velocity of  $1.0m/s$ . We chose this average velocity estimate through trial and error to yield appropriate time estimates for the test environments of these field experiments.

The corridor applies restrictions that make it unnecessary to plan the entire path to the goal explicitly at each time step. Instead we split the complete problem into two segments - a near-term discretionary portion  $\mathcal{P}_a$  and a fixed remainder  $\mathcal{P}_b$  that each receive a proportional amount of the total time budget.

The local path planner adjusts the near-term path  $\mathcal{P}_a$  within the fixed-distance horizon. Execution begins with the time budget allocated evenly to segments in  $\mathcal{P}_b$ ; these segments and their budgets are added to the collective discretionary portion as they fall into the rover's replanning horizon. This ensures that the rover does not linger too long in any one location, permitting short-term adaptive resource allocation while enforcing an even allocation of time resources in the long term. Altering the length of the planning horizon controls the tradeoff between planning complexity and long term optimality. The local horizon path planning utility function is the following:

$$U(\mathcal{P}) = R_{\text{MAXENT}}(\mathcal{P}_a) \text{ S.T. } L_O(\mathcal{P}_a) \leq (m_O - L_O(\mathcal{P}_b)) \quad (4.42)$$

We will omit the subscripts for clarity in our discussion of the planning algorithm; all the relevant planning aims to find a path  $\mathcal{P}_a$  for a time budget  $m_O - L_O(\mathcal{P}_b)$ .

The subpath time cost is a function of the average rover velocity  $\nu_{\text{avg}}$  and the Euclidean subpath lengths  $l(\mathcal{P}_i)$  associated with each leg  $\mathcal{P}_i$  of a path segment.

$$L_O(\mathcal{P}) = \nu_{\text{avg}} \sum_{\mathcal{P}_i \in \mathcal{P}} l(\mathcal{P}_i) \quad (4.43)$$

The path selection criteria attempts to maximize the joint entropy of observations in the path conditioned on all previous observations. We derive the path reward from the set  $A^*$  of evenly-spaced sample locations along a path  $\mathcal{P}$ . We compute the entropy of the multivariate Gaussian conditioned on existing collected observations. For a current set of observations  $A = \{a_1, \dots, a_n\}$  associated with input variables  $X = \{x_1, \dots, x_n\}$ , the reward function is given by substituting the conditional covariance matrix into Equation 3.15:

$$R_{\text{MAXENT}}(\mathcal{P}) = H(\mathcal{P}|A) \quad (4.44)$$

$$\equiv |\Sigma_{A^*, A^*}| \quad (4.45)$$

$$\equiv |K_{A^*, A^*} + K_{A^*, A}(K_{A, A} + \sigma^2 I)^{-1} K_{A, A^*}| \quad (4.46)$$

The path planner generates a near-term plan  $\mathcal{P}_a$  for the discretionary budget by considering lateral offsets for each forward position along the corridor (Figure 4.16). We search the space of lateral offsets using a recursive greedy path planning algorithm first formulated by Chekuri and Pal [119] and elaborated by Singh et al. [120]. The algorithm evaluates each legal midpoint, breaking the path planning problem into two smaller segments. For each midpoint it evaluates several potential assignment ratios of the remaining resource budget among the halves. It calls itself recursively on the left subpath and then the right, finding the high-scoring path for the left half and then taking that portion as fixed for the purposes of finding the right half — hence, the algorithm is “Greedy” (See the algorithm of Figure 4.17). Its branching factor is equal to the number of waypoints considered multiplied by the number of potential budget splits.

For submodular functions (such as the entropy objective we consider here) Chekuri and Pal prove a  $\mathcal{O}(\log \text{OPT})$  performance guarantee for general directed graphs [119]. In other words, for general walks on directed graphs, greedy selection ensures a reward value of at least  $\mathcal{O}(\text{OPT} / \log |\mathcal{P}_{\text{opt}}|)$  where  $\text{OPT}$  is the optimal reward and  $|\mathcal{P}_{\text{opt}}|$  the number of waypoints in the associated path. Increasing the number of waypoints yields only logarithmic detriment to the approximation.

Our algorithm selects lateral waypoints from five possibilities. For each waypoint we evaluate 3 potential budget splits ascribing 25%, 50%, and 75% of the budget to the first segment. Candidate subpaths must never exceed the budget allocation for their segment; if there is no valid option the planner resorts to a default path driving directly to that segment’s endpoint. Thus, if the robot overspends its discretionary budget (due to navigation difficulties) it drives directly to the end of its discretionary exploration area as quickly as possible. During the next replanning, the discretionary portions are reassigned and it receives some of the reserve budget originally assigned to the nondiscretionary portion of the plan. It then resumes planning within the bounds of the new, larger discretionary budget.

The result of this approach is steady progress down the corridor since the planner can never commit more time than the current discretionary portion to local operations within the local planning horizon. Nevertheless, one could imagine science investigations in which radically uneven

```

Objective:  $\max_{\mathcal{P}_a} R_{\text{MAXENT}}(\mathcal{P}_a)$ 
Given: existing observations  $A$ 
Algorithm: GreedySelect(start, end, budget,  $\mathcal{P}_{\text{fixed}}$ )
if  $L_O(\text{start}, \text{end}) > \text{budget}$  then
    return infeasible;
if base case: no intermediate waypoints then
     $\mathcal{P} \leftarrow \{\text{start}, \text{end}\};$ 
else
    for  $b$  in splits do
        for middle in lateral offsets do
             $\mathcal{P}_1 \leftarrow \text{GreedySelect}(\text{start}, \text{middle}, b, \mathcal{P}_{\text{fixed}});$ 
             $\mathcal{P}_2 \leftarrow \text{GreedySelect}(\text{middle}, \text{end}, \text{budget} - b, \mathcal{P}_{\text{fixed}} \cup \mathcal{P}_1);$ 
            if  $H(\mathcal{P}_{\text{fixed}} \cup \mathcal{P}_1 \cup \mathcal{P}_2 | A) > \text{best}$  then
                 $\text{best} \leftarrow H(\mathcal{P}_{\text{fixed}} \cup \mathcal{P}_1 \cup \mathcal{P}_2 | A);$ 
                 $\mathcal{P} \leftarrow \mathcal{P}_1 \cup \mathcal{P}_2;$ 
return  $\mathcal{P};$ 

```

Figure 4.17: Recursive greedy algorithm for selecting lateral offsets.

time expenditures were appropriate. For example, a single extremely rare science feature (such as the Meridiani meteorite [121]) might warrant extensive measurement with multiple instruments. But in our site survey task, information gain will drive the agent to improve spatial coverage in any case; the basic objective of fast long-range survey is congruent with a “keep moving” approach.

The short planning horizon permits fast real-time replanning; for a modern single-core laptop processor running unoptimized code the entire replanning procedure requires anywhere from 5 – 30 seconds. Planning is fastest in situations where the agent has very little time remaining in its discretionary budget; this allows it to dramatically reduce the breadth of the search by pruning non-conforming paths. Of course, if it is behind schedule there is no time remaining in the discretionary portion and it need not perform any search at all. It simply drives forward from one middle waypoint to the next until it passes close enough to the non-discretionary segments that it begins to receive some of their reserved time allowance. The relevant portion is then assigned to the discretionary planning budget during the next replanning opportunity.

We investigated alternative strategies that planned more-or-less regularly, over longer distances, and with more lateral waypoint options. Ultimately, we settled on parameter settings that provided qualitatively reasonable behavior while maintaining a high ratio of movement to planning time. In the physical experiments that follow the rover replans its path at 3 minute intervals, evaluating discretionary paths that each simulated 40 future observations. Thus, the rover might spend anywhere from 2 – 5 minutes planning over the course of a 30 minute run. This was a practical compromise between path optimality and the time spent planning; it was necessary that the time spent planning “pay for itself” with better paths and superior reconstruction error. Our parameter settings provide

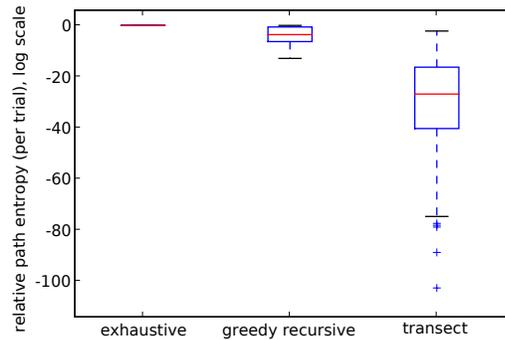


Figure 4.18: Comparison of path entropies for 70 planning trials using exhaustive and greedy path planning methods. The plotted values show the path entropy relative to the optimal (exhaustive) strategy. The control group (in the rightmost column) uses a fixed path running directly from the start to the end of the corridor. All paths incorporate an identical number of observations spaced at regular intervals along the traverse.

a point example of a reasonable alternative to periodic sampling but we do not claim that they are globally optimal or that other planning allowances might also yield improvements over fixed-path sampling strategies.

We evaluate the greedy path planning’s performance by comparing the entropy of the resulting paths to those produced by exhaustive and fixed planning strategies. We perform 70 planning trials using data from the Amboy Crater traverses. As in the “lookahead prediction” experiments of Chapter 3, each trial corresponds to a specific time step of the traverse. At each rover location we construct a Gaussian process model based on remote sensing data and previously collected data. We then generate a 4-waypoint path forward down the corridor.

We enforce a time constraint that precludes the longest paths, preventing both exhaustive and greedy methods from simply choosing the obvious high-entropy paths that jump from one side of the corridor to the other. Figure 4.18 shows the entropy of the resulting paths compared to an exhaustive search strategy and a nonadaptive fixed transect path that runs directly from the start to the end of the corridor. Unsurprisingly the greedy algorithm sacrifices some optimality but performs significantly better than the static alternative.

The exhaustive search is optimal with respect to a fixed Gaussian process model, but an even more general solution might consider the potential of future observations to change the Gaussian process model itself. This is a Partially-Observable Markov Decision Process (POMDP). The general POMDP exploration task is intractable, but recent work with off-line optimization has made progress using approximate solutions and cleverly-factorized state spaces [122, 123]. In this sense our algorithm is strictly myopic; it ignores POMDP aspects of the problem and presumes that model parameters are correct at each time step. However, at regular time intervals our agent re-estimates model parameters using all collected data and creates a new MES observation plan

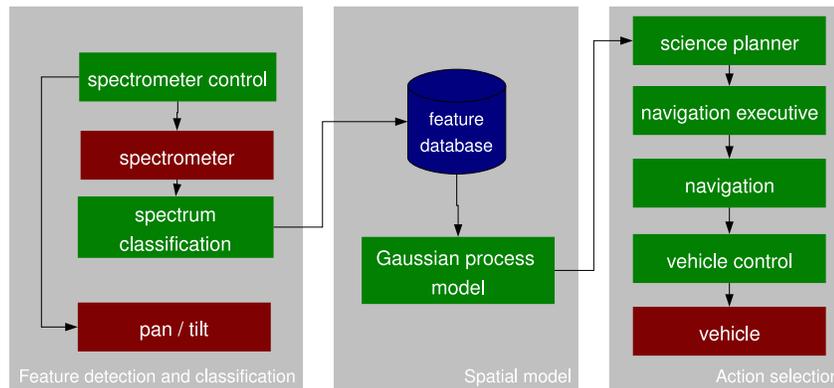


Figure 4.19: Software architecture for corridor exploration.

for the remaining resource budget. This is tantamount to Loredó’s Bayesian adaptive exploration approach [115].

#### 4.4.2 Software Architecture and Execution Strategy

The corridor exploration software architecture is similar to that of chapter 2.3.1. A schematic diagram of the software modules appears in Figure 4.19. We recognize the three basic system components of feature detection and classification, spatial modeling, and action selection. In this case there is no explicit feature detection — each spectral measurement constitutes an independent feature. We classify spectra as in chapter 3 using five-band averaging followed by Principal Component Analysis.

A latent-input Gaussian process model maps the latitude and longitude location of each sample, along with any available remote sensing data, onto the scalar geologic class prediction at future observation points. This informs the “science planner” module that uses it at each stage of the planning process to evaluate the entropy of candidate paths. The science planner and related modules correspond roughly to the action selection portion of the architecture.

Every three minutes of travel time the science planner pauses to re-plan the rover path. This generally corresponds to 50 – 100m of forward travel, subject to delays from spectrometer calibration and obstacle avoidance. We recompute the principal component axis for evaluating spectroscopy, and re-train covariance function hyperparameters of the Gaussian process on all the available data. We re-assign the discretionary portion of the budget to take account of elapsed time. The science planner chooses a path with the recursive Chekuri and Pal algorithm, and sends the resulting waypoints to a navigation executive.

The navigation executive breaks the complete trajectory to the next waypoint into lower-level subgoals spaced no closer than 30m apart. It sends these subgoals to a local navigator that uses  $D^*$  path planning to produce drive arcs toward the goal. Note that the rover does not necessarily

track the original straight-line path between locations as simulated by the science planner, but instead attempts to find the closest path to the next waypoint from the rover's current location. The navigation executive also handles interrupt activities such as recovery maneuvers and fault handling.

Finally, the vehicle controller module performs active motor-level control by supplying appropriate drive speeds to move the rover in the direction commanded by the local  $D^*$  path planner. We set the nominal target rover velocity to a fixed value of 1.0 meters per second. In practice the lower-level navigator exercised a velocity-limiting function based on its evaluation of the terrain, and the vehicle traveled more slowly to prevent excessive vibrations. Moreover, the rover occasionally detoured slightly from the intended interpolating path to avoid rough terrain and obstacles. Again, the planner's average velocity estimate  $\nu$  of 0.33 meters per second accounted for delays due to spectrometer calibration, replanning, and execution uncertainty in the low-level navigator.

The science planner serves one other important purpose: it re-calibrates the spectrometer at periodic intervals using the scripted axle-straightening procedure described in section 2.3.1. Every 7 minutes during rover operations, the science planner pauses the vehicle's forward motion by a direct communication to the vehicle controller, and sends a message to the spectrometer controller to initiate a spectrometer calibration action. This overrides the entire navigation procedure, which waits until the spectrometer calibration is finished. When calibration is completed — with a success or a failure — the spectrometer control module broadcasts a status message causing the science planner to release its lock on the navigation systems and resume forward motion.

Occasionally the rover would find itself perched on precarious or rocky terrain during this recalibration procedure, in which case it could not successfully straighten the rover axles. Because this would cause the spectrometer to miss the white reference target, it would abort the calibration procedure and re-try again 60 seconds later. These 60-second attempts would continue until the rover succeeded in collecting a white reference spectrum.

During the traverse, the spectrometer control module attempts to collect spectra at regular  $2m$  distance intervals. However, the rate of data capture and processing limited spectrometer acquisitions to a maximum rate of around one spectrum per 7 seconds. This corresponds roughly to a new spectrum for each  $2 - 5m$  of forward progress. The module schedules these spectral measurements independently by monitoring rover wheel odometry. Since the rover must stop for each calibration, the procedure overrides any spectrometer acquisition activity for its complete duration. This prevents the spectrometer acquisition and calibration routines from attempting to use the pan-tilt unit simultaneously. Using rover motion as an implicit communication channel resolved the pan-tilt contention issue without any explicit interprocess communication. Moreover, it ensured that no spectra collection occurred while the rover was stationary during planning activities.

### 4.4.3 Field Test Procedure

This section describes field tests at the Amboy Crater site. The tests demonstrate autonomous surficial mapping of basaltic lava units with dynamic Maximum Entropy Sampling that exploits surface and orbital data. The rover system begins with virtually no prior information about the

correlations between these data products; instead, it learns the appropriate model parameters on the fly. The main instrument in use by the rover is its Visible Near-Infrared Reflectance Spectrometer, which it learns to correlate to remote sensing data.

Here we focus on the “Bonnie and Clyde” site previously considered in Chapter 3. The entire exploration corridor measures 300m in length, and contains several distinct basaltic platforms. Figure 4.22 shows DOQ overflight imagery of the traverse area together with a rover plan produced after 27 spectral samples (colored dots). The path planning algorithm chooses a strategy that covers the principal units of surface material within the corridor while respecting the time budget. The blue line indicates one isocontour of the marginal prediction entropy; in this case entropy of the black basalt is high because no sample from this patch has been collected.

The entire suite of tests compared several mapping techniques:

- **Fixed Transect:** An exploration scheme consisting of a straight drive across the exploration corridor (Figure 4.20 Left). We derive class labels for individual map pixels by binarizing the final Gaussian process predictions into two classes using k-means quantization ( $k = 2$ , see section 4.4.4 for more detail).
- **Fixed Coverage:** An exploration scheme consisting of a zig-zag coverage pattern to uniformly cover the exploration corridor, subject to time constraints (Figure 4.20 Right). We alternate beginning with a right or a left turn. K-means produces pixel labels.
- **Adaptive ASTER:** Adaptive exploration using ASTER VIS/NIR orbital image with 15m/pixel resolution as a latent input. (Figure 4.21 Right). K-means produces final pixel labels.
- **Adaptive DOQ:** An exploration scheme utilizing the 1m/pixel overflight data in one visible band as a latent input (Figure 4.21 Right). K-means produces the final pixel labels.
- **Synthetic Random:** A map that does not correspond to any physical traverse; here both pixel class labels are chosen with equal probability. This is not a practical strategy, but we use it as a control group for the experiment.
- **Synthetic Uniform:** Another synthetic map, filled with pixel labels corresponding to the predominant class. Again, we use this as a control group for our comparison.

Time budgets were 24 minutes for each physical exploration method except the trials with DOQ overflight data. Two of these DOQ trials used a budget of 30 minutes, and a third (the examples appearing in Figure 4.22) was handicapped with a budget of only 16 minutes. The expedition ended before we could perform time-equivalent DOQ tests, so we encourage caution in interpreting the statistical significance of the DOQ result.

The experimental trials began by initializing the rover’s position estimate using onboard differential GPS and calibrating the spectrometer. The trials each began at the same location and ended when the rover reached the far end of the corridor or when the allotted time was exhausted. The

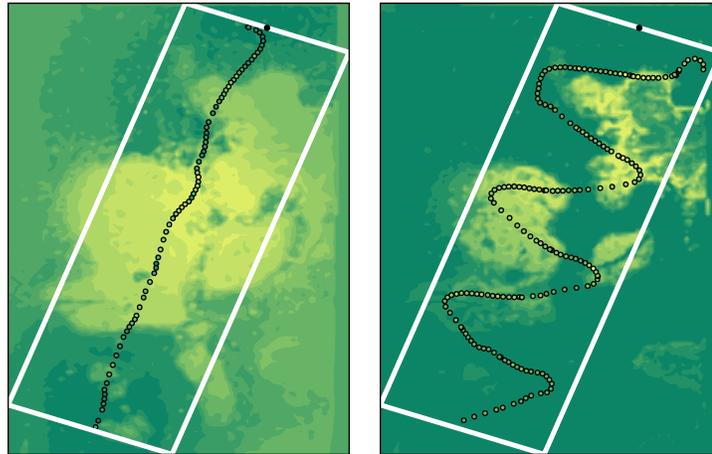


Figure 4.20: Examples of fixed navigation strategies with the resulting map reconstructions. Left: Transect-based navigation. Right: Farming pattern for maximal coverage. The Farming pattern is subject to the same time constraints as the adaptive strategies.

initial plan consisted of a straight-line drive towards the goal; only after enough spectra were collected to initialize the Gaussian process model did the science planner begin producing adaptive plans. This usually occurred within the first 10-20 seconds after initiating rover motion.

In order to reduce inter-trial variance we discarded trials that encountered unusually catastrophic navigation errors. We discounted experimental trials in which a single obstacle avoidance maneuver lasted longer than 60 seconds, or when navigational difficulties required a manual operator intervention. This occurred in about 1/3 of the trials we conducted in this environment. Despite these efforts, variations in low-level navigation behavior would still result in different paths between runs even when the planner's selected waypoints were identical.

Additional sources of inter-trial variance are lighting conditions and the varying period between spectrometer calibration actions. In order to control for varying light levels we performed tests when the sun was at high elevation. During our field season this limited our operating hours to the approximate interval of 9:30-14:30. Recall that our architecture paused every 7 minutes during the traverse in order to re-calibrate the spectrometer, but would make one or more additional attempts in succession until a calibration procedure succeeds.

Another source of inter-trial variance was the failure of a calibration routine due to uneven ground. This required multiple consecutive attempts and often resulted in poor traverse results. Not only would the rover waste time performing multiple calibrations, but the extra time between white references often produced bad spectrometer data as the old calibration became increasingly obsolete due to lighting changes. To separate these implementation issues from those of path planning, we disregarded those traverses in which the rover missed more than one calibration opportunity in a row, or when it obviously missed the reference target during calibration and produced bad data as a result.

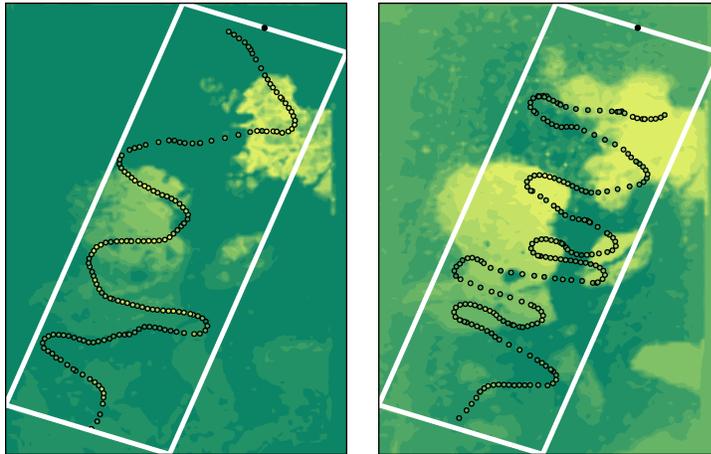


Figure 4.21: Examples of adaptive navigation strategies with the resulting map reconstructions. Left: Adaptive navigation using the low-resolution ASTER data product as a latent input. Right: Adaptive navigation using the high-resolution DOQ data product as a latent input. The low-level navigator is not constrained to travel forwards, and this run evidences some backwards travel during turns and obstacle avoidance.

A final confounding factor concerns registration and localization error. In practice we did not exclude runs from the experimental data on account of localization error, but it is worth noting it in any case as an additional source of inter-trial variance.

#### 4.4.4 Field Test Results

The accuracy score measures the fidelity of a reconstructed map based on the samples collected by each method. We evaluate each dataset by retraining a Gaussian process model using the high resolution remote sensing image, and binarizing the result into two geologic types with k-means vector quantization. Figure 4.23 shows one example of the map inference and the resulting binary geologic unit classification. We create a ground-truth map by manually labeling both classes of surface material in a high-resolution DOQ orbital image. Then, we associate the unsupervised classes with their most numerous matches in the ground-truth image. This yields a pixel-wise classification accuracy score for the reconstructed map.

The onboard adaptive exploration algorithm adapted to identify informative paths along the exploration corridor. Figure 4.22 shows the end of the trial after the rover has traveled 0.46 kilometers. In this trial, which lasted approximately 24 minutes, the rover chose a modified coverage pattern that found an appropriate compromise between boundary-following and coverage. This image illustrates the inference result with areas of dense basalt in yellow and clay sediment in green. The map prediction is uncertain in areas that are far from the rover path (consider the unvisited basalt patch in the lower right), but more accurate within the corridor. Samples generated by

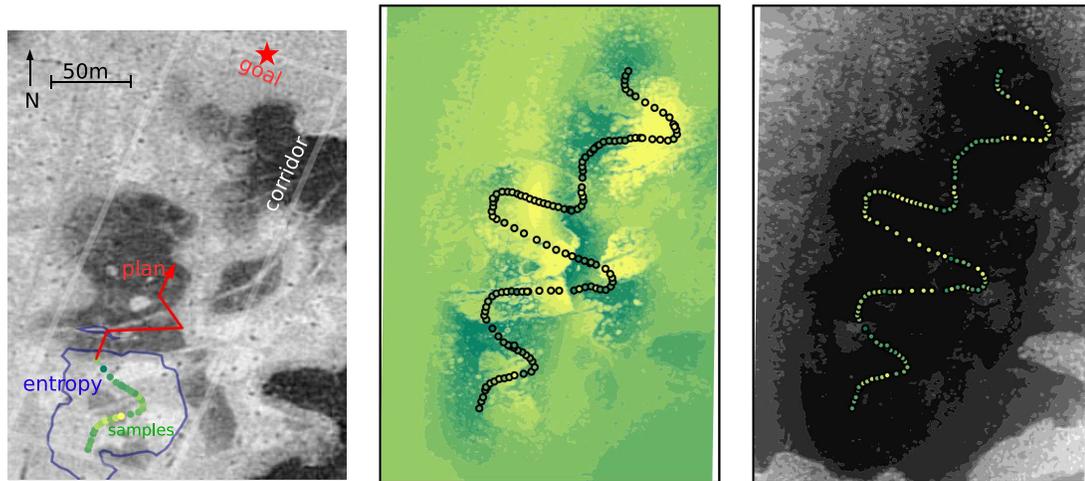


Figure 4.22: Left: Overflight imagery of the field site. Colored dots show spectral samples collected along the rover path. The blue line indicates one isocontour of the marginal prediction entropy. The red line shows waypoints in the current navigation plan. Center: Geologic map of surface material resulting from the traverse. Right: Marginal prediction variance associated with the reconstructed map.

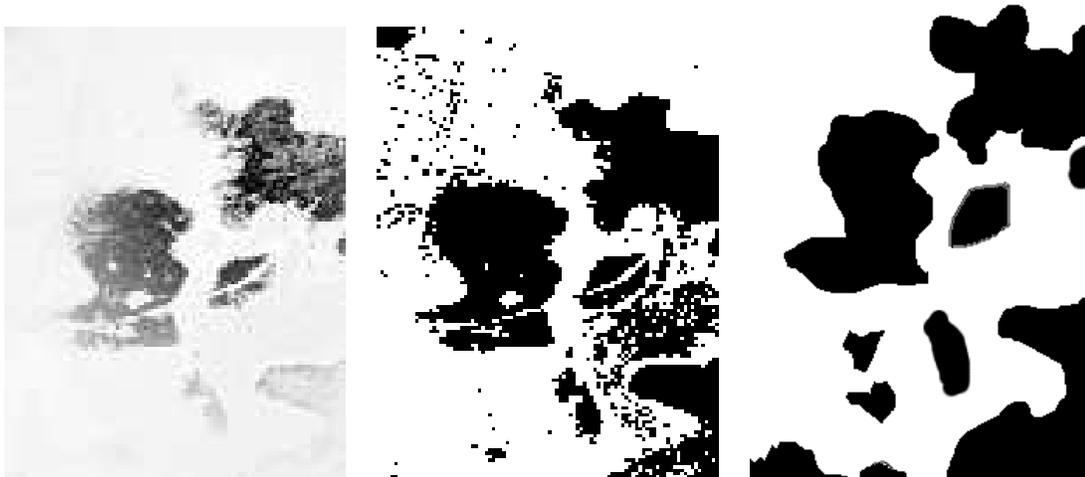


Figure 4.23: Left: Typical map inference result. Center: Binarized version resulting from k-means quantization. Right: ground truth labels. The map projection differs slightly in this image, but individual pixels were aligned using georeferencing data to calculate reconstruction error.

Method	Accuracy (std)	# Spectra (std)	p-value (n)
Adaptive			
DOQ	0.87 (0.01)	199.3 (45.3)	<0.01 (3)
ASTER	0.81 (0.03)	225.0 (32.8)	<0.05 (4)
Fixed			
Coverage	0.75 (0.05)	206.0 (14.4)	n/a (4)
Transect	0.74 (0.09)	75.9 (15.0)	0.45 (7)
Synthetic			
Uniform	0.57 (< 0.01)	n/a	0 (5)
Random	0.50 (< 0.01)	n/a	<0.01 (5)

Table 4.2: Results of the Amboy Crater Traverses: map reconstruction accuracy, the average number of spectra collected during each experimental run, and the p-value of a one-tailed T-test against the methods’ equivalence to a fixed coverage pattern.

the path planning algorithm offer a representative sample from which a human could more easily construct a complete interpretation of the environment.

Table 4.4.4 shows performance statistics resulting from 18 trials in this traverse area. Reported p-values result from a one-tailed T-test against the hypothesis that performance is equivalent to a static coverage pattern. Figure 4.24 plots the number of features (collected spectra) against reconstruction error. In general the adaptive methods outperform static methods, which in turn outperform the uniform and random cases. Traverses based on DOQ overflight data perform best of all.

DOQ trials are subject to the timing issues mentioned above, but we include this data here for several reasons. The low variance among the DOQ accuracy implies that the time budgets were not different enough to pose a clear handicap or advantage. Moreover, due to navigation irregularities between runs, the average number of measurements collected was actually *smaller* for DOQ trials than for the coverage patterns. The coverage patterns happened to drive through open terrain, resulting in fewer obstacle-induced pauses.

Several trends are apparent. First, with respect to the hand-labeled ground truth image, adaptive methods significantly outperform static coverage patterns for  $\alpha = 0.05$ . The variance in the number of collected spectra is greater for adaptive methods than for fixed methods, while the accuracy variance is reduced. This discrepancy — greater variation in path choice, and less variation in reconstruction results — might be related to the adaptive methods’ responsiveness to navigation obstacles and spectral imaging conditions.

Attention to the remote sensing data product ensures that the rover sees representative samples of diverse terrain from all parts of the corridor, without missing significant patches of basalt. While half of the fixed coverage patterns miss the final basalt patch (Figure 4.20), all the adaptive methods target this area. They tend to avoid homogeneous regions of sediment such as the Southeast and Northwest corners of the corridor.

Finally, Maximum Entropy Sampling’s border-affinity is evidenced in the rate of “crossings,”

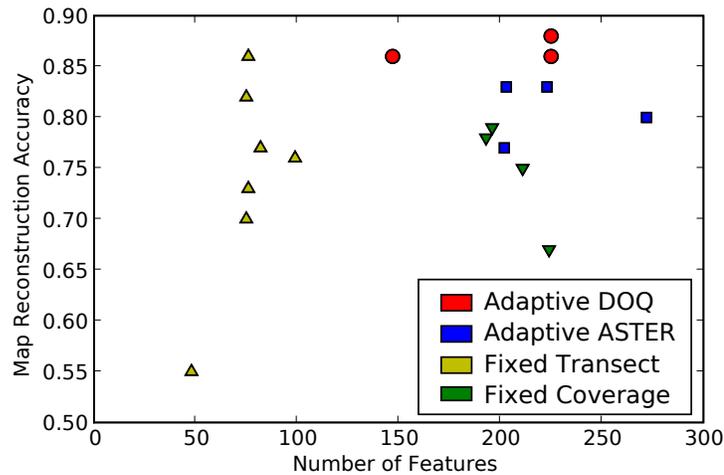


Figure 4.24: Results of the experimental trials from Amboy Crater: the number of spectral samples collected in each run and the accuracy of the reconstructed map.

that is, transitions onto or off of a basalt mound. Adaptive sampling utilizing the high-resolution orbital data crosses boundaries at a significantly higher rate than the fixed coverage pattern. The result is a higher number of samples from near boundary regions where the data in the augmented input space is sparser. In other words, adaptive sampling from orbital data permits extrapolation within each unit, rendering additional samples from the interiors redundant. The adaptive methods spend more time, on average, in basalt patches due to this regular border crossing (Table 4.3).

These experiments control for radical navigation errors, and we should not let this blind us to other advantages offered by adaptive sampling. Exploration robots operate in uncertain, dynamic environments. One can never be certain that activities like navigation or instrument deployment will succeed or take the expected amount of time and energy resources. However, the adaptive replanning allows the rover to recover from resource shortfalls with “next-best” exploration plans. Similarly, it can move opportunistically to exploit unexpected sample sites or resource surpluses. The power of active sampling lies not only as a device for active learning, but also as a hedge against execution uncertainty.

The catastrophic navigation failures from the Amboy trials are instructive; they show the value of adaptive observation planning to recover from execution error. Consider the traverse illustrated in Figure 4.25. Here the rover encounters a line of bushes that force it far outside the exploration corridor. This causes a delay of several minutes, and the rover falls behind schedule. Adaptive replanning permits the rover to recover a reasonable mapping strategy that satisfies the mission time constraints. It compromises in favor of the truncated coverage pattern shown in Figure 4.25.

The Amboy Crater experiments demonstrate a system that performs automated site survey on kilometer scales. It begins with virtually no *a priori* knowledge of surface materials. Instead it

Method	Time in basalt (std)	Crossings (std)	Crossing rate (std)	p-value (n)
Adaptive				
DOQ	60% (2%)	14 (4)	0.08 (0.01)	0.04 (3)
ASTER	56% (9%)	11 (2)	0.05 (0.01)	0.76 (4)
Fixed				
Coverage	48% (9%)	10 (3)	0.05 (0.02)	n/a (4)
Transect	53% (15%)	6 (1)	0.08 (0.02)	0.04 (7)

Table 4.3: Results of the Amboy Crater Traverses' boundary crossing behavior: The percentage of time spent on basalt platform units as determined by hand-labeled orbital imagery. The second column shows the number of crossings, *e.g.* the number of times the rover moved onto or off of a platform. The third column shows the proportion of boundary crossings as a fraction of the total number of observations. Finally we provide the p-value of a one-tailed T-test against the crossing rate's equivalence to a fixed coverage pattern.

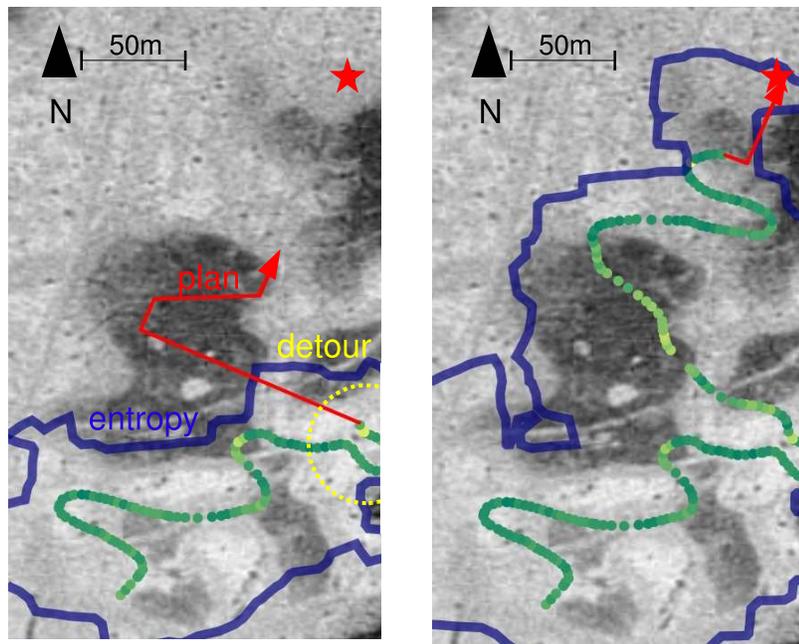


Figure 4.25: In this run, a navigation error delayed the rover for several minutes. The adaptive replanning was able to produce a reasonable, alternative exploration schedule in light of the delay.

learns the relevant distinctions on the fly and improves its exploration efficiency by leveraging learned trends in spatial proximity and correlations with remote sensing. The corridor exploration task exhibits some artificial constraints, but demonstrates the essential features of autonomous site survey.

## Chapter 5

# Selective Transmission with Proxy Features

And so on into winter / Till even I have ceased / To come as a foot printer, / And only  
some slight beast / So mousy or so foxy / Shall print there as my proxy. — Robert  
Frost, “Closed for Good” [124]

Remote science favors phenomena at the limits of perception, and scientists will inevitably consider features of the data that are too abstract or subtle for automatic pattern recognition. For example, geologists analyzing geomorphology of MER rover imagery have considered sediment structure [125], the shape and size attributes of rocks [37], and outcrops [82]. These present significant pattern recognition challenges; it is unlikely that the remote exploration agent can detect them all reliably. When specific feature detectors are possible, they require intensive development effort and may not generalize across environments [24].

Fortunately hidden or noisy phenomena do not preclude meaningful autonomous science. Remote agents can exploit correlations with *proxy features* that are simpler to detect. These proxy features consist of compact numerical descriptors of the data products that are easy to extract. The output of supervised pattern recognition, an unsupervised classification, or a dimensionality reduction procedure can all provide candidate proxy features. Proxy methods optimize the exploration objective by exploiting learned relationships between these descriptors and the actual science content of the data products.

Figure 5.1 shows a stark example based on a 2003 survey of scientists at the Jet Propulsion Laboratory [81]. The survey aimed to characterize agreement among scientists’ preferences for images in a mock selective data return scenario. Researchers presented scientists a diverse set of 25 images taken by the FIDO rover at a Mars-analog field site in Arizona. Each scientist ranked the images according to the order in which they would like to have received them from a selective subsampling algorithm.

The original study compares different scientists’ preferences and finds statistically significant correlations. While several outliers exhibited dramatically different preferences due to differences

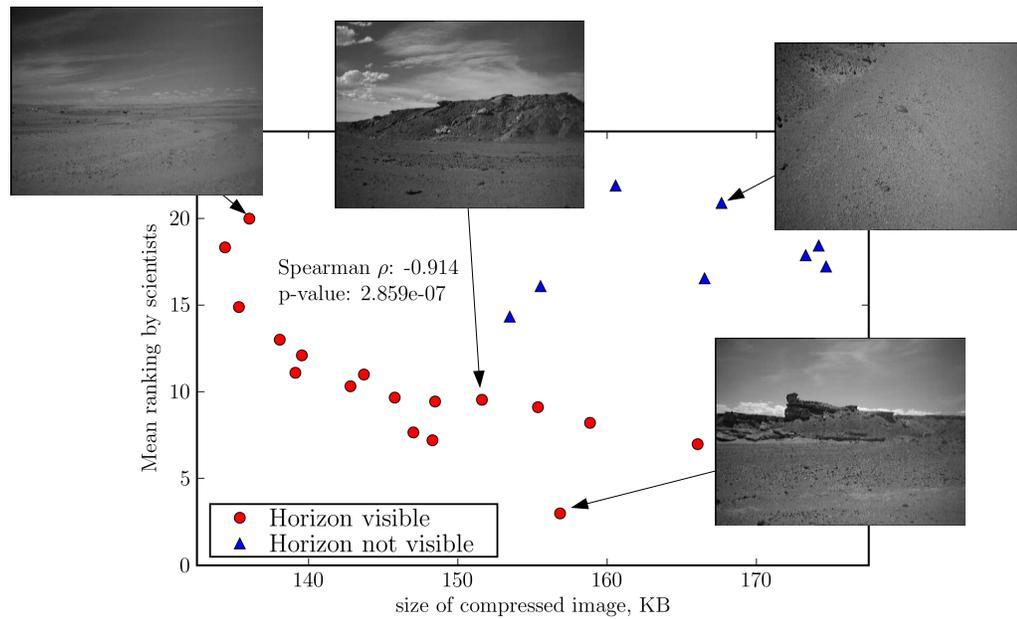


Figure 5.1: Scientist rank preferences correlate strongly with the compression rate for farfield images in the FIDO survey. Scientists broadly prefer images that compress poorly; these are more likely to contain interesting geomorphology. Here the  $\rho$  and  $p$  values refer only to those images which display the horizon.

Scientist	Spearman $\rho$	p-value
1	-0.69	0.002
2	-0.55	0.02
3	-0.82	0.00005
4	-0.79	0.0001
5	-0.39	0.11
6	-0.62	0.007
7	0.10	0.68
8	-0.28	0.26
9	-0.88	0.000002

Table 5.1: The Spearman rank correlation coefficient measures the correlation between preference ordering and compressed image size for each individual scientist. Preferences show a significant correlation in 6 of 9 cases.

in background or their ideas about mission goals, the overall picture was one of general agreement among different individuals. However, our post-analysis of these results reveals another significant trend. When compressed using the JPEG compression standard [126], the resulting *image size* exhibits a strong negative correlation with the scientists’ mean preference ranking. In other words, image compression rates appear to predict the science value of the image.

That JPEG compression correlates with scientist preferences is not especially surprising. The wavelet compression routine exploits regular structures in natural images, the compression rate is a rough measure of the “visual information” that is present (Figure 5.1). We divide the FIDO images into two groups: those that show the horizon and those that do not. For images with a horizon, a low compression rate signifies irregular structure like broken outcrops and other geomorphological features of interest. Scientists broadly prefer these over easily-compressed horizon images.

We can quantify the relationship between the compressed image size and scientists’ preferences using the Spearman rank correlation coefficient. The Spearman coefficient measures the dependence between rank-ordered variables. The statistic, commonly represented by  $\rho$ , is a non-parametric analog of the classic Pearson correlation coefficient for linear relationships. It measures of the ability of an arbitrary nonparametric function to describe the relationship between two variables. It is computed as follows, where  $d_i$  is the difference between rank of each of the values of the “compressed size” and “scientists” mean preference ranking” variables and  $n$  is the size of the dataset.

$$\rho = \frac{6 \sum d_i^2}{n(n^2 - 1)} \quad (5.1)$$

In this case the mean preference ranking’s value of  $-0.91$  is highly significant (p-value  $< 1e-6$ ). The compression rate accounts for a significant portion of the variance among preference rankings.

The correlation is also significant for the individuals’ independent preference rankings for 6 out of 9 individuals (Table 5). This suggests that an explorer agent could use the image compression rate to determine the best images for transmission. Regardless of whether or not the images are

actually compressed for downlink, the JPEG reduction constitutes a “proxy feature” that is simple to compute and germane to the science content of the images.

Proxy features can improve autonomous science systems’ flexibility by leveraging generic feature classification and pattern analysis in potential application to a wide range of science tasks. This contrasts with explicit feature detection that relies on the abstract language of science features to define the objective. Explicitly detecting “high-level” science features like outcrops requires new custom-made detectors for each specific science goal. But the cost of development and testing these detectors is prohibitive; detecting rocks in images is still an open research topic after years of research [24]. On the other hand, proxy methods can exploit correlated attributes at any level of description. With appropriate training data, automatic learning can utilize them to benefit many different environments and science phenomena.

In addition, proxy methods can provide quantitative insight into scientists’ field technique. They represent statistical relationships between quantifiable properties of data products and empirical sampling strategies. They lend quantitative insight to heuristic methods; for example, representative and periodic sampling can both be viewed as special cases of proxy exploration that make specific assumptions about the structure of the environment. Proxy methods provide an information-theoretic grounding for these heuristics while suggesting the circumstances under which they are most reasonable.

This section examines two proxy methods in the context of the intelligent mapping approach. We consider selective data return with information-driven reward criteria applied to generative spatial models. We use proxy features as a latent input to a regression model, where Gaussian process predicts a continuous scalar value in a time series. Experiments show that the proxy features are sufficient to improve data transmission actions when the true scalar value is unobserved by the agent. Next we consider an alternative approach models the joint probability density of appearances of proxy features and the real science contents, treating the proxy features as noisy observations of a hidden state. Here a Hidden Markov Model describes the geologic classes of images collected along a linear transect.

We show that classic representative sampling and periodic sampling strategies are simply examples of proxy feature data collection under different assumptions about sensor noise and the spatial continuity in the environment. Different degrees of correlation between proxy features and science content create situations for which the resulting policy approaches periodic or targeted sampling behavior.

## 5.1 Proxy Features as Covariates in a Regression Model

Recall that we can represent a set of collected data products as  $A = \{a_1, \dots, a_n : a_i \in \mathcal{A}\}$ . These are associated with independent experimental variables  $X = \{x_1, \dots, x_n : x_i \in \mathcal{X}\}$ , such as the samples’ physical locations. We represent the hidden science content of the data products with the set  $S = \{s_1, \dots, s_n : s_i \in \mathcal{S}\}$  in the space of science contents  $\mathcal{S}$ . Each  $s_i$  is apparent to a scientist looking at the data product, but in general it cannot be detected by the remote agent. However, the

agent does have a corresponding set of proxy features  $Y = \{y_1, \dots, y_n : y_i \in \mathcal{Y}\}$  with elements corresponding to each of the collected data products.

We employ a decoupled objective function and assume that the set of collected data is fixed during selective data transmission (Section 3). The explorer can only return a size  $m_T$  subset of data products. This leads to the following utility function for selective downlink:

$$U(A, A') \equiv R_T(A') \quad \text{for } A' \subseteq A, \quad |A'| \leq m_T \quad (5.2)$$

Proxy features are relevant for a wide range of information gain reward functions, including the mutual information and maximum entropy criteria. Here we will focus on mutual information. A reasonable objective is to maximize the mutual information between the science content of untransmitted data products (denoted by  $S \setminus S'$ ) and the science content of transmitted data products (denoted by  $S'$ ):

$$R_T(A') = I(S \setminus S'; S') \quad \text{for } A' \subseteq A, \quad |A'| \leq m_T \quad (5.3)$$

We omit the constraints in future equations for clarity.

The proxy features serve the dual purpose of indicating important data products for return and simultaneously providing cues to Earthbound scientists about the content of other data products that are not returned. This role only makes sense if the proxy features are more compact than the original data product. Practical proxy features are simple numerical attributes such as discrete classifications or numerical descriptors in  $\mathbb{R}^n$ . We can safely assume for our task that both proxy features and independent variables will require trivial bandwidth to transmit. We include the *entire* set of  $X$  and  $Y$  associated with all data products are included in the downlink and absorb their cost into our transmission budget. The discretionary portion remains a constant. We can rewrite the objective function:

$$R_T(A, A') = I(S \setminus S'; S', X, Y) \quad \text{for } A' \subseteq A, \quad |A'| \leq m_T \quad (5.4)$$

Proxy features let us choose  $A'$  to improve this objective without observing  $S$  or  $S'$  directly.

### 5.1.1 Shared Hyperparameter Learning

In our Gaussian process formulation we consider the hidden science content to be an underlying function  $f(x)$  of the independent variables. The resulting regression problem learns the mapping from independent variables to science content:

$$f(x) : \mathcal{X} \mapsto \mathcal{S} \quad (5.5)$$

Here for simplicity we represent science content as a single scalar value so that  $s_i = f(x) \in \mathcal{S}, \mathcal{S} = \mathbb{R}$ . Recall that for a set of input locations  $X$  we define the Gaussian process prior using a mean vector and covariance matrix. Following convention, we take the mean vector to be zero so that the prior is defined solely in terms of the covariance matrix  $K$ . For a vector  $f(x)$  representing  $f(X)$  evaluated at multiple locations, the prior can be written:

$$P(f(X)) = \frac{1}{(2\pi)^{N/2} |K|^{1/2}} \exp \left\{ -\frac{1}{2} f(X)^T K^{-1} f(X) \right\} \quad (5.6)$$

Each element of the prior covariance matrix is given by the covariance function  $\kappa(x_i, x_j)$  evaluated between two observation sites.

$$f(X) \sim \mathcal{N}(0, \Sigma) \text{ where } \Sigma = K_{A,A} \text{ and } K_{a_i, a_j} = \kappa(x_i, x_j) \quad (5.7)$$

Here we take  $\kappa$  to be the squared exponential covariance function, parameterized by coefficients  $\psi_1, \psi_2$ , and “length scale” hyperparameters  $w_k$  for each dimension  $k$  of the input space.

$$\kappa(x_i, x_j) = \psi_1 + \psi_2 \exp \left\{ -\frac{1}{2} \sum_{k=1}^d \frac{(x_{ki} - x_{kj})^2}{w_k^2} \right\} \quad (5.8)$$

Thus, the covariance with respect to  $S$  between input points  $X$  is given by the covariance matrix  $K_{A,A}$  which we can separate into submatrices associated with the transmitted portion  $K_{A',A'}$  and the remainder  $K_{A \setminus A', A \setminus A'}$ :

$$K_{A,A} = \begin{bmatrix} K_{A \setminus A', A \setminus A'} & K_{A \setminus A', A'} \\ K_{A', A \setminus A'} & K_{A', A'} \end{bmatrix} \quad (5.9)$$

The mutual information of the science content of the returned data with respect to the unobserved datapoints can be written in terms of differential entropies defined in terms of the submatrix determinants [127].

$$R_{\text{MUTINF}}(A, A') = H(S'|X, Y) + H(S \setminus S'|X, Y) - H(S|X, Y) \quad (5.10)$$

$$= \frac{1}{2} \left[ \log(2\pi e^{|A'|} \det(K_{A', A'})) + \right. \quad (5.11)$$

$$\left. \log(2\pi e^{|A \setminus A'|} \det(K_{A \setminus A', A \setminus A'})) - \right. \quad (5.12)$$

$$\left. \log(2\pi e^{|A|} \det(K_{A,A})) \right] \quad (5.13)$$

Note that  $S$  does not appear in the final decomposition. Similarly, the MES reward function can be written without reference to  $S$ .

$$R_{\text{MAXENT}}(A, A') = H(S'|X, Y) \quad (5.14)$$

$$= \frac{1}{2} \log(2\pi e^{|A'|} \det(K_{A', A'})) \quad (5.15)$$

$$(5.16)$$

For a fixed covariance function, the hyperparameters and independent variables completely determine elements of  $K_{A,A}$ , which is itself a function of input points  $X$ . Thus, the information-optimal sampling strategy is independent of the actual science content (the values taken by  $S$ ). This is true for any information gain metric based on a fixed covariance function; the independent variables alone define the sampling locations that will provide the most information about the science content of untransmitted data.

We use the proxy features to augment the independent variables, yielding a combined latent input space:

$$f(x, y) : \mathcal{X} \times \mathcal{Y} \mapsto \mathcal{S} \quad (5.17)$$

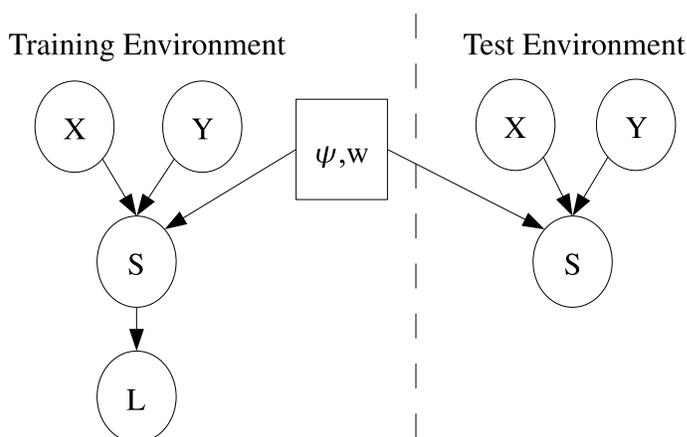


Figure 5.2: A graphical representation illustrates learning shared hyperparameters  $\psi$  and  $w$ . Along with these hyperparameters the observed proxy features  $Y$  and independent variables  $X$  define a Gaussian distribution over science contents  $S$ . Hand-chosen training values  $L$  permit maximum-likelihood learning of the hyperparameters. In a new test environment, the observed values and hyperparameters define a prior Gaussian distribution over science contents. The prior covariance matrix gives the best subsampling strategy.

This changes the number of input dimensions but not the form of the covariance function, so we can still compute the reward functions of Equations 5.13 and 5.16 without observing  $S$ . However, the result will now reflect samples' similarity not just in the spatiotemporal domain but also in the domain of computed proxy features.

This procedure requires an appropriate covariance function and hyperparameters. Here we will learn the hyperparameters from a set of labeled training data whose covariance relationships are presumed to generalize to the test set. If the science content is known in the training case, the agent can use this dataset to learn the covariance function parameters for the general augmented input space using classical gradient ascent of the likelihood or Markov Chain Monte Carlo. Figure 5.2 illustrates the process, in which hand-labeled contents  $L$  inform parameter learning on the hyperparameters  $\phi$  and  $k$ . These produce a prior covariance matrix on  $S$  to permit experimental design in the test environment.

This shared hyperparameter learning is appropriate given the natural interpretation of the covariance function as a prior over regression functions. The covariance functions capture similarities in the structure of the distributions of science contents in the test and training environments. They reflect the degree of correlation between input dimensions (or proxy features) and science contents. As long as the covariance function is determined by relative locations in the input space, the assumption of a common covariance function fits naturally with the hierarchical model. Researchers have investigated similar ideas under the heading of hierarchical Bayesian hyperparameter learning or multi-task learning [128].

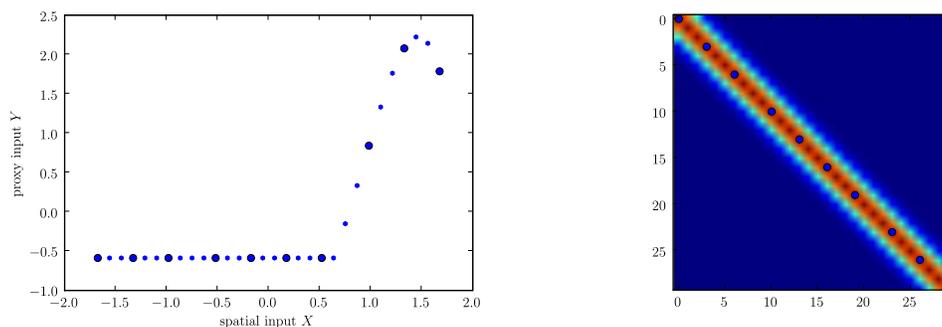


Figure 5.3: Data return with Maximum Entropy Sampling, based solely on the independent spatial variables  $X$ . Large blue dots indicate transmitted data. Left: collected datapoints. Right: Covariance matrix, with red showing large values and blue small values. Covariance relationships are determined by data points’ proximity, resulting in a periodic sampling strategy.

Some examples illustrate proxy features’ influence on sampling behavior. Here we consider the two-part function portrayed in Figure 5.3. This simulated environment consists of a flat homogeneous region and “bump” of high change. When the agent only has access to the the spatial dimension, our experimental design reward criterion seeks to fill the available space. This amounts to periodic sampling in the spatial domain. Figure 5.3 shows the covariance matrices and sampling behavior for this simple data return scheme.

Figure 5.4 shows a similar scenario with the Gaussian process trained on a two-dimensional input including an arbitrary real-valued proxy feature  $Y$ . Here, the signal-to-noise ratio is high; according to labeled training data,  $Y$  is an excellent correlate of the science content. The length-scale parameters associated with each input dimension, and the resulting covariance matrix, reflect this correlation. The Gaussian process determines that one can observe the contents of the large, homogeneous portion of the input space with just a few measurements. Science content correlates strongly along this dimension, obviating the need for dense sampling. The selective return strategy allocates a larger portion of the data return budget to sampling in the proxy features’ domain. This is tantamount to representative sampling in feature space that de-emphasizes physical proximity.

A key difference between treating proxy features as latent inputs rather than, for example, a noisy observation of  $S$  (as one would employ in classical Gaussian process regression) lies in the the resulting information-optimal observations. By using them as an additional input we produce nonstationary sampling behavior that reflects the appropriate kinds of diversity the dataset should exhibit in order to be most informative about the true  $S$ .

In Figure 5.5 we train the Gaussian process hyperparameters on a hand-labeled dataset that exhibits a high-noise relationship between proxy features and science content. This results in behavior akin to periodic sampling. The length-scale hyperparameters reduce the emphasis on the unreliable proxy features and encourage representative sampling in the spatial domain.

By abstracting from science content labels, proxy feature methods can improve robustness

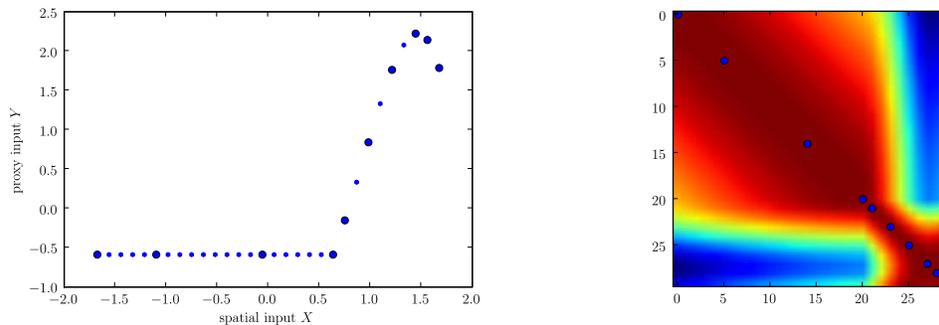


Figure 5.4: Data return with the Maximum Entropy criterion, based on both independent spatial variables  $X$  and low-noise proxy features  $Y$ . Covariance parameters were trained on science contents that correlated well with the proxy features. Due to the high signal-to-noise ratio, the Gaussian process determines that one can observe the contents of the large, homogeneous portion of the input space with just a few measurements. It devotes a larger portion of the data return budget to sampling in the proxy features' domain.

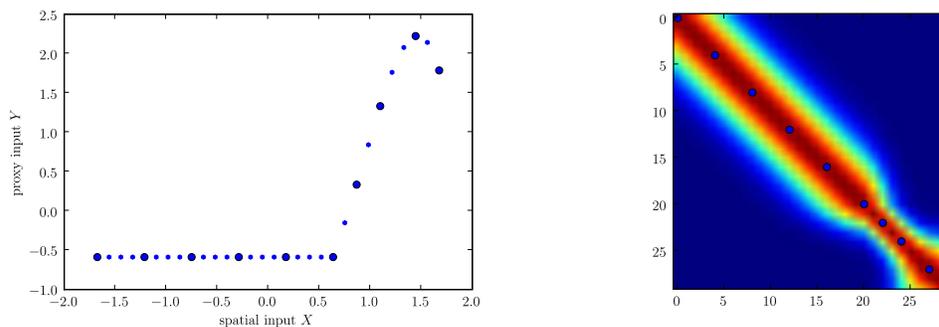


Figure 5.5: Data return with Maximum Entropy Sampling, based on a training set that exhibits a poor correlation (high noise) relationship with the underlying  $f(x)$ . Note that the covariance function favors relationships based on locality in the  $X$  dimension. The result is a more even, periodic spacing of transmitted data points in space.

and generality of selective data return. They enable the robot to adapt its experimental design without detecting the features of interest directly. One disadvantage of the method we describe is the requirement that scientists specify science contents in terms of a single scalar value. More sophisticated science content labels are possible; multiple independent scalar quantities could form an orthogonal basis for more refined distinctions. Alternatively, co-Kriging strategies [19] could extend the Gaussian processes to model multivariate science content.

### 5.1.2 Selective Return of Image Sequences

Image sequences are a compelling test case because they often contain complex, high-level science content that is difficult to recognize automatically. Cameras for engineering and science activities are pervasive in remote exploration and can generally capture data at a very high rate. However, the images themselves are large data products and complete sequences may be infeasible to transmit. Thus intelligent image subsampling would be immediately useful to many remote exploration domains.

Image analysis can help to classify the geomorphology of entire images. Previous analyses of rover panoramas have involved detecting rocks and clustering them into groups according to hue, intensity, size, and texture attributes [129, 18]. The proportions of each rock class constitute an “image signature” that reflects the local geology of the site (Figure 5.6). These site signatures can be reduced through techniques like Principal Component Analysis (PCA) to produce lower-dimensional visualizations of the terrain types and their distinctions. Similar efforts have classified the geomorphology of image regions using texture analysis [33].

Our mission scenario involves navigation image sequences drawn from an autonomous rover traverse in the Atacama Desert, a Mars-analog site in Chile [51]. The rover platform used for the experiments is “Zoë,” described previously in Chapter 2. Zoë conducted a series of multiple-kilometer autonomous traverses at the “Guanaco” site (referenced elsewhere as Site F), a location in the dry interior of the Atacama [51]. The Atacama terrain, visible in image 5.7, is barren with varying soil and rock morphology ranging from pebble-strewn plains to fields of large ( $> 30\text{cm}$ ) blocks.

At the Guanaco site expedition, four traverses crossed boundaries that remote geologists had previously identified as possible contact points between units of surface material. Figure 5.8 shows one such boundary in an ASTER remote sensing image (band 1). Surface images demonstrate clear morphological differences in surface material observed at each location (Figure 5.10). In this experiment we employ simple proxy features to interpret images for downlink. This results in datasets that improve the reconstruction of traverse geomorphology over those that focus simply on position.

Zoë’s navigation cameras provide a 60-degree field of view at  $320 \times 240$  pixel resolution. They point downward to view the terrain directly in front of the rover, and capture images at the rate of 1 per 2 meters of forward travel. We consider four traverse segments that each contain 400 distinct navigation images. These would be invaluable for determining relationships between surface material and orbital data, but the complete sequences are extremely bandwidth-heavy. Instead, we

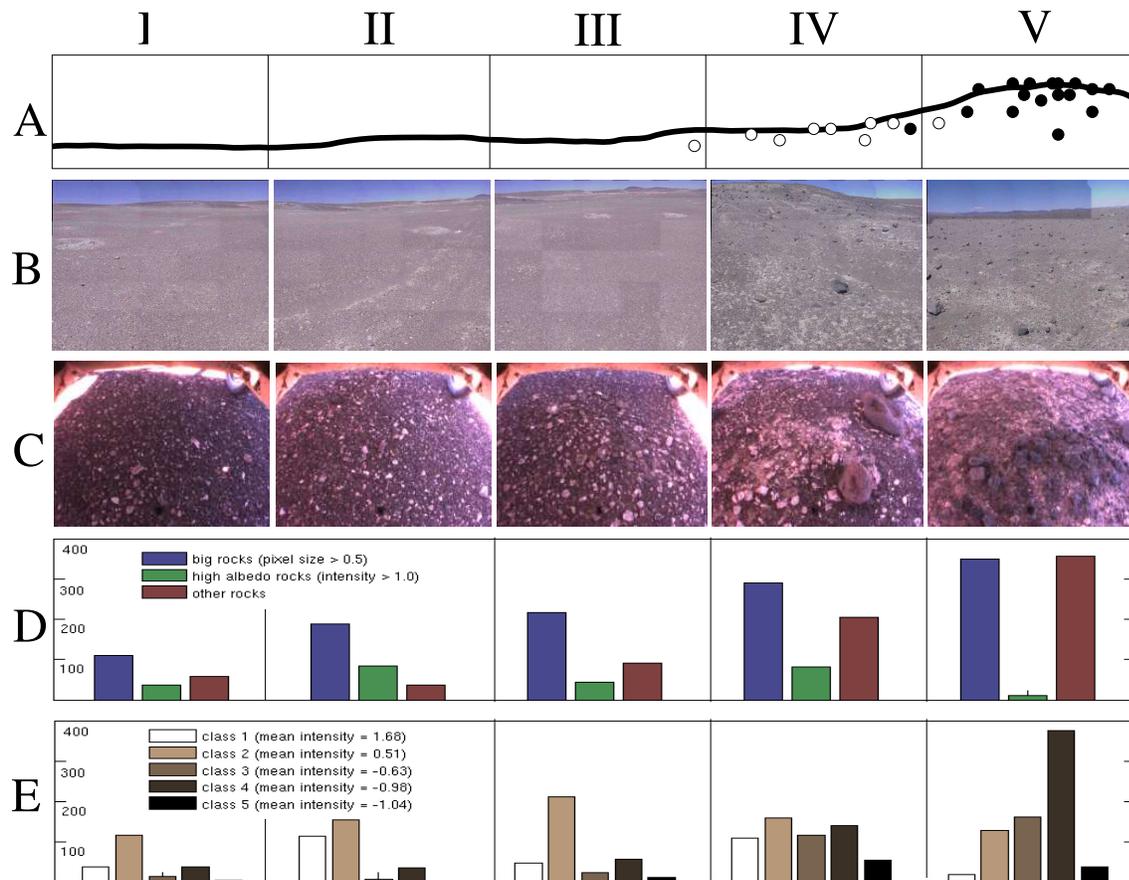


Figure 5.6: Five rover panoramas at different locales and the distributions of detected rock types [129]. Each column corresponds to one of five locales the field experiment. Row A: Qualitative human interpretation. The peak of the hill has the highest concentration of rocks. Row B: A portion of the panorama from each locale. Row C: Underbody images from each locale. Row D: Rock detection with manually defined “large rock,” “bright rock,” and “outlier” classes. Row E: Unsupervised model using rock detection; classification employs EM clustering over color and size features. The mean intensity of each class is shown.



Figure 5.7: Zoë in the Atacama desert, Chile.

desire a “best  $n$ ” subset that is most informative about the geomorphology of the entire traverse. Finally, we desire a rank ordering to ensure that the best images are transmitted first in case of communication faults or interruptions.

The flowchart of Figure 5.9 summarizes the training and subselection procedure for our selective image return scenario. We used a cross-validation procedure, training each of the four traverses using one of the others from a similar region. A human operator manually labels the contents of the training sequence by examining the image contents and choosing the most obvious visual transition between surface materials. We ascribed scalar values of 0 and 1 to images on either side of this border. In practice this usually corresponded to an obvious difference in soil texture or the quantity and size of rocks that were present. However, the Gaussian process representation can also accommodate more finely-tuned distinctions.

Image texture has been found in the laboratory to be a good correlate of geomorphology [38]; our analysis uses simple local texture descriptors extracted from the foreground of each image. We describe local texture using the *texton* method of Varma and Zisserman [130]. Textons are “archetypal textures;” pixel-level texture classifications that together fill the space of observed image textures.

We generate textons by convolving all images in a training set with the Maximum Response 8 filter bank. This results in an 8-dimensional response vector for each pixel. We then cluster the responses from a reserved set of training images using k-means to form a set of 16 universal textons. These textons, in turn, produce an image map that assigns each individual pixel to its

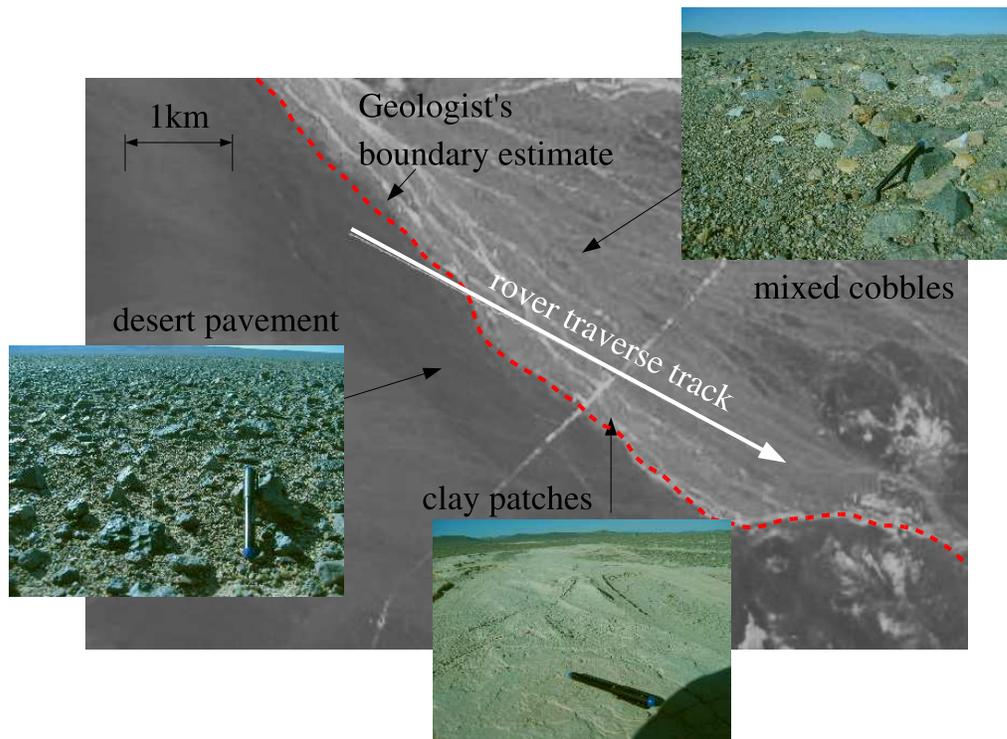


Figure 5.8: ASTER orbital imagery corresponding to a 4.3km traverse. The initial terrain consists of homogeneous desert pavement covered with small rocks ( $< 10$  cm). After the first kilometer the terrain changes (yellow dashed line) to a mixture of larger block sizes with occasional bare clay. Scientists evaluating the orbital data hypothesized a second boundary (red dashed line) but no change was clearly visible from the ground. In this annotated image a thin white line shows the traverse locations where navigation imagery was captured.

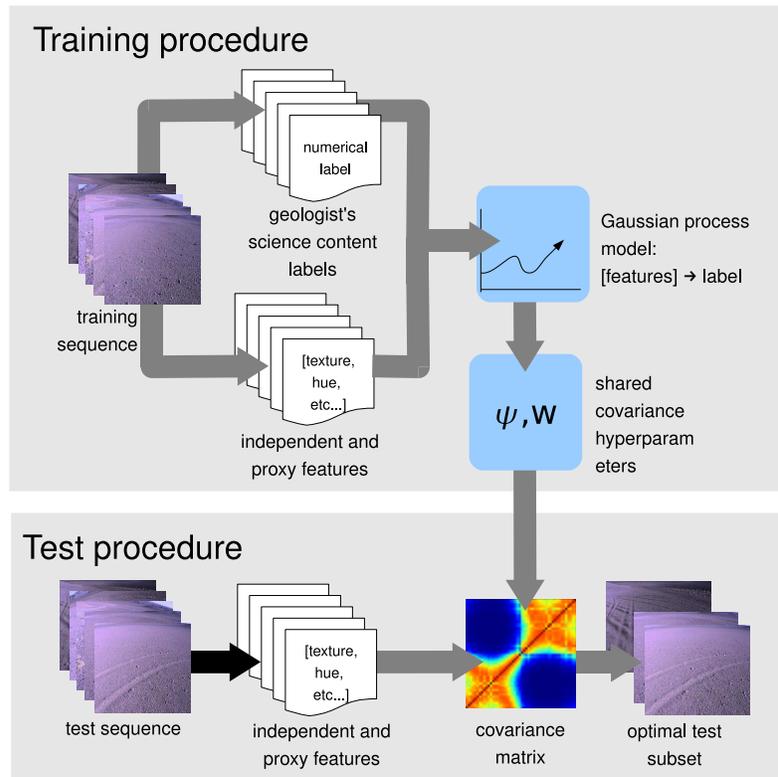


Figure 5.9: Procedure for subselection of an image sequence using proxy features.

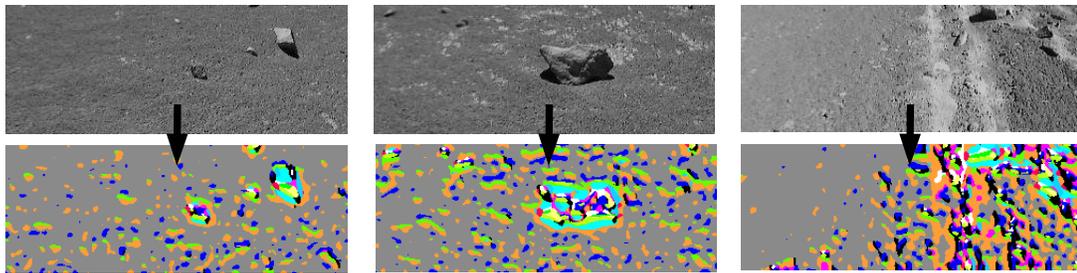


Figure 5.10: Original navigation images (top) and false color image showing texton values of individual pixels (bottom). To limit horizon and sky artifacts, we use only the bottom half of each image to compute its summary descriptor.

Euclidean-nearest texton (Fig. 5.10). Different textons approximately capture the different kinds of surface materials such as rocks, shadows, sediments of varying texture, and salt deposits.

The histogram of texton counts in each image provides a cue about its geomorphology. We normalize texton counts and compress the result to a  $1D$  representation using Principal Component Analysis as in chapter 3. The first principal component produces a latent input that augments the input space of a Gaussian process. We use a squared exponential covariance function, with length scales that determine the importance ascribed to differences in the spatial and morphological domains.

We determine the appropriate length scales by maximum likelihood on the training traverses. For a new traverse, we use the same covariance function to generate the prior covariance matrix as per Equation 5.7. This facilitates selective subsampling using the MES criterion of Equation 5.16. Greedy optimization identifies an ordered subset for downlink, adding images to the dataset one-by-one in the order that maximizes the determinant of the submatrix  $K_{A',A'}$ . The final data set  $A'$  is the best  $m$  images that comprise the transmission downlink. The remote scientist also receives the independent variables and proxy features associated with all of the other images that the agent collected. The complete set of proxy features is a required so that the posterior distribution over science contents matches that which was computed by the remote agent. In practical terms, the proxy features point out images that have similar contents. This assists the scientist to draw inferences about those images that were not returned.

Figure 5.11 shows an example of a covariance matrix for one such image sequence. This traverse was drawn from an image sequence near site C, waypoint 34 [62]. This traverse entailed a brief sprint across an open basin filled with uniform pebbles and occasional, sparsely distributed cobble-sized rocks greater than  $10cm$  in size. Near image 120, the rover approaches a hill where the density and size of rocks is thicker. In addition, the terrain texture changes, revealing patches of white sediment that become increasingly prominent. Finally, the last 50 images of the traverse show a drastic increase in the number of rocks as the rover nears the top of the hill. The result is that the images in the sequence exhibit varying rates of change.

We train the Gaussian process on unit labels from a similar traverse, and apply the covariance function hyperparameters to sample from this sequence. Figure 5.11 shows images returned from this traverse by greedy selection using the mutual information criterion. The numerals printed on each image indicate the rank ordering of the first 10 returned images from the dataset. The plot beneath shows a time series with the time step number on the horizontal axis, and the texture-based proxy feature on the vertical axis.

Here the selective data return strategy allocates a disproportionate number of images to the final third of the traverse, spending more than half of its downlink budget on those images. The result is a more representative sampling of the complete spatio-morphological domain. Note that at no point has the algorithm ever explicitly represented image morphology. Instead it draws cues from the learned degree of correlation between image labels, positions, and texton-based proxy features. The proxy features suggest that science content is likely to change rapidly in the final third of the traverse.

Selective data return with proxy features can improve the reconstruction of unobserved dat-

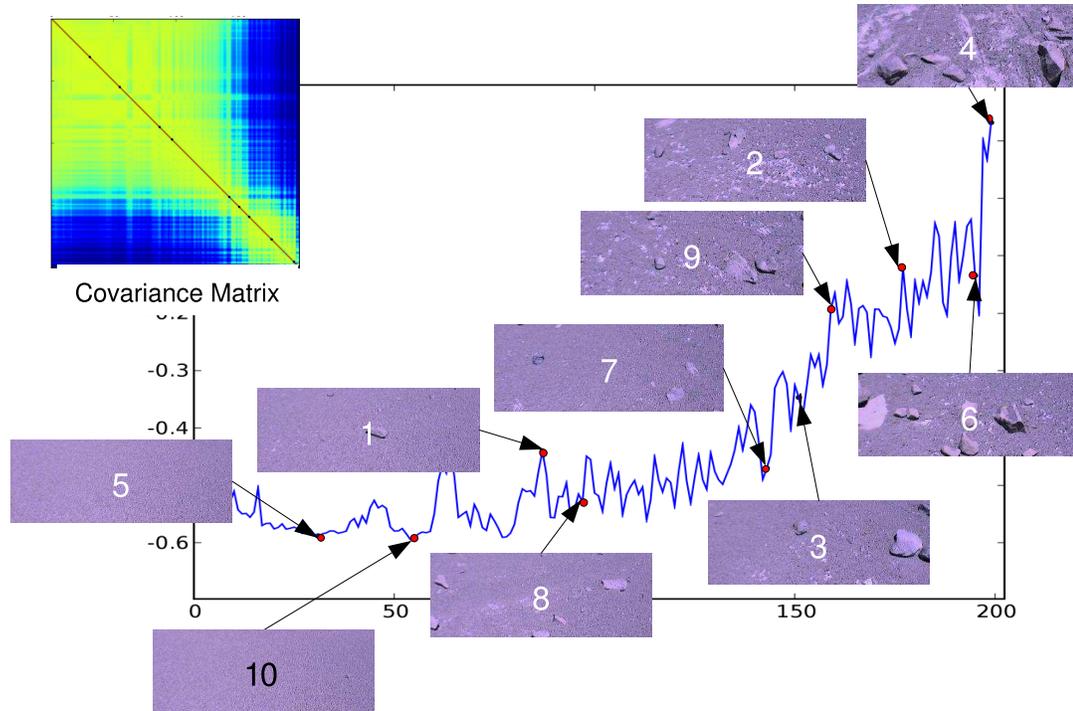


Figure 5.11: Sampling of a transect consisting of rover traverse imagery, based on a covariance matrix based on supervised training with proxy features. The horizontal axis shows the order in which images were acquired, while the vertical axis shows the texture feature used as a proxy feature. The numbers associated with each image show the order in which it was chosen for addition to the downlink.

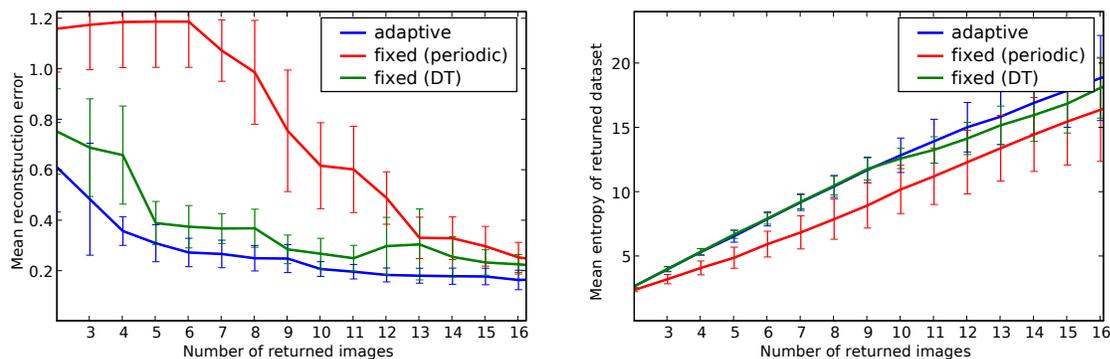


Figure 5.12: Left: Reconstruction error for hand-labeled images using adaptive and fixed selective return strategies. The plot shows the mean result of the four Atacama trials (confidence intervals for  $\alpha = 0.05$ ). Right: Differential entropy of the returned observation locations, calculated using the “true” covariance matrix fit to the actual observations.

apoints over naïve selection strategies. Figure 5.12 shows the mean reconstruction error for the four Atacama traverses with 95% confidence intervals. We evaluate each algorithm by revealing the true labels of returned data products and training a Gaussian process model on the result. This allows inference about the science contents at the remaining unobserved locations. Proxy features are provided to assist with the reconstruction *even* for those subselection methods that do not use them to compute the selective data return.

We compare three sampling strategies. The first is a *sequential* strategy that transmits images at evenly spaced temporal intervals in the order the images were collected. This performs poorly; only with the final few transmitted images does it provide representative samples from the entire spatial domain. Moreover, it ignores the actual content of the images themselves.

A second fixed strategy uses a time-filling heuristic: a distance transform that returns images that are temporally farthest from any other. It begins with the middle image; the next two returned images bifurcate each of the remaining halves, the next four images bifurcate the resulting quarter-traverse segments, *et cetera*. This reduces error far faster, but it still ignores information in the image domain.

Finally, an adaptive strategy utilizing maximum entropy returned images according to the order in which they were chosen through greedy selection. The adaptive strategy tends to fill the space as well, but it fills the joint position/texture space and in doing so provides representative samples of more different terrain types sooner. Indeed, the reconstruction exhibits a marked improvement, with approximately half the reconstruction error of the fixed selection strategies by the third returned sample. The use of mutual information instead of conditional entropy produced no significant change in reconstruction error.

The improvement from adaptive sampling is most significant for small subsets; here the adaptive data return ensures that the remote scientist receives representative samples from each terrain

type. Adaptive sampling reduces reconstruction error by nearly 50% when the agent returns just 4 data products. For larger subsets the average benefit accrued from adaptive sampling is still a meaningful percentage of the total reconstruction error. However, in these cases the overall reconstruction error decreases to the point where it is overwhelmed by between-trial variance, we cannot claim a statistically significant difference for these large subsets.

In summary, it is useful to formulate selective data return as an experimental design problem with a direct parallel to adaptive data collection. Here we can leverage proxy features as a latent input, which alters adaptive sampling behavior in a Gaussian process model. Experiments suggest that learned correlations can improve the overall performance of selective data return of image sequences.

## 5.2 Proxy Features as Noisy Observations

An alternative to latent input regression is to explicitly model the joint probability distribution of science contents and proxy features. This lets the explorer agent maximize any experimental design criterion that is based on the posterior distribution over science contents. Here we consider this approach with an alternative model that represents hidden science contents with discrete classification labels. Like the Gaussian process model it reflects both spatial and attribute-based dependencies, alternating between periodic and targeted sampling in response to various levels of detection noise.

We use a discrete-time Hidden Markov model to describe traverse content. As before, each observation is associated with one or more independent experimental variables. Here we will limit ourselves to a discrete time one-dimensional case in which the independent variable is the observation’s location in an ordered series. These are indexed by time or position along a transect. Since the data products are ordered, we can use subscripts to note that the observation  $a_t$  is associated with time step  $t$ .

To review, the images  $A = \{a_1, a_2, \dots, a_n\}$  in the traverse are each associated with hidden science content  $S = \{S_1, S_2, \dots, S_n\}$  and proxy features  $Y = \{Y_1, Y_2, \dots, Y_n\}$ . During data transmission phase the agent computes a downlink transmission  $A'$  to optimize a reward function  $R(A, A')$ . This reward should reflect the downlink’s science value subject to a constraint on the total number of returned images. We have advocated using a measure of information gain, and have previously used mutual information and conditional entropy objectives. Here we use the closely related sum of residual marginal entropies  $\sum_t H(S_t|A')$ , which facilitates an easy solution to the discrete-time optimization.

Choosing the return subset requires modeling correlations between proxy features  $Y_t$  and the actual label  $S_t$ . However, we also aim exploit spatial relationships between hidden variables themselves. Consecutive images are likely to have similar science contents; we capture these correlations using a Hidden Markov Model (HMM) [131] where each image in the sequence is conditionally independent of the others given its immediate neighbors (Fig. 5.13). This describes an underlying process composed of a hidden science “state” with occasional discrete transitions.

Hidden Markov Model parameters consist of initial state probabilities  $P(S_1)$ , transition proba-

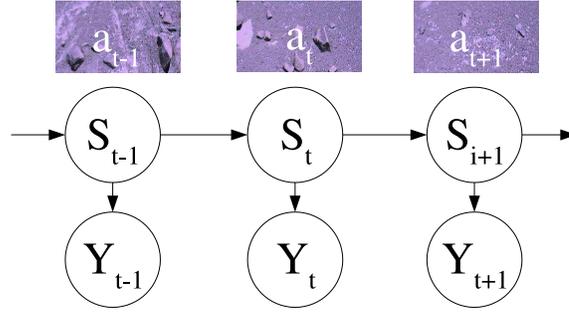


Figure 5.13: Hidden Markov Model of rover traverse imagery. Images available for downlink comprise the set  $A = \{a_1, a_2, \dots, a_T\}$ . Each is associated with a science content label represented here by hidden variables  $\{S_1, S_2, \dots, S_T\}$ . The agent examining the images produces noisy classifications resulting in proxy features  $\{Y_1, Y_2, \dots, Y_T\}$ .

bilities  $P(S_{t+1}|S_t)$  representing the probability of state transitions from one time step to the next, and emission probabilities  $P(Y_t|S_t)$  representing correlations between images' true labels and the automatic classifications.

Designers have several options for estimating HMM parameters. If scientists can enumerate all the image features of interest then standard supervised learning techniques can estimate HMM parameters from labeled training data. The Maximum Likelihood parameter estimates for discrete labels are given by the empirical counts:

$$\hat{P}(S_{t+1} = i | S_t = j) = \frac{\sum_t \delta(S_{t+1}=j, S_t=i)}{\sum_t \delta(S_t=i)} \quad (5.18)$$

$$\hat{P}(Y_t = k | S_t = j) = \frac{\sum_t \delta(Y_t=k, S_t=i)}{\sum_t \delta(S_t=i)} \quad (5.19)$$

In the following experiments we will assume that previous training data is available so that the HMM parameters are known in advance.

On the other hand, if the investigation is a general site survey the scientists might not know what to expect and the phenomena of interest represented by  $S_t$  may be ambiguous. In this case it might be better to estimate labels automatically at runtime. For example, the explorer agent could use the Baum-Welch algorithm [131] to estimate Maximum Likelihood HMM parameters. The best number of state labels is a model selection problem to be addressed by cross validation or measures such as the Akaike Information Criterion [132].

Following Krause and Guestrin [133] we base our reward function on the sum of posterior marginal entropies. The complete reward decomposes into the sum of individual rewards for each  $S_t$ .

$$R_T(A') = \sum_{a_t \in A} R(a_t, A') \quad (5.20)$$

$$= - \sum_{S_t} H(S_t | A') \quad (5.21)$$

Expected reward hinges on the conditional entropies of the hidden variables. We can calculate the expectation  $E[R_T(A')]$  in the following manner, where  $S$  represents the possible assignments to hidden variables associated with the image set  $A$ , and  $S' \subseteq S$  represents science content associated with downlink images  $A'$ .

$$E[R_T(A')] = -E\left[\sum_{S_t} H(S_t|A')\right] \quad (5.22)$$

$$\equiv -\sum_{S'_t} P(S'_t) \sum_{S_t} P(S_t, S'_t) \log P(S_t|S'_t) \quad (5.23)$$

Krause and Guestrin present a dynamic programming algorithm for optimal subset selection that exploits a decomposition of the local reward function [133]. It decomposes the reward into subchains that become conditionally independent when hidden nodes are revealed. Ultimately we need only evaluate combinations of possible assignments to two variables in  $B$  - those at either end of the current subchain. The resulting recursive algorithm can be used to compute  $R(A')$  in polynomial time.

In summary, selective return consists of the following procedure:

1. The agent performs a traverse and collects images whose science content is a sequence of hidden variables  $S = \{S_1, S_2, \dots, S_T\}$ .
2. Onboard pattern recognition processes the images to yield the sequence of noisy classifications  $Y = \{Y_1, Y_2, \dots, Y_T\}$ .
3. The agent chooses the subset of images  $A'$  to optimize  $R_T(A')$  from equation 5.23.
4. Scientists receive the image subset and observe the science content  $S'$  in  $A'$ .

A series of simulated trials investigated the effects of different environments and classification noise on data return behavior. We synthesize data from a traverse sequence in which  $S_t \in \{0, 1\}$  and  $Y_t \in \{0, 1\}$ . Each  $S_t$  maps onto its corresponding binary-valued classification with a small probability of error. Similarly, the environment has some small probability of a state transition between sampling locations.

Fig. 5.14 illustrates typical data return behavior with a simulated traverse containing a single state transition. The agent mitigates observation uncertainty by returning images surrounding the likely transition; this helps the receiver to identify its precise location. Fig. 5.15 shows a traverse with equivalent states and observations. In this second case the agent's model parameters imply frequent state transitions and a poor correlation between science content and the automatic classification. The agent considers its proxy classifications to be less informative and opts for a more conservative policy. Its strategy is closer to periodic sampling with returned images spaced regularly throughout the sampling domain.

We computed data return strategies for a series of simulated traverses that presented a discrete-time analog to the noisy step function. Here a sequence of 30 collected images contained a single state transition; these were subsampled using greedy selection to a 6-image downlink dataset. Fig.

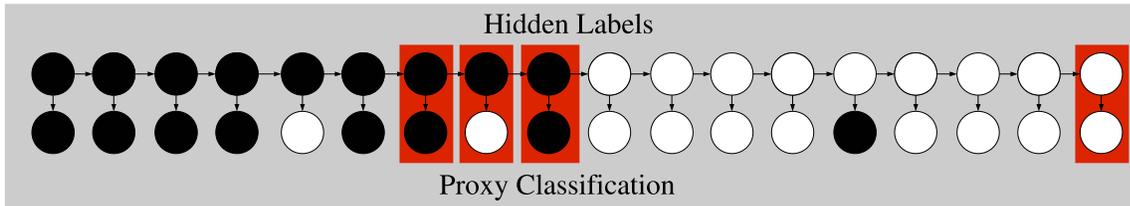


Figure 5.14: Typical behavior of the data return algorithm in the noisy binary case, with nodes colored according to their value. Only the the proxy classifications  $Y_i$  on the bottom row are visible to the agent. Rectangles indicate images selected for return. This example uses transition and error probabilities of 0.1. Subsamples consist of images clustered near likely state transitions.

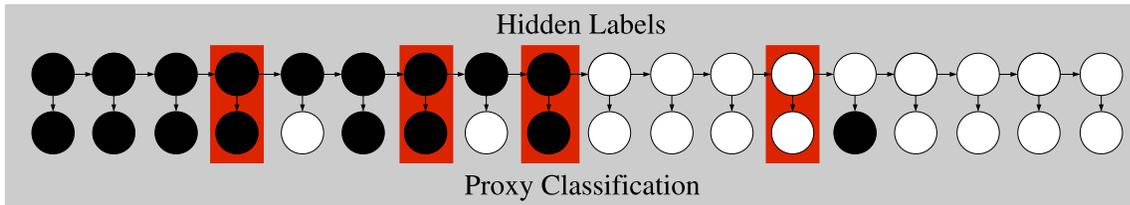


Figure 5.15: An identical segment, but presuming transition and error probabilities of 0.33. This increased noise means the automatic classifications are less informative. It results in a more conservative subsampling strategy.

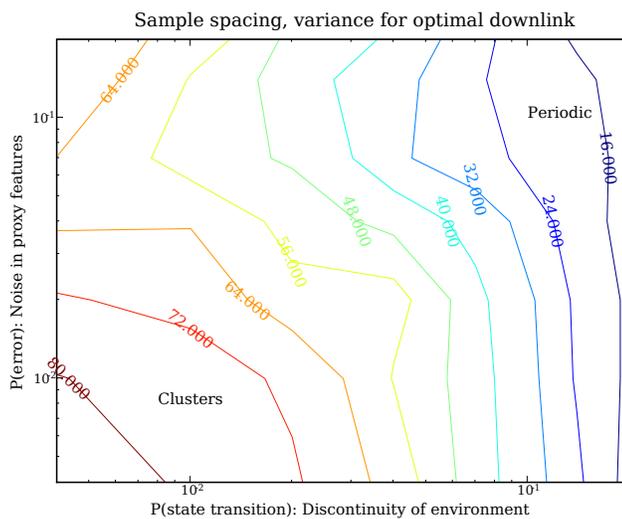


Figure 5.16: Variance in spacing of image return as a function of noise in the environment and onboard classification.

5.16 characterizes the resulting downlinks for various levels of transition and observation noise. Isocontours show the variance in the distance between samples; high variance suggests irregular clusters of samples such as those in Fig. 5.14 while low variance indicates evenly-spaced sampling as in Fig. 5.15.

The case of high detection noise is analogous to the low signal-to-noise Gaussian process model (Figure 5.5). Both favor a periodic sampling strategy. Alternatively, low noise favors downlinks that are clustered in space, focusing instead on boundary transitions in the proxy features. Formulating selective data return as compression prescribes a return strategy that reflects both the context of local observations and the fidelity of onboard pattern recognition to the underlying phenomena of interest.

### 5.3 Boundary detection with HMMs

The HMM provides a natural framework for mapping geologic state during long-distance autonomous traverse. We take the science content  $S_i$  to be a label of the geologic unit for image  $a_t$ . We apply the HMM selective return strategy to three of the traverses from the previous experiment (Fig. 5.17). The image analysis scheme classifies these images according to the class of surface material that is present.

Recall that the proportions of each image texton form a 16-dimensional feature vector which serves as a feature space for predicting the images' geologic unit type. We produce training data

automatically by indexing rover position against a georeferenced map of predicted unit boundaries drawn previously by geologists using orbital images. When the rover is very far from any possible boundary, we can be certain about the appropriate image label in spite of inevitable localization and registration errors. As long as geologists are fairly certain about the boundary position we can assign an unambiguous unit label to those images. This training data, drawn from the peripheries of the traverse, is set aside to train the system for application to segments where the boundary locations are uncertain. The image classification uses a support vector machine with a radial basis kernel function [134].

Table 5.2 shows the results of the classification procedure. The first column shows 10-fold cross validation training error, while the second shows test accuracy on over 1000 images from the Atacama image dataset. Test accuracy is substantially lower than cross-validation would suggest due to several factors that complicate generalization across traverses. The geologic units themselves are not completely homogeneous, and the geologic content of two traverses may differ despite having the same map label. Moreover, changing lighting conditions create shadows and cause systematic changes in the visual appearance of terrain. Two traverses may therefore exhibit markedly different features if they occur at a different time of day. However, for the purposes of this experiment it is not necessary that the image classification be perfect; rather, our goal is to account for this error to improve selective downlink.

The final column of table 5.2 shows the image label accuracy on test results after posterior smoothing of the HMM. The Viterbi algorithm [136] produces the Maximum Likelihood state sequence, which is a better correlate of the true unit labels than the independent image classifications. We include this score for interest only, to suggest that the HMM structure helps state estimation by

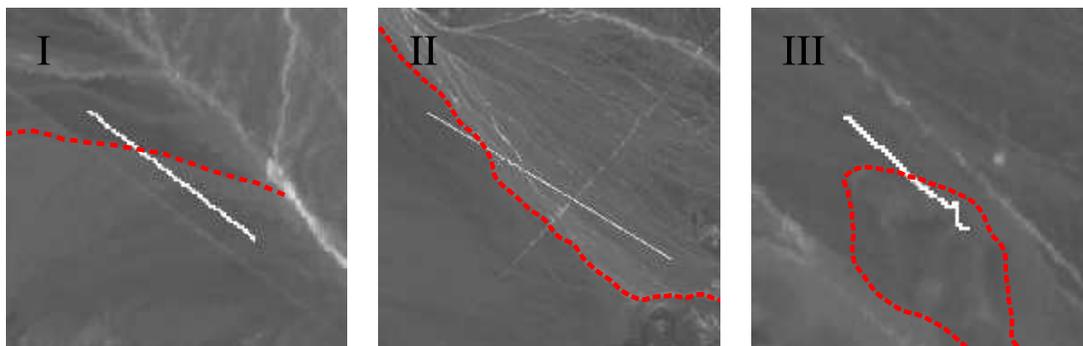


Figure 5.17: Three of the boundary-crossing traverses from the Atacama expedition. These Visible/Near Infrared orbital images were collected by the ASTER instrument [135] and used by geologists to draw preliminary geologic maps. The bright white line in each image represents the position track of the rover during the traverse, while the red dashed line shows the boundary hypothesized by the project geologist. Each image pixel represents approximately 15 square meters of terrain. Desert roads, drainages, and other high-albedo features are also visible.

Table 5.2: Classification accuracy for texture-based image features.

Traverse	Training Accuracy	Test Accuracy	State Estimation
I	96.8	82.0	84.0
II	100.0	88.5	96.2
III	97.4	88.3	98.1

capturing the continuity of the real environment. The Viterbi parse is an estimate of  $X_t$  and not  $Y_t$ ; it does not necessarily reflect the accuracy of the image classification procedure itself.

Each of the three test traverses consisted of a navigation sequence of 50 images; selective return produced a subset of 10 images for downlink. We used conservative HMM parameters: a transition probability of 0.9, and a classification accuracy of 0.8. The transition probability was fixed through trial and error to produce reasonable results for the training traverse. Note that this value is sensitive to the imaging rate; a denser image sequence would imply a lower state transition probability. Figure 5.18 shows the resulting locations of the returned images, the location of the boundary hypothesized by geologists, and the rover’s posterior probability estimate of the hidden state.

Of the three segments, the state estimation in traverse I exhibits the poorest match to the geologist-hypothesized boundary. We attribute this to particularly poor classification scores (Table 5.2). The distinctions between the geologic units was subtle, and the units themselves were highly heterogeneous (each containing mixtures of small of rock-strewn and open areas). Thus, the actual location of the true boundary is somewhat ambiguous. Moreover, field conditions complicated the selective return procedure. The traverse occurred late in the day with rapidly changing light conditions that independently altered image appearance over time.

Traverse II’s state estimation accurately identifies the geologic unit boundary. The return procedure fixes any remaining ambiguity with an image that corresponds exactly to the anticipated border. Other images selected for return include three locations later in the traverse. Here several images were classified as belonging to the first unit, which increased the probability of an improbable state transition. The resulting uncertainty over hidden state in these areas motivated the selective return procedure to spend extra bandwidth there.

In traverse III geologists had hypothesized a homogeneous traverse finished by a single discrete transition at the base of a hill. In fact, the approach to the hill was itself strewn with rocks similar to those present at the peak. This was reflected by a gradual transition rather than a clear boundary. It is unsurprising that image labels in this area were mixed, resulting in an uncertain, gradually changing state estimate. The selective return procedure reacted by allocating samples more-or-less evenly throughout this transition area. A clear difference in sedimentology is evident between the first and last images from the downlinked dataset (Fig. 5.19).

The field experiments demonstrate that time sequence models are appropriate for remote transect surveys. The HMM permits context-sensitive state inference and calculation of informative

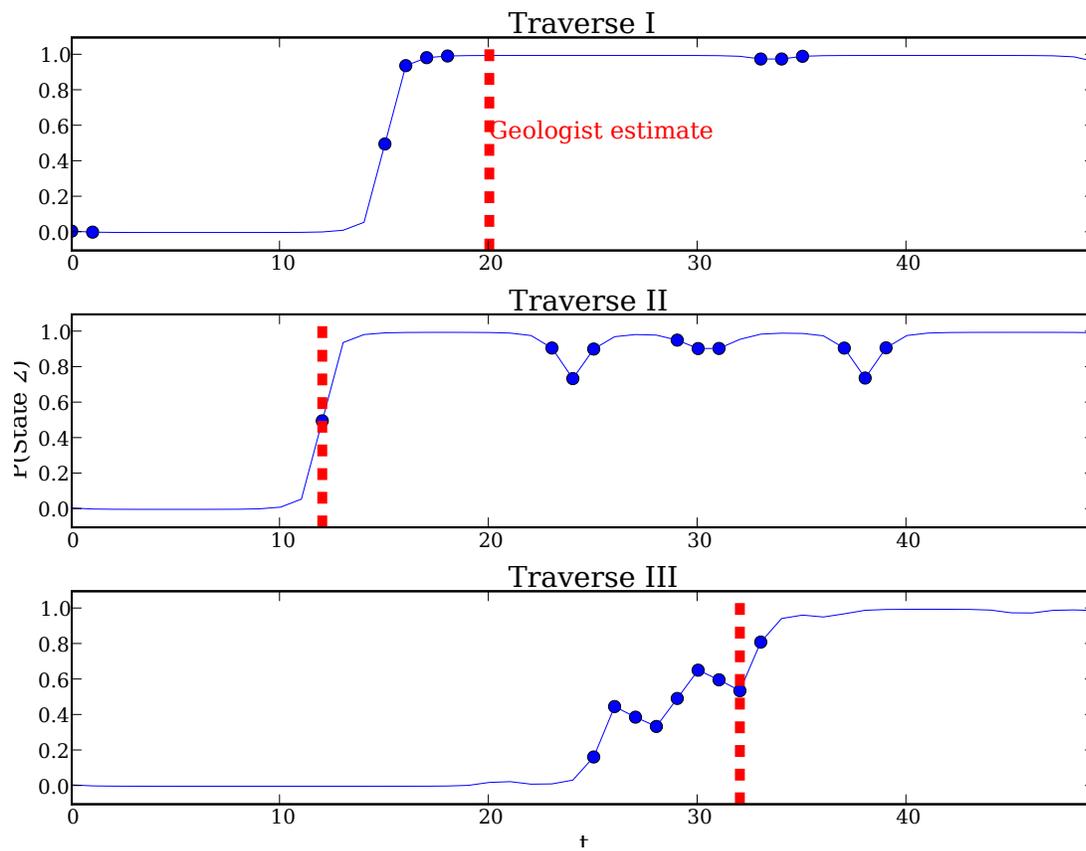


Figure 5.18: Returned images from the three traverses, along with the posterior state estimate. Selective return allocates images to areas where the hidden state is most ambiguous, such as border areas (top and bottom) and anomalous observations (middle)

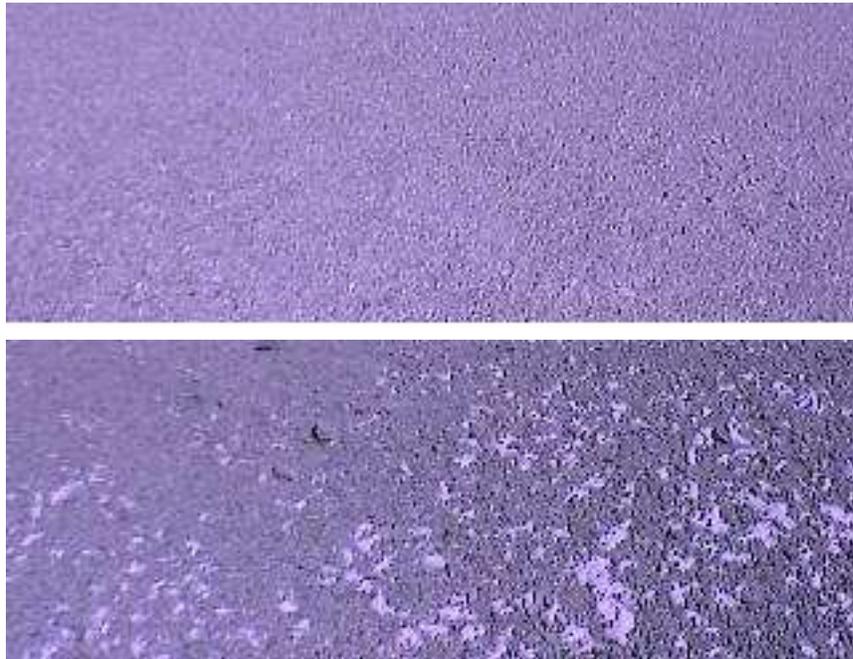


Figure 5.19: The first (top) and last (bottom) images returned from traverse III. Only the analyzed foreground portion of each image is shown.

observations for selective data return. Both the HMM and Gaussian process models we consider exhibit significant similarities in their behavior. Depending on the amount of noise in the proxy features they reach some compromise between periodic sampling in the spatial domain and targeted observations chosen based on these proxy attributes.

However, there are significant differences between the two approaches as well. The HMM model explicitly models a joint probability distribution over the proxy features and hidden science content. Thus it could also be used to predict the probability, not only of a particular hidden science content, but also future proxy features at other locations in the environment. The price for this modeling power is the Markov assumption, and limited potential for modeling higher-dimensional spatial phenomena that does not lie on a transect. Alternative topologies may still be tractable if the posterior is Gaussian; Gaussian Markov Random Fields [19] are promising candidates.

In summary, we have formulated selective data return as another mapping task — specifically, a experimental design problem in which we aim to communicate information about the state of the environment at unobserved locations. Here onboard analysis of data products can improve the data return objective. Additionally, recognizing the tenuous relationship between onboard perception and science content can improve robustness. Most strategies for science value assessment, such as the target signatures approaches mentioned in section 4, presume perfect feature detection. In practice it may be preferable to optimize mapping tasks using proxy features instead.



## Chapter 6

# Conclusions

Modern systematic maps rely on a standardized form of knowledge which establishes a prescribed set of possibilities for knowing, seeing and acting. They create a knowledge space within which certain kinds of understandings and of knowing subjects, material objects and their relations in space and time are authorized and legitimated. — David Turnbull, *Cartography and Science in Early Modern Europe: Mapping the Construction of Knowledge Spaces* [5]

This dissertation has advocated an *intelligent mapping* approach to remote exploration in which the agent learns an predictive spatial model, or map, of science phenomena and leverages it to improve exploration decision making. This model incorporates sensor data at multiple scales, such as local spectroscopy and orbital measurements. Map predictions that exploit spatial and cross-sensor correlations can improve exploration performance over static methods grounded in static observation plans.

### 6.1 Contributions

We have made significant contributions to the field of autonomous remote exploration. We demonstrated repeatable real-time automatic detection and classification of geologic features, and presented a new *spectrometer profile* data product that leverages rock detection to facilitate real-time instrument targeting. Our approach permits tracking multiple targets simultaneously, and resolves the correspondence problem to create a metric map of unique science features. The use of a dynamically targeted reflectance instrument and solar illumination produces considerable improvements in the quantity of gathered spectra. The method carries relevance for transect studies during mission plans that mix autonomy with scripted actions.

We have investigated Gaussian process spatial models for remote exploration that integrate remote sensing with surface measurements. Field data from the Amboy crater region has shown that the addition of remote sensing can offer a significant advantage in the reconstruction and prediction of spectrometer observations. We have investigated Gaussian process models for adaptive

exploration, and evaluated latent input spaces as a practical means for modeling nonstationary environments. Rover tests in field conditions demonstrate adaptive mapping that leverages on-the-fly spatial and cross-sensor inference.

In the action selection chapter we advocated an information-theoretic experimental design approach to remote exploration. We treat geologic site survey as spatial experimental design; a decoupled strategy treats data collection and downlink as separate sequential tasks. We have introduced latent input spaces as a new method for producing useful adaptive sampling behavior, and have explored qualitative and quantitative interactions between different latent inputs and experimental design criteria. We demonstrate how the inclusion of remote sensing data can significantly improve adaptive exploration decisions.

We have demonstrated an approach to rover exploration that employs intelligent mapping to guide adaptive navigation under time constraints. It begins with virtually no *a priori* knowledge of surface materials, and learns the relevant distinctions on the fly. The result is a new *corridor exploration* operational mode that holds promise as a practical means for introducing science-driven navigational autonomy into rover mission plans. The result is the first ever fully-autonomous geologic maps of surface material, and the first reported science interpretation of remote sensing data by a field robot. Our approach permits adaptive kilometer-scale site surveys in a sub-hour time frame.

Finally, we have used information-driven learning in spatial models as a unified conceptual framework to explain a number of existing exploration heuristics. *Proxy features* permit principled adaptive data return based on noisy observations of hidden science phenomena. The result, depending on prior assumptions, ranges from periodic to targeted sampling. Proxy features are a principled, robust, and generalizable method for introducing autonomy into selective sampling and return decisions. They provide a means of making summary and downlink decisions without direct detection of science content.

## 6.2 Future Work

There are several promising avenues for future work in feature detection and classification. Geologic feature detection could still be improved. There is a need for formal validation efforts involving planetary scientists, especially with regards to image features such as morphology for which no “gold standard” datasets currently exist. Richer models of surface morphology could employ non-linear dimensionality reduction or explicit segmentation of images into different surface types. Classification decisions could incorporate multiple sensors as in the work of Pedersen [25]. Here the option to forgo comprehensive measurements becomes increasingly important; many informative sensors such as thermal-infrared spectrometers suffer from long integration times.

Feature classification offers a chance to incorporate more specialized geologic knowledge, both in the interpretation of spectroscopic data and the analysis of images. In this work we rely exclusively on unsupervised data interpretation using dimensionality reduction. This benefits robustness because the robot is capable of uncovering the principal distinctions in whatever data it collects

without requiring any close matches with a training set. Alternatively, supervised classifications could direct the investigation toward particular distinctions that the scientists find most interesting. Pedersen has explored the use of geologically-relevant classifications such as those pertaining to mineralogical type. If example data products are available from the site, and instrument calibration is consistent across multiple command cycles, one could perform autonomous geologic classification using deconvolution with onboard spectral libraries or geologic feature analysis with supervised classifiers.

Additional work in spatial modeling could address a practical “weak link” in our approach: the assumption that the science question of interest concerns a phenomenon represented by a single scalar value at each observation site. While this assumption serves a wide range of site survey tasks, scientists might find it easier to define science goals in terms of ideal “representative subsets” that they would like to receive. These cases favor a subset learning approach similar to that of DesJardins *et al* [96]. One could learn covariance relationships in a similar fashion without the use of explicitly-labeled training data. Under the assumption that the scientists themselves were performing information-driven subsampling, these subsets constitute training data for comparing the feasibility of different covariance relationships.

In the long run more sophisticated geologic models might permit studies that go beyond superficial mapping to chart the positions of true geologic units. These analyses would consider what might lie *under* the ground in addition to the surface material immediately visible to the robot. The kernel-based smoothing offered by Gaussian processes alone would not be sufficient; such investigations would require models capable of representing the ages and formation processes of layered subsurface structures. One could augment spatial inference on the surface with symbolic inference to determine overlaps and other vertical relationships. These might suggest the relative ages of the different materials.

There are several potential changes in the action selection routines that future work might consider. First, it might be advantageous to formulate path cost in terms not a cost per unit time rather than a hard “end-of-day” position constraint. With this alternative formulation the utility of a path would be equal to its associated entropy minus a coefficient-weighted term proportional to the time cost of the path [106]. This would encourage the rover to use less than its entire discretionary planning budget for portions of the traverse corridor where entropy was low. An explicit time cost reflects the expected value of extra time remaining when the rover reaches the end of the discretionary path. It would function as a heuristic to mitigate the myopic nature of the path planner.

It would also be useful to introduce some notion of navigability into the high-level path planning routine. The current science planning module is free to choose waypoints that are known to be unreachable by lower-level navigation components. A simple solution might ascribe a variable time cost to each path based on the navigability of the terrain.

Finally, there is wide scope for expanding the intelligent mapping strategy to other platforms and exploration environments. One-dimensional models such as time series are relevant to any temporal subsampling application, such as spacecraft orbital flybys, lander image sequences, or engineering imagery from a specific spacecraft maneuver. Two-dimensional spatial models have

obvious applicability for surface exploration with mobile platforms. These include rovers such as the Mars exploration scenario from the introduction, but also aerial remote sensing from balloons, airships, or gliders. Airborne platforms also hold out the possibility of atmospheric sampling leading to three-dimensional maps of dynamic environments. The Gaussian process model generalizes readily to these cases with extra dimensions to account for altitude and time.

Remote sensing data will precede surface exploration, and it is important that autonomous agents utilize this data effectively. Automatic onboard image analysis has already benefited terrestrial science in the detection and classification of transient phenomena [41]. Intelligent mapping is a useful approach for observation planning in this domain as well. Often an orbiter might provide low-resolution context imagery along with an occasional high-resolution, small-footprint sensor or image. Because the orbiter might not return to the same location in a reasonable time, remote spacecraft could improve data return by analyzing the context image and deciding where to target their high-resolution instruments. Information gain with respect to a spatial model is a principled criterion for resolving these decisions.

### 6.3 Maps and Exploration

Maps and science have always had a close relationship. Cartographers of the first half of the 20th century generally believed that the continuing refinement of maps paralleled and supported a relentless advance of science toward ever-more-accurate representations of reality.

More recently, scholars have questioned the role of maps as purely passive records of data [5, 2]. Maps identify the objects of an investigation, exclude others, and define the relevant relationships between these objects that the scientist should consider. In short, the map presupposes specific procedures for data acquisition and analysis. It is embodied within a social network of scientists and conventions that control its interpretation and evaluate its authority. The postmodernist philosophers of science often emphasize the oppressive aspects of such artifacts: the ways in which they restrict discourse or unjustly favor the institutional status quo.

The tendency of maps to define scientific study is especially evident in remote exploration. Our discussion of intelligent mapping has highlighted the significant influence of the model definition on sampling and transmission behavior. But such tools need not be inherently restrictive; they are also a means for improving the behavior and efficiency of remote exploration systems. The model definition constitutes a code through which the scientist communicates appropriate behavior to the agent. It allows the scientist to specify exactly what they would like the agent to do in terms that are intimately connected to the data and the physical environment in which the agent operates. In this sense the map is fundamentally an enabling technology.

Planetary exploration is entering a new era in which mission capabilities are constrained not simply by our ability to deploy physical sensors remotely, but also by time and bandwidth resources to gather data during the vehicle lifetime. Long range *over-the-horizon* travel by surface rovers is a poignant example. However, the implications of autonomous science extend far beyond surface rovers. In future years the Deep Space Network will be forced to service more spacecraft, with longer

lifetimes and higher-resolution instruments, that visit ever-farther corners of the solar system.

Mapmaking offers a general criterion for first-pass survey. Scientists may not anticipate the most interesting features of over-the-horizon terrain. By discovering structure in data the autonomous mapmaker identifies these unexpected novelties and trends. Adaptive mapmaking provides a natural framework for determining science value by combining prior expectations with discovered trends. This will be increasingly important as we explore the outer bodies of the solar system, where our advance understanding is weak even as communications constraints are more severe.



# Bibliography

- [1] T. S. Kuhn, *The Structure of Scientific Revolutions*. Chicago: University of Chicago Press, 2 ed., 1970.
- [2] D. Turnbull, *Maps are Territories: Science is an Atlas*. Chicago: University of Chicago Press, 1989.
- [3] R. O. Kuzmin, R. Greeley, R. Landheim, N. A. Cabrol, and J. D. Farmer, "Geologic map of the mtm -15182 and mtm -15187 quadrangles, eusev crater- ma'adim vallis region, mars," *U.S. Geological Survey Geologic Investigations Series I-2666*, 2000.
- [4] A. F. C. Haldemann and M. P. Golombek, "Mars pathfinder near-field rock distribution re-evaluation," in *Lunar and Planetary Sciences Conference*, 2003.
- [5] D. Turnbull, *Cartography and Science in Early Modern Europe: Mapping the Construction of Knowledge Spaces*, vol. 48. 1996.
- [6] T. W. Fong, M. Allan, X. Bouyssounouse, M. G. Bualat, M. Deans, L. Edwards, L. Fleckiger, L. Keely, S. Y. Lee, D. Lees, V. To, and H. Utz, "Robotic site survey at haughton crater," in *9th International Symposium on Artificial Intelligence, Robotics, and Automation in Space*, February 2008.
- [7] D. Wettergreen et al., "Second experiments in the robotic investigation of life in the atacama desert of chile," in *ISAIRAS*, 9 2005.
- [8] D. W. et al, "Long-distance autonomous survey and mapping in the robotic investigation of life in the atacama desert," in *8th International Symposium on Artificial Intelligence, Robotics and Automation in Space*, 2008.
- [9] V. C. Gulick, R. L. Morris, M. A. Ruzon, and T. L. Roush, "Autonomous Image Analysis During the 1999 Marsokhod Rover Field Test," *J. Geophysical Research*, vol. 106, no. E4, pp. 7745–7764, 2001.
- [10] R. Castaño, R. C. Anderson, T. Estlin, D. DeCoste, F. Fisher, D. Gains, D. Mazzoni, and M. Judd, "Rover Traverse Science for Increased Mission Science Return," in *IEEE Aerospace*, March 2005.
- [11] Robert J. Cesarone and Douglas S. Abraham and Leslie J. Deutsch, "Prospects for a Next-Generation Deep Space Network," *Proceedings of the IEEE*, vol. 95:10, pp. 1902–1915, October 2007.
- [12] R. Castaño, T. Estlin, D. Gaines, C. Chouinard, B. Bornstein, R. C. Anderson, M. Burl, D. R. Thompson, A. CAstaño, and M. Judd, "Onboard autonomous rover science," in *IEEE Aerospace Conference Proceedings, Big Sky, Montana*, March 2007.
- [13] B. Cox, "Updated avionics electronic trends," Tech. Rep. OPFMS:033, 2007.

- [14] L. Pedersen, M. Deans, D. Lees, S. Rajagoplan, and D. E. Smith, "Multiple target single cycle instrument placement," *ISAIRAS*, 2005.
- [15] L. Pedersen, W. J. Clancey, M. Sierhuis, N. Muscettola, D. E. Smith, D. Lees, K. Rajan, S. Ramakrishnan, P. Tompkins, A. Vera, and T. Dayton, "Field demonstration of surface human-robotic exploration activity," 2006.
- [16] P. McGuire, J. Ormö, E. Martínez, J. Manfredi, J. Elvira, H. Ritter, M. Oesker, and J. Ontrup, "The Cyborg Astrobiologist: first field experience," *International Journal of Astrobiology*, vol. 3, no. 03, pp. 189–207, 2005.
- [17] M. J. Woods, A. Shaw, E. Honary, D. P. B. D. Long, and D. Pullan, "CREST Autonomous Robotic Scientist: Developing a Closed-Loop Science Capability for European Mars Missions," 2008.
- [18] T. Smith, D. R. Thompson, D. Wettergreen, N. Cabrol, K. Warren-Rhodes, and S. Weinstein, "Life in the Atacama: Autonomous Science for Improved Data Quality," *Journal of Geophysics Research*, vol. 112:G04S03, 2007.
- [19] N. A. Cressie, *Statistics for Spatial Data*. Wiley and Sons, New York, 1991.
- [20] C. Williams and C. Rasmussen, "Gaussian Processes for Regression," *Advances in Neural Information Processing Systems*, vol. 8, pp. 514–520, 1996.
- [21] F. Bacon, "Preparative toward a Natural and Experimental History," 1620.
- [22] V. Gulick, R. Moris, M. Ruzon, and T. Roush, "Autonomous science analyses of digital images for mars sample return and beyond," *Lunar and planetary science XXX*, 1999.
- [23] H. Dunlop, D. R. Thompson, and D. Wettergreen, *Multi-scale Features for Detection and Segmentation of Rocks in Mars Imagery*. 2007.
- [24] D. R. Thompson and R. Castaño, "A performance comparison of rock detection algorithms for autonomous geology," in *IEEE Aerospace Conference Proceedings, Big Sky Montana*, March 2007 [in press].
- [25] L. Pedersen, *Robotic Rock Classification and Autonomous Exploration*. PhD thesis, Robotics Institute, Carnegie Mellon University, Pittsburgh, PA, November 2000.
- [26] M. D. Wagner, D. Apostolopoulos, K. Shillcutt, B. Shamah, R. G. Simmons, and W. Whittaker, "The science autonomy system of the nomad robot," in *ICRA*, vol. 2, pp. 1742–1749, 2001.
- [27] J. Fox, R. Castaño, and R. C. Anderson, "Onboard autonomous rock shape analysis for mars rovers," in *IEEE Aerospace Conference Proceedings, Big Sky, Montana*, March 2002.
- [28] L. Pedersen, "Science target assessment for mars rover instrument deployment," in *IEEE International Conference on Intelligent Robotics and Systems*, vol. 1, pp. 817–822, 2002.
- [29] R. Castaño, "Personal communication," 2006.
- [30] T. Estlin, D. Gaines, I. Nesnas, R. Castano, B. Bornstein, M. Judd, and R. C. Anderson 2008.
- [31] D. C. J. MacKay, *Information Theory, Inference, and Learning Algorithms*. Cambridge, UK: Cambridge University Press, 2003.

- [32] D. R. Thompson, S. Niekum, T. Smith, and D. Wettergreen, "Automatic detection and classification of geological features of interest," in *IEEE Aerospace Conference Proceedings, Big Sky Montana*, March 2005.
- [33] D. R. Thompson, T. Smith, and D. Wettergreen, "Autonomous detection of novel biologic and geologic features in atacama desert rover imagery," in *Lunar and Planetary Sciences Conference*, March 2006.
- [34] R. A. Craddock, M. Golombek, and A. D. Howard, "Analyses of rock size-frequency distributions and morphometry of modified hawaiian lava flows: Implications for future martian landing sites," 2000.
- [35] J. F. B. I. et. al, "Pancam multispectral imaging results from the spirit rover at gusev crater," *Science*, vol. 305, pp. 800–806, 2004.
- [36] J. A. Grant, R. Arvidson, J. F. B. III, N. A. Cabrol, and M. H. C. et al., "Surficial deposits at gusev crater along spirit rover traverses," *Science*, vol. 305, pp. 805–810, 2004.
- [37] M. B. G. et. al., "Geology of the gusev cratered plains from the spirit rover transverse," *Journal of Geophysical Research*, vol. 111 EO2S07, 2006.
- [38] H. Dunlop, *Automatic Rock Detection and Classification in Natural Scenes*, vol. CMU-RI-TR-06-40. August 2006.
- [39] R. S. Crofts, "A visual measure of shingle particle form for use in the field," *Journal of Sedimentary Petrology*, vol. 44, pp. 931–934, 1974.
- [40] B. Bornstein and R. Castaño, "Creation and testing of an artificial neural network based carbonate detector for mars rovers," in *IEEE Aerospace Conference Proceedings, Big Sky, Montana*, March 2005.
- [41] S. Chien, R. Sherwood, D. Tran, B. Cichy, C. R. G., R. Castaño, A. Davies, D. Mandl, S. Frye, B. Trout, *et al.*, "Using Autonomy Flight Software to Improve Science Return on Earth Observing One," *Journal of Aerospace Computing, Information, and Communication*, vol. 2, pp. 196–216, 2005.
- [42] P. Viola and M. Jones, "Rapid object detection using a boosted cascade of simple features," *Proc. CVPR*, vol. 1, pp. 511–518, 2001.
- [43] G. Bradski, A. Kaelher, and V. Pisarevsky, "Learning-Based Computer Vision with Intel's Open Source Computer Vision Library," vol. 9, pp. 119–130, May 2005.
- [44] P. Viola and M. Jones, "Robust real-time object detection," *International Journal of Computer Vision*, vol. 57, no. 2, pp. 137–154, 2002.
- [45] Y. Freund and R. E. Schapire, "Experiments with a new boosting algorithm," in *International Conference on Machine Learning*, pp. 148–156, 1996.
- [46] M. Deans, D. Wettergreen, and D. Villa, "A Sun Tracker for Planetary Analog Rovers," *The 8th International Symposium on Artificial Intelligence, Robotics and Automation in Space*, p. 76, 2005.
- [47] P. Backes et al., "Automated rover positioning and instrument placement," *IEEE Aerospace*, 2005.
- [48] e. a. M. Bajracharya, "Target tracking, approach, and camera handoff for automated instrument placement," *IEEE Aerospace*, 2005.

- [49] R. Madison, "Improved target handoff for single cycle instrument placement," *IEEE Aerospace*, 2006.
- [50] D. Lowe, "Distinctive Image Features from Scale-Invariant Keypoints," *International Journal of Computer Vision*, vol. 60, no. 2, pp. 91–110, 2004.
- [51] N. A. C. et. al, "Life in the Atacama: Searching for Life with Rovers," *Journal of Geophysics Research*, vol. 112:G04S02, 2007.
- [52] S. Singh, R. Simmons, T. Smith, A. Stentz, V. Verma, A. Yahja, and K. Schwehr, "Recent Progress in Local and Global Traversability for Planetary Rovers," in *ICRA*, 2000.
- [53] R. Hartley and A. Zisserman, *Multiple View Geometry in Computer Vision*. Cambridge Univ. P., 2nd ed. ed., 2004.
- [54] F. C. Peralta, D. R. Thompson, and D. Wettergreen, "Autonomous Rover Reflectance Spectroscopy with Dozens of Targets," 2008.
- [55] S. A. Hutchinson, G. D. Hager, and P. I. Corke, "A tutorial on visual servo control," *IEEE Trans. Robotics and Automation*, vol. 12, pp. 651–670, Oct, 1996.
- [56] Mustard, J.F. and C. M. Pieters and S. F. Pratt, "Deconvolution of spectra for intimate mixtures," in *LPSC XVII*, 1986.
- [57] P. J. Diggle, *Statistical Analysis of Spatial Point Patterns*. London: Arnold Publishers, 2003.
- [58] C. O. Sauer, "The education of a geographer," *Annals of the Association of American Geographers*, vol. XLVI:3, 1950.
- [59] "Learning metric-topological maps for indoor mobile robot navigation," *Artificial Intelligence*, vol. 99:1, pp. 21–71, February 1998.
- [60] S. Moorehead, R. Simmons, and W. Whittaker, "A multiple information source planner for autonomous planetary exploration," in *ISAIRAS*, 2001.
- [61] N. Fairfield, G. A. Kantor, and D. Wettergreen, "Real-time slam with octree evidence grids for exploration in underwater tunnels," *Journal of Field Robotics*, 2007.
- [62] K. Warren-Rohodes, S. Weinstein, J. Dohm, J. Piatek, E. Minkley, A. Hock, C. Cockell, D. Pane, L. A. Ernst, G. Fisher, S. Emani, A. S. Waggoner, N. A. Cabrol, D. S. Wettergreen, D. Apostopolous, P. Coppin, E. Grin, G. C. Diaz, J. Moersh, G. G. Oril, T. Smith, K. Stubbs, G. Thomas, M. Wagner, and M. Wyatt, "Searching for microbial life remotely: Satellite-to-rover habitat mapping in the Atacama Desert, Chile," *Journal of Geophysics Research*, vol. 112:G04S05, 2007.
- [63] B. Kuipers and Y. T. Byun, "A robot exploration and mapping strategy based on a semantic hierarchy of spatial representations," 1993.
- [64] L. Wasserman, *All of Nonparametric Statistics*. Springer, 2006.
- [65] L. Pedersen, "Autonomous characterization of unknown environments," in *2001 IEEE International Conference on Robotics and Automation*, vol. 1, pp. 277 – 284, May 2001.
- [66] C. E. Rasmussen and C. K. I. Williams, *Gaussian Processes for Machine Learning*. MIT Press, 2006.
- [67] M. Seeger, "Gaussian processes for machine learning," *International Journal of Neural Systems*, vol. 14:2, pp. 69–106, 2004.

- [68] B. Scholkopf, R. Herbrich, and A. Smola, "A generalized representer theorem," *Proceedings of the Annual Conference on Computational Learning Theory*, pp. 416–426, 2001.
- [69] C. M. Bishop, *Neural Networks for Pattern Recognition*. Oxford University Press, 1995.
- [70] C. Paciorek and M. Schervish, "Nonstationary Covariance Functions for Gaussian Process Regression," *Advances in Neural Information Processing Systems*, vol. 16, pp. 273–280, 2004.
- [71] P. D. Sampson and P. Guttorp, "Nonparametric Estimation of Nonstationary Spatial Covariance Structure," *Journal of the American Statistical Association*, vol. 87, no. 147, 1992.
- [72] T. Pfingsten, M. Kuss, and C. E. Rasmussen, "Nonstationary Gaussian Process Regression using a Latent Extension of the Input Space," 2006.
- [73] R. Greeley and J. D. Iversen, "Field Guide to Amboy Lava Flow, San Bernardino County, California," 1978.
- [74] D. A. Christian, D. Wettergreen, M. Bualat, K. Schwehr, D. Tucker, and E. Zbinden, "Field experiments with the ames marsokhod rover," in *Proceedings of the 1997 Field and Service Robotics Conference*, December 1997.
- [75] A. Winterholler, M. Roma, T. Hunt, and D. Miller, "Design of a High-Mobility Low-Weight Lunar Rover," *i-SAIRAS 2005'-The 8th International Symposium on Artificial Intelligence, Robotics and Automation in Space*. Edited by B. Battrick. ESA SP-603. European Space Agency, 2005. Published on CDROM., p. 4.1, 2005.
- [76] A. W. Hatheway, *Lava Tubes and Collapse Depressions*. PhD thesis, 1971. Amboy map reprinted in Greeley 1978.
- [77] United States Geological Survey, *Digital Orthophoto Quadrangles Fact Sheet 057-01*.
- [78] A. S. McEwen, W. A. Delamere, E. M. Eliason, J. A. Grant, V. C. Gulick, C. J. Hansen, K. E. Herkenhoff, L. Keszthelyi, R. L. Kirk, M. T. Mellon, S. W. Squyres, N. Thomas, and C. Weitz, "Hirise: The high resolution imaging science experiment for mars reconnaissance orbiter," in *Lunar and Planetary Institute Conference Abstracts*, vol. 33, p. 1163, Mar 2002.
- [79] K. Wagstaff, R. C. no, S. Molinar, and M. Kilmesh, "Science-based region-of-interest image compression," in *LPSC*, 2004.
- [80] C. W. Ko, J. Lee, and M. Queyranne, "An Exact Algorithm for Maximum Entropy Sampling," *Operations Research*, vol. 43, no. 4, pp. 684–691, 1995.
- [81] R. Castaño, K. Wagstaff, L. Song, and R. C. Anderson, "Validating rover image prioritizations," tech. rep., 2005.
- [82] S. Squyres, J. Grotzinger, R. Arvidson, J. Bell, W. Calvin, P. Christensen, B. Clark, J. Crisp, W. Farrand, K. Herkenhoff, *et al.*, "In Situ Evidence for an Ancient Aqueous Environment at Meridiani Planum, Mars," *Science*, vol. 306, no. 5702, pp. 1709–1714, 2004.
- [83] S. C. R. C. B. B. A. F. A. C. J. B. R. G. P. Whelleyand and M. Lemmon, "Results from Automated Cloud and Dust Devil Detection ONboard the MER Rovers," in *ISAIRAS*, 2008.
- [84] T. Smith, D. R. Thompson, S. Weinstein, and D. Wettergreen, "Autonomous rover detection and response applied to the search for life via chlorophyll fluorescence in the atacama desert," March 2006.

- [85] D. Apostolopoulos, L. Pedersen, B. Shamah, K. Shillcutt, M. Wagner, and W. Whittaker, "Robotic Antarctic Meteorite Search: Outcomes," *IEEE International Conference on Robotics and Automation*, vol. 4, 2001.
- [86] T. Smith, S. Niekum, D. R. Thompson, and D. Wettergreen, "Concepts for science autonomy during robotic traverse and survey," March 2005.
- [87] R. Neapolitan and K. Naimipour, *Foundations of Algorithms using C++ Pseudocode*. London: Jones and Barnett, second edition ed., 1998.
- [88] C. S., S. R., T. D., C. R., B. Cichy, A. Davies, G. Rabideau, N. Tang, M. Burl, D. Mandl, *et al.*, "Autonomous Science on the EO-1 Mission," *International Symposium on Artificial Intelligence, Robotics, and Automation in Space (i-SAIRAS 2003)*. Nara, Japan. May, 2003.
- [89] S. Chien, R. Sherwood, D. Tran, B. Cichy, G. Rabideau, R. Castaño, A. Davies, D. Mandl, S. Frye, B. Trout, *et al.*, "Lessons learned from autonomous sciencecraft experiment," *Proceedings of the fourth international joint conference on Autonomous agents and multiagent systems*, pp. 11–18, 2005.
- [90] A. Castaño, A. Fukunaga, J. Biesiadecki, L. Neakrase, P. Whelley, R. Greeley, M. Lemmon, R. Castaño, and S. Chien, "Automatic detection of dust devils and clouds on Mars," 2007.
- [91] J. Bresina, K. Golden, D. Smith, and R. Washington, "Increased flexibility and robustness of mars rovers," (Noordwijk, The Netherlands), June 1999.
- [92] D. E. Smith, "Choosing objectives in oversubscription planning," 2004.
- [93] T. Smith and R. Simmons, "Heuristic search value iteration for pomdps," *Proceedings of the 20th conference on Uncertainty in artificial intelligence*, pp. 520–527, 2004.
- [94] M. Markou and S. Singh, "Novelty detection: a reviewpart 1: statistical approaches," *Signal Processing*, vol. 83, no. 12, pp. 2481–2497, 2003.
- [95] C. Campbell and K. Bennett, "A Linear Programming Approach to Novelty Detection," *Advances in Neural Information Processing Systems 13*, 2001.
- [96] M. desJardins, E. Eaton, and K. L. Wagstaff, "Learning user preferences for sets of objects," in *International Conference on Machine Learning*, 2006.
- [97] C. E. Shannon, "A mathematical theory of communication," in *The Bell System Technical Journal*, vol. 27, pp. 379–423, July 1948.
- [98] K. Chaloner and I. Verdinelli, "Bayesian Experimental Design: A Review," *Statistical Science*, vol. 10, no. 3, pp. 273–304, 1995.
- [99] D. Lindley, "On the measure of information gain provided by an experiment," vol. 27, pp. 199–218, 1956.
- [100] D. Lindley, *Bayesian statistics*. Society for Industrial and Applied Mathematics Philadelphia, 1972.
- [101] K. Ryan, "Estimating Expected Information Gains for Experimental Designs With Application to the Random Fatigue-Limit Model," *Journal of Computational & Graphical Statistics*, vol. 12, no. 3, pp. 585–603, 2003.
- [102] L. Pedersen, M. D. Wagner, D. Apostolopoulos, and W. R. L. Whittaker, "Autonomous robotic meteorite identification in antarctica," in *ICRA*, May 2001.

- [103] R. Simmons, D. Apfelbaum, W. Burgard, M. Fox, D. an Moors, S. Thrun, and H. Younes, “Coordination for multi-robot exploration and mapping,” *AAAI*, 2000.
- [104] R. Zlot, A. Stentz, M. Dias, and S. Thayer, “Multi-robot exploration controlled by a market economy,” vol. 3, 2002.
- [105] A. Makarenko, S. Williams, F. Bourgault, and H. Durrant-Whyte, “An experiment in integrated exploration,” pp. 534–539, 2002.
- [106] S. Moorehead, *Autonomous Surface Exploration for Mobile Robots*. PhD thesis, Robotics Institute, Carnegie Mellon University, Pittsburgh, PA, August 2001.
- [107] C. Stachniss, G. Grisetti, and W. Burgard, “Information gain-based exploration using rao-blackwellized particle filters,” in *Proceedings of Robotics: Science and Systems*, (Cambridge, MA, USA), 2005.
- [108] T. J. Loredo and D. F. Chernoff, “Bayesian Adaptive Exploration,” *Statistical Challenges in Astronomy*, pp. 275–97, 2002.
- [109] M. Shewry and H. Wynn, “Maximum Entropy Sampling,” *J. Applied Statistics*, vol. 14, no. 2, 1987.
- [110] J. Lee, “Constrained maximum-entropy sampling,” vol. 46:5, pp. 655–664, Sep. - Oct. 1998.
- [111] P. Sebastiani and H. Wynn, “Experimental design to maximize information,” *AIP Conference Proceedings Workshop on Bayesian Inference and Maximum Entropy Methods in Science and Engineering*, vol. 568, pp. 192–203, May 2001.
- [112] P. Sebastiani and H. P. Wynn, “Maximum Entropy Sampling and Optimal Bayesian Experimental Design,” in *Journal of the Royal Statistical Society*, 2000.
- [113] C. Guestrin, A. Krause, and A. Singh, “Near-optimal Sensor Placements in Gaussian Processes,” School of Computer Science, Carnegie Mellon University, 2005.
- [114] D. MacKay, “Information-Based Objective Functions for Active Data Selection,” *Neural Computation*, vol. 4, no. 4, pp. 590–604, 1992.
- [115] T. Loredo, “Bayesian adaptive exploration in a nutshell,” *PHYSTAT*, 2003.
- [116] T. Loredo, “Bayesian Adaptive Exploration,” *23 rd International Workshop on Bayesian Inference and Maximum Entropy Methods in Science and Engineering*, pp. 330–346, 2003.
- [117] V. V. Fedorov, *Theory of Optimal Experiments*. academic press, 1972.
- [118] A. Krause and C. Guestrin., “Nonmyopic Active Learning of Gaussian Processes — an Exploration-Exploitation Approach,” *International Conference on Machine Learning*, 2007.
- [119] C. Chekuri and M. Pal, “A Recursive Greedy Algorithm for Walks on Directed Graphs,” in *IEEE Symposium on the Foundations of Computer Science*, 2005.
- [120] A. Singh, A. Krause, C. Guestrin, W. Kaiser, and M. Batlin, “Efficient Planning of Informative Paths for Multiple Robots,” in *IJCAI*, 2007.
- [121] P. E. Arvidson and S. W. Squyres, “Recent Results from the Mars Exploration Rover Opportunity Mission,” *American Geophysical Union, Spring Meeting, abstract P31A-02*, May 2005.

- [122] Trey Smith and David R. Thompson and David Wettergreen, "Generating exponentially smaller pomdp models using conditionally irrelevant variable abstraction," *ICAPS*, 2007.
- [123] A. Meliou, A. Krause, C. Guestrin, and J. M. Hellerstein, "Nonmyopic informative path planning in spatio-temporal models," Tech. Rep. UCB/EECS-2007-44, UC Berkeley.
- [124] R. Frost, *Closed for Good*. New York: Spiral Press, 1948.
- [125] N. A. Cabrol, J. D. Farmer, E. A. Grin, L. Richter, L. Soderblom, R. Li, K. Herkenhoff, G. A. Landis, and R. E. Arvidson, "Aqueous processes at Gusev crater inferred from physical properties of rocks and soils along the Spirit traverse," *Journal of Geophysics Research*, vol. 111:E02S20, 2006.
- [126] G. K. Wallace, "The jpeg still picture compression standard," *IEEE Transactions on Consumer Electronics*, vol. 38:1, 1992.
- [127] T. M. Cover and J. A. Thomas. Wiley and Sons, New York, 1991.
- [128] A. Schwaighofer, V. Tresp, and K. Yu, "Learning Gaussian Process Kernels via Hierarchical Bayes," *Advances in Neural Information Processing Systems*, 2004.
- [129] D. R. Thompson, T. Smith, and D. Wettergreen, "Data mining during rover traverse: From images to geologic signatures," in *8th International Symposium on Artificial Intelligence, Robotics and Automation in Space*, September 2005.
- [130] M. Varma and A. Zisserman, "A Statistical Approach to Texture Classification from Single Images," in *International Journal of Computer Vision: Special Issue on Texture Analysis and Synthesis*, vol. 62:1-2, pp. 61–81, 2005.
- [131] R. L. Rabiner, "A Tutorial on Hidden Markov Models and Selected Applications in Speech Recognition," *Proceedings of the IEEE*, vol. 77:2, 1989.
- [132] L. Wasserman, *All of Statistics: A Concise Course in Statistical Inference*. Springer, 2004.
- [133] A. Krause and C. Guestrin., "Optimal Value of Information in Graphical Models - Efficient Algorithms and Theoretical Limits," *Proc. of the International Joint Conference on Artificial Intelligence (IJCAI)*, 2005.
- [134] C.-C. Chang and C.-J. Lin, *LIBSVM: A Library for Support Vector Machines*, 2001.
- [135] Y. Yamaguchi, A. B. Kahle, H. Tsu, T. Kawakami, and M. Pniel, "Overview of advanced spaceborne thermal emission and reflection radiometer (aster)," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 36:4, 1998.
- [136] G. Forney Jr, "The Viterbi Algorithm," *Proceedings of the IEEE*, vol. 61, no. 3, pp. 268–278, 1973.