Adaptive Spectroscopic Exploration Driven by Science Hypotheses for Geologic Mapping

Alberto Candela Garza

CMU-RI-TR-17-62

August 2017

Robotics Institute
School of Computer Science
Carnegie Mellon University
Pittsburgh, PA 15213

Thesis Committee:

David Wettergreen, Chair
Nathan Michael
Daniel Maturana

Submitted in partial fulfillment of the requirements for the degree of Master of Science.

Abstract

As exploration goes to further extremes, communication becomes more constrained. Planetary rover operations are limited in bandwidth, delayed due to distance and resource restricted to just a few communication cycles per day. Despite being one of the closest planets from Earth, Mars rovers spend the vast majority of their time isolated, awaiting instructions. Clearly the exploration of much farther celestial bodies, such as the moons of Jupiter or Saturn, will be even more difficult. At the same time, planetary exploration involves frequent scientific reformulation and replanning. This problem suggests the need for a new paradigm where rovers can communicate and operate more efficiently by having a deeper understanding of the evolving scientific goals and hypotheses guiding their missions, rather than just collecting and sending data to human scientists for interpretation and planning.

This work establishes the science hypothesis map as a spatial probabilistic structure in which scientists initially describe their abstract beliefs and hypotheses, and in which the state of this belief evolves as the robot makes raw measurements. It discusses how to incorporate path planning for maximizing scientific information gain, which is efficiently computed. As proof of concept, this thesis describes a geologic exploration problem where a robot uses a spectrometer to infer the geologic composition of regions. This research shows that the science hypothesis map can be used to infer geologic units with high accuracy, and that exploration using information gain-based path planning has better performance than exploration with conventional science-blind algorithms.

Acknowledgments

Both my master's studies and this thesis were only possible through the help of many others, so I would like to dedicate a few words to acknowledge some of this support.

First of all, I would like to thank my committee members, Nathan Michael and Daniel Maturana, for taking the time of reading this thesis and providing highly useful comments.

I am very grateful to the Fulbright and CONACYT organizations for their financial support. Their fellowships allowed me to cover most of my educational expenses and pursue this graduate degree. Additionally, this research was supported by the National Science Foundation National Robotics Initiative Grant #IIS-1526667.

I would like to thank Greydon Foil. Although our mutual stay at CMU did not overlap much, his guidance as the senior graduate student in the lab turned out to be incommensurable.

I would also like to thank Gregg Swayze and Eldar Noe Dobrea for their invaluable help during the project, especially in the geologic aspect, as well as for their excellent disposition and sense of humor during our Cuprite field trip.

I owe a great debt of gratitude to David R. Thompson. Working by his side and under his mentoring has been an incomparable learning experience and a true honor; and granting me the opportunity to work at the NASA Jet Propulsion Laboratory twice, has literally been a dream come true.

The person that I probably need to thank the most is my advisor David Wettergreen. From the first time we met, he knew right away that this project would suit my interests. His matchless guidance and insight both on research and in life will be forever appreciated.

Finally, I would like to thank Sharon for her unconditional support, patience, and affection throughout these last couple of years.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

Modern planetary robotic exploration is guided by scientists specifying actions, such as specific instrument measurements, on a path that they believe will best address the investigation questions. The exploration path is formed from expert knowledge of the site and expectations about where to gather mission-critical information. To illustrate this, Figure 1.1 shows an artist's concept of Curiosity collecting instrument measurements on Mars. Scientists reinterpret their measurements with growing contextual knowledge of the environment, so real exploration is characterized by a frequent reformulation and replanning throughout the mission lifetime [30]. Replanning occurs on large strategic scales, bypassing or favoring geographic locales, as well as local tactical scales, lingering at an anomalous feature for additional measurements [21, 56]. However, planetary exploration always occurs with low bandwidth and high latency communications, leaving limited opportunities to revise the exploration plan once committed, or causing scientists to devise an incremental plan in which the robot waits inactive between command cycles.

This thesis describes an approach to overcome the communication bottleneck in robotic exploration, where the command to the remote explorer is based on an evolving model of what the initially scientist believes, rather than a single prescribed route, enabling the robot to take more adaptive and efficient actions based on real-time information, improving the rate and productivity of discovery.

This work develops the *science hypothesis map* as the spatial probabilistic structure in which scientists communicate their belief about the world and in which the belief state evolves as the robotic explorer collects information. This article develops the groundwork for science hypothesis maps, detailing how they can guide the robot to improve the rate and productivity of discovery.

In general, scientists define their exploration questions in terms of abstract concepts far removed from the raw sensor data available to a robot. It is usually unfeasible to encode this rich,

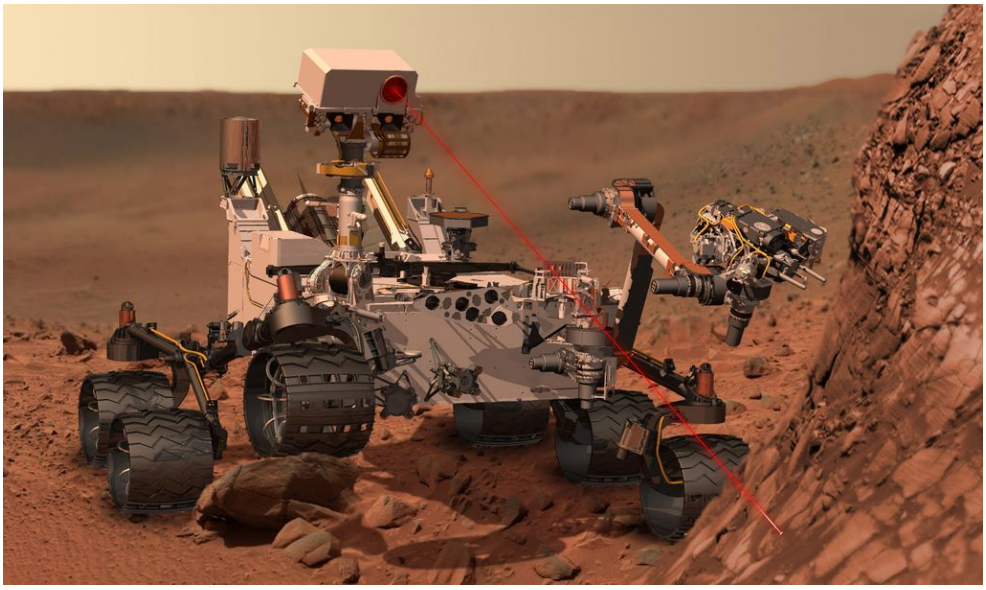

Figure 1.1: This image is an artist's concept of Curiosity, the NASA rover deployed on Mars. It uses its Chemistry and Camera (ChemCam) instrument to investigate the composition of a rock. The ChemCam consists in two different instruments combined as one: a laser-induced breakdown spectroscopy (LIBS), and a Remote Micro Imager (RMI). Image courtesy of NASA [2].

Table 1.1: Examples of simplified Science Traceability Matrices.

Investigation objectives	**Physical properties**	**Measurements**
Determine geology	Abundance of key mineral classes	Surface reflectance in 450-2500 nm
Determine the presence and diversity of marine life	Salinity of the water	Water conductivity in μS/cm
Monitor air quality	Abundance of particles such as carbon monoxide, sulfur dioxide, etc.	Particle concentration in μg/m^3

abstract knowledge in a direct way that enables true robotic understanding. However, it is possible to construct simpler hierarchical probabilistic models relating these representations to measurable data. Conventions for mission design show ways of quantifying these relationships. For example, NASA missions represent the relationship between abstract investigation objectives and raw measurements with a Science Traceability Matrix (STM). This is done through a tripartite division into *investigation objectives*, *physical properties*, and *instrument measurements*. Table 1.1 shows some examples of Science Traceability Matrices for geologic, oceanographic [8], and air quality investigations [4].

The science hypothesis map extends the idea of a Science Traceability Matrix to a probabilistic model with spatial extent, granting robustness under uncertainty. It also determines the

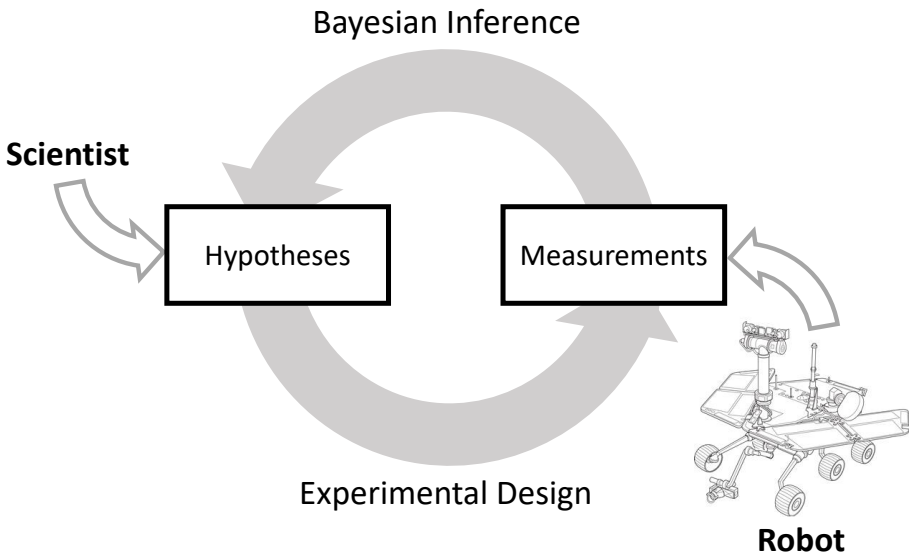

Figure 1.2: Adaptive robotic exploration as a cycle of two processes: Bayesian inference and experimental design (path planning). Scientists can guide the process by defining hypotheses and constraints (in the form of resource models, goal waypoints, etc.).

information that collected measurements, directly interpretable by the robot, provide with respect to the fundamental investigation objective. The most beneficial measurements produce a large expected reduction in uncertainty. The previous notion serves as the foundation of *Bayesian experimental design* [11]. It can be decomposed into two processes that form a cycle: *Bayesian inference* and *experimental design*. In the first one, the hypotheses' probabilities are updated when new measurements are collected. In the second one, the updated hypotheses dictate actions that minimize uncertainty, which translate to *path planning* in robotic exploration (Figure 1.2). With each communication event, the scientist can define the variables of interest and the probabilistic models relating them to instrument data. In real time, a robot can continually calculate navigation plans that optimize scientific discovery.

The presented formulation has several specific benefits. First, science hypothesis maps improve operational efficiency. Current operations require meticulous direction of each robot. Science hypothesis maps communicate objectives simply and intuitively, and define the appropriate behaviors without low-level action planning. Second, they can improve science yield by forcing exploration to be driven by quantitative, formal hypotheses that are related mathematically to the measurements made by the robot. Too often, such hypotheses are left implicit or only heuristically tied to the activities selected. Robotic activity plans can be optimized, allowing the system to react to unanticipated events that occur while the robot is out of communication. For example, delays in navigation are common due to hazard avoidance [56] or unanticipated environmental influences. Robots that understand the hypotheses under study can recover from these effects by

recomputing the measurement strategy. They can exploit resource surpluses in a similar fashion [56].

1.1 Overview of the Thesis

This thesis first describes and discusses work that is related to the alluded problem of autonomous robotic exploration in Chapter 2.

Chapter 3 presents a general description of the science hypothesis map and the way it can be used to perform Bayesian inference of high-level scientific investigation objectives from low-level measurements.

Chapter 4 explains the application of the concepts of the science hypothesis map to a geologic mapping problem of a well-studied site: the Cuprite Mining District in Nevada.

Next, Chapter 5 describes information gain-based path planning strategies that can be integrated with the science hypothesis map to optimize the rate and productivity of discovery.

Chapter 6 first tests the accuracy of the science hypothesis map independently, and then evaluates the performance of diverse science-blind planners vs. science-aware planners; both in simulations with high resolution data of Cuprite, Nevada.

Finally, Chapter 7 concludes this thesis with a discussion of the findings on this work, as well as relevant subsequent avenues for research.

Chapter 2

Related Work

This chapter describes and discusses both classic and state-of-the-art work that is highly related to the tackled problem on this thesis. First, in the area of adaptive robotic exploration. Additionally, spectroscopy is a very important element in planetary and geologic exploration, therefore a whole subchapter is devoted to explaining diverse strategies and techniques for the analysis of spectroscopic data.

2.1 Adaptive Robotic Exploration

Solving the problem of robotic exploration requires multiple steps. First, a model of the environment is typically obtained using prior information from diverse sources and extracting meaningful features. Once a model is learned, a sampling reward is predicted for each feasible sampling point. Following that, a path that seeks to maximize the cumulative reward is planned through the terrain. However, the main difference between conventional exploration and adaptive exploration is the following: in adaptive exploration, the information provided by previously collected measurements can be used to adaptively improve the model, and more importantly, to modify the path itself.

First of all, building a model of the environment typically involves extracting features from imagery collected either through orbital flybys or through rover traversal. Thompson et al. extract rock classes from the rover imagery using a modified Viola/Jones rock classifier [55], which can then be used as features for terrain classification in a spatial model [58]. Classification in the pixel space, using robust machine learning techniques like Random Forests, yields much more precise classification [25].

Once a model is built, the reward calculation proceeds by estimating the gain from each sampling opportunity. Both Foil and Thompson consider iteratively modeling the sample space to

assess the reward of future samples [24, 55]. Additionally, Thompson et al. provide a spatial model that predicts future readings based on previous observations using a Gaussian process reliant on information gain based sampling. Foil et al. models the sample space using modifications to Gaussian mixture model and relies on different metrics to calculate rewards. For instance, [24] uses an adaptive Gaussian mixture model in sample space to assign rewards to points where sampling is feasible. Many methods in exploration robotics assign an entropy or information gain metric to every exploration opportunity and define an objective function which usually serves to maximize the total information gain.

Several approaches exist to generate paths using the estimated reward of a location. Such planning involves formulating an objective function and optimizing it with budget constraints. Maximum-Entropy Sampling (MES) strategies select the next sampling point based on the sole criterion of maximizing the entropy within the set of sampled points. Such a method is therefore restricted by selecting when a single sample is being selected. However, the problem at hand needs to select a subset of all available sampling opportunities. Spatial design is applied by Thompson et al., which consists in selecting samples that maximize the differential entropy of the sample space [55]. On a different approach, Thompson et al. optimize the reconstruction error and analyzes greedy Least Squares and non-negative Least Squares optimization techniques [58].

Robots have become excellent at detecting objects [25, 44, 55], environmental phenomena [21], or finding representative measurement locations by only using low-resolution remote data [58]. To illustrate this, Figure 2.1 shows an example of autonomous exploration carried out by Zoë in the Atacama Desert, Chile. Zoë is a rover that was developed at Carnegie Mellon University (CMU). It is capable of performing onboard science analysis for various tasks such as life detection in extreme environments [30].

However, these methods only work under certain static conditions, goals, and assumptions. Hence the motivation of using evolving science goals to drive adaptive exploration. A few deployed robots can perform automatic science data analysis [8, 15, 21, 62], however they do not use science hypotheses as an integral part of their reasoning. Instead, they pursue static objectives that are fixed at the outset. These simplified tasks, such as mapping scalar fields (e.g., ocean temperature [8]), or detecting transient features (dust devils on Mars [10]), are defined long in advance.

Bayesian networks have been used by some robots for tasks such as mineral classification [26], or meteorite identification [44], but they also operate under predefined static objectives that ignore the evolution of the robot's overall knowledge of a scene throughout the mission.

Researchers have only recently begun to investigate the complete Bayesian experimental de-

Figure 2.1: Example of the CMU rover Zoë navigating autonomously in the Atacama Desert, Chile. Zoë is collecting data with its spectrometer from a location of interest for the science team.

sign loop (Figure 1.2), and consequently build probabilistic models connecting high-level concepts with low-level observations [5, 27]. However, these approaches have not investigated the influence of the scientist's prior knowledge of a scene during exploration, which is critical in realistic scenarios. They have also not integrated nor compared state-of-the-art informative path planners with their systems, which are discussed with more detail in chapter 5.

The main innovation in this work is an approach that incorporates low-level measurements and discoveries during exploration and dynamically responds with evolving high-level objectives and plans.

2.2 Analysis of Spectroscopic Data

Planetary missions tend to heavily rely on the analysis of spectroscopic data. Spectroscopy is a discipline that studies the interaction between matter and electromagnetic radiation. By observing unique patterns in the reflectance and absorption of light throughout the different wavelengths of the electromagnetic spectrum, the chemical composition of an object may be derived, such as the presence of water or certain minerals [17, 39]. There are many different spectroscopic sensors and techniques, both for remote and in situ sensing. In here, we focus on imaging spectroscopy, which is the acquisition of orbital images where each pixel stores information from many wavelengths of the electromagnetic spectrum, instead of just the three bands that comprise the standard RGB color model. Figure 2.2 shows an example of a spectral image, containing

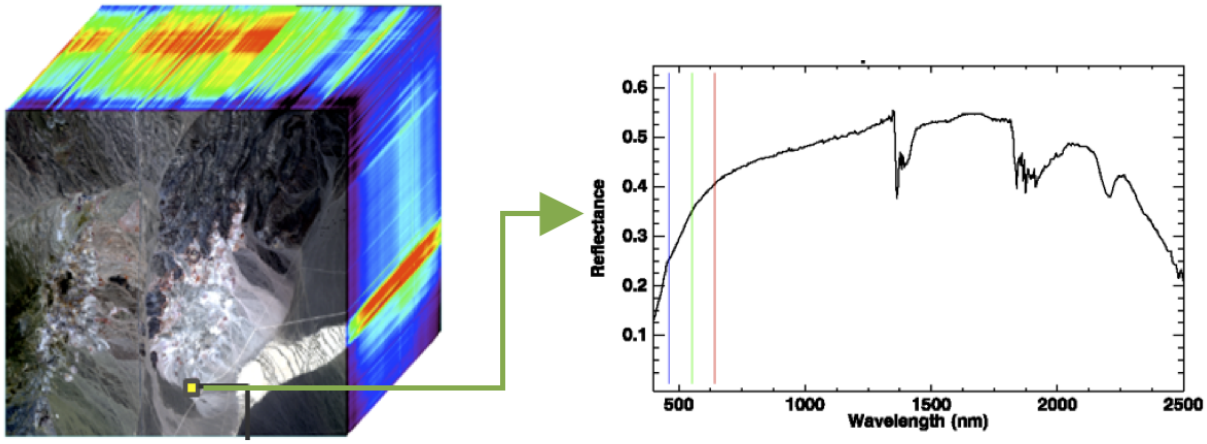

Figure 2.2: Example of a spectral image. 3-D representation (left) and a spectrum corresponding to a single pixel (right). Information provided by imaging spectroscopy is richer and far more complex than the one that is obtained from standard RGB imagery.

much richer information that a conventional RGB image, yet a similar spatial structure.

The analysis of this kind of data is a laborious task and often requires the expertise and intuition of scientists, as well as some knowledge of the specific context in which the spectra was measured. However, the analysis becomes increasingly hard on new planets or environments. Spectral images tend to be high-dimensional large files, so adequate preprocessing and analysis methods are necessary to obtain meaningful and efficient results. There are many different types of spectroscopic analysis methods, depending on the task at hand and the required or desired level of complexity. Figure 2.3 shows a high-level organization of the existing spectral analysis strategies, which are explained next with more detail. Table 2.1 shows a summarized comparison of these strategies.

First of all, a spectral signal is the measured response of the interaction between light and an object, such as a rock. Most of the times, it is difficult to find a truly pure material in the environment, so typically the analysis has to deal with mixtures. Due to simplicity, sparse linear mixing models are the most common for explaining the material abundances in each pixel. However, linear models may be too simplistic in some scenarios. There exist more complex nonlinear models that may be capable of fitting the data with more accuracy, but also have the drawback of potential overfitting [17, 38, 39].

Spectral analysis can be performed by just taking into account relevant features. In this context, geoscientists tend to use the absorption bands of spectra as characteristic features for material identification [17, 35]. Figure 2.4 shows an example of a spectral absorption feature. As it may be observed, it comprises basic elements such as position, depth, and width. This have been

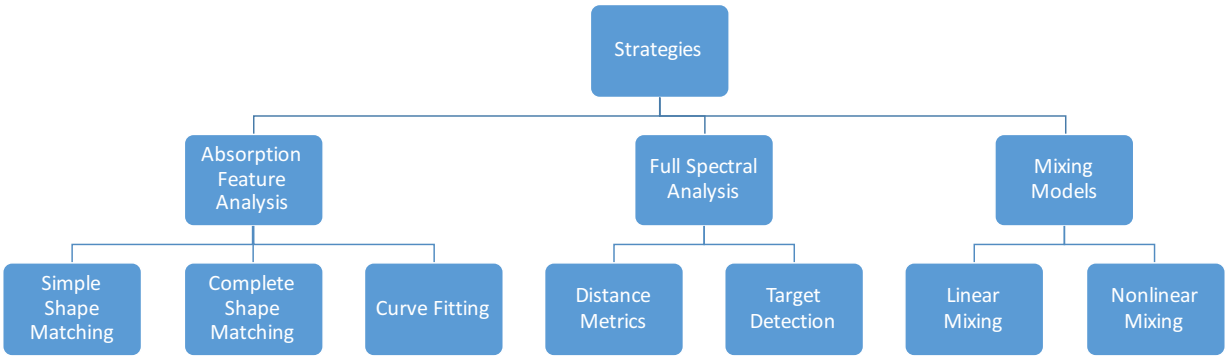

Figure 2.3: Existing spectral analysis strategies, as organized by Rencz and Ryerson [46].

used to construct expert systems for simple shape matching [36]. More sophisticated methods perform a complete shape matching of the absorption features using a least squares criterion in conjunction with encoded expert knowledge, most notably the Tetracorder system [18]. Finally, some techniques attempt to model the features by fitting functions such as Gaussian, Lorentzian, or Voight curves to spectral signals [9]. In general, these methods and no to implement, are highly sensitive to noise, but their results and way of operation are better in terms of scientist interpretability.

Another wide approach for spectral analysis is to look at the whole spectrum and compare it to spectra of known materials. A very popular method consists in comparing spectra using distance metrics. Probably the most common is the Spectral Angle Mapper, which is a direct analog of the cosine distance function or the normalized cross-correlation operation [35]. Other examples of distance metrics within this context are the Spectral Information Divergence (a variation of the Kullblack-Leibler divergence) [12], and the Hamming distance (by encoding spectra as binary signals) [42]. However, these methods are not very useful when dealing with mixtures. There exist target detection techniques for finding present materials of interest, such as the spectral matched filter [39, 48], but they work under the assumption of linear mixtures. They are relatively simple to automate and they are robust robust in the presence of noise, but lose physical interpretability.

In principle, many image processing, machine learning, and computer vision tools can be applied to spectral images. Algorithms ranging from Gaussian Classifiers and K-means [59] to Deep Neural Networks [45] have been applied to these. An important thing to underscore is that spectral images are very large files, therefore dimensionality reduction can be applied

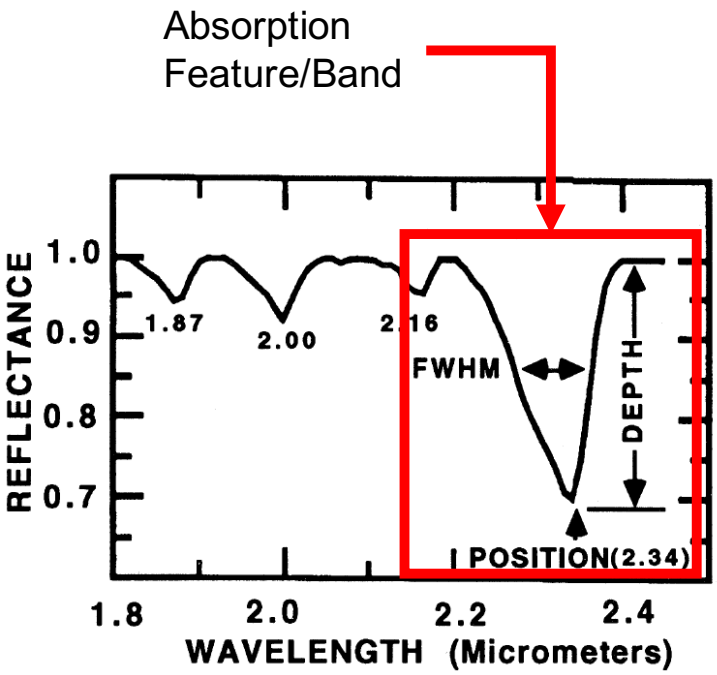

Figure 2.4: Example of a spectral absorption feature [35].

Table 2.1: Summarized comparison of the spectral analysis techniques.

Strategy	**Computational Simplicity**	**Physical Interpretability**	**Data fitting**
Absorption feature analysis	Low	High	Potential overfitting
Full spectral analysis	High	Low	Potential underfitting
Linear mixing	Low	High	Potential underfitting
Nonlinear mixing	High	Low	Potential overfitting

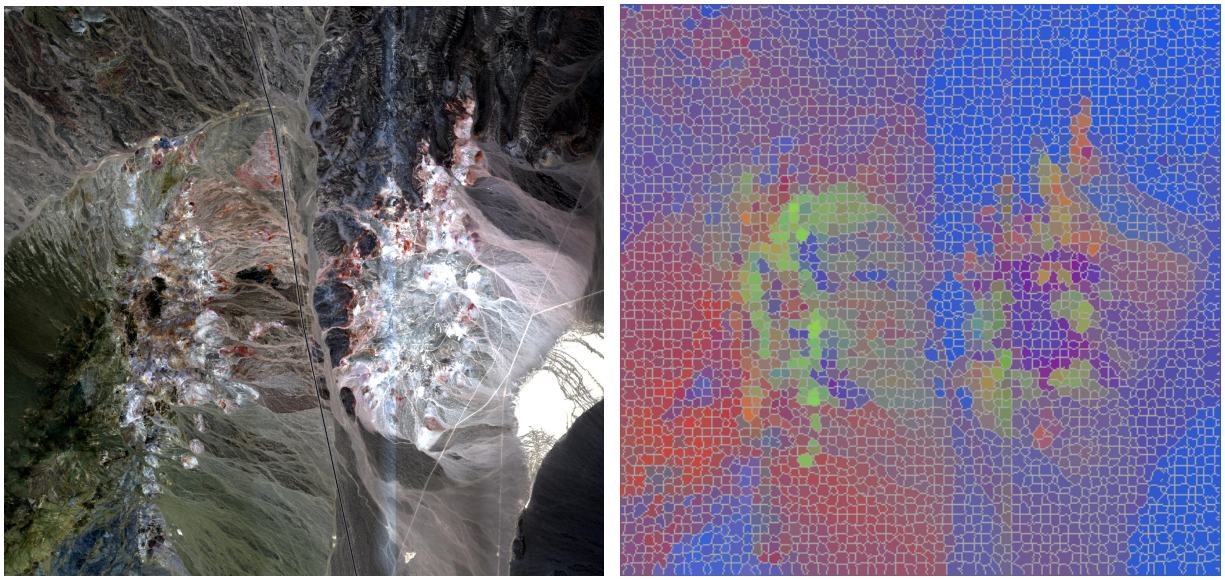

(a) RGB channels.

(b) Spectral superpixel segmentation using SLIC. Visualization is on a false RGB image.

Figure 2.5: Example of the superpixel segmentation of a spectral image.

in order to get meaningful results with reasonable computational resources. Since there is a high correlation between adjacent channels, methods such as Principal Component Analysis (PCA) are common for spectral dimensionality reduction [38, 39]. Another relatively recent strategy consists in taking advantage of the existing high correlation between neighboring pixels. In standard RGB images, superpixel algorithms are popular for finding features and reducing processing complexity [1]. They have also been used for various spectral data analysis tasks [20, 22, 54]. Figure 2.5 shows an example of the Simple Linear Iterative Clustering (SLIC) algorithm applied to a spectral image.

In conclusion, there exist multiple spectroscopic analysis techniques. Among them, there is an overall trade-off between computational simplicity (both in terms of implementation and data processing), physical interpretability, and accuracy. For an autonomous rover scenario involving probabilistic processes, which are robust in the presence of noise and discrepancies, full spectral analysis and linear mixing techniques seem to be the most appropriate to apply. Although physical interpretability may be lost, the robot's goal is to efficiently communicate high-level results to the scientist, and spectral analysis lies as a low-level measurement processing operation. Additionally, preprocessing steps such as PCA and superpixels turn out to be quite useful in this context, especially in terms of data processing efficiency.

Chapter 3

Science Hypothesis Map

Recapitulating, Chapter 1 explains the hierarchical structure of the STM, as well as the basics of Bayesian experimental design. It also mentions their potential to improve autonomous robotic exploration.

Consequently, this chapter first develops the notion of the science hypothesis map by merging these two concepts, specifically the STM with Bayesian inference: a spatial probabilistic framework that allows to infer high-level investigation objectives from low-level sensor measurements on a map. It later describes how to perform Bayesian inference with the science hypothesis map, and finally discusses limitations and potential future developments.

3.1 Structure

First of all, there is a strong spatial component involved when dealing with remote exploration [55]. This process may be seen as a region labeling task. To illustrate this, Figure 3.1 shows a geologic map of Mars. Note that the labeling is uniform within craters, valleys, plateaus, et cetera.

Another fundamental ingredient in the formulation of the science hypothesis map is using a probabilistic model. Stochastic models have proven superior over straightforward deterministic models when dealing with uncertainty and noise. An important thing to consider is flexibility in the model by being able to operate with both quantitative and qualitative (also known as categorical) random variables. Most methods solely work with quantitative variables, but many scientific concepts in geology, biology, chemistry, etc. are actually better modeled as categorical variables.

The *science hypothesis map* is a spatial probabilistic structure that allows to perform inference of abstract scientific concepts from raw sensor measurements. It has the following hierar-

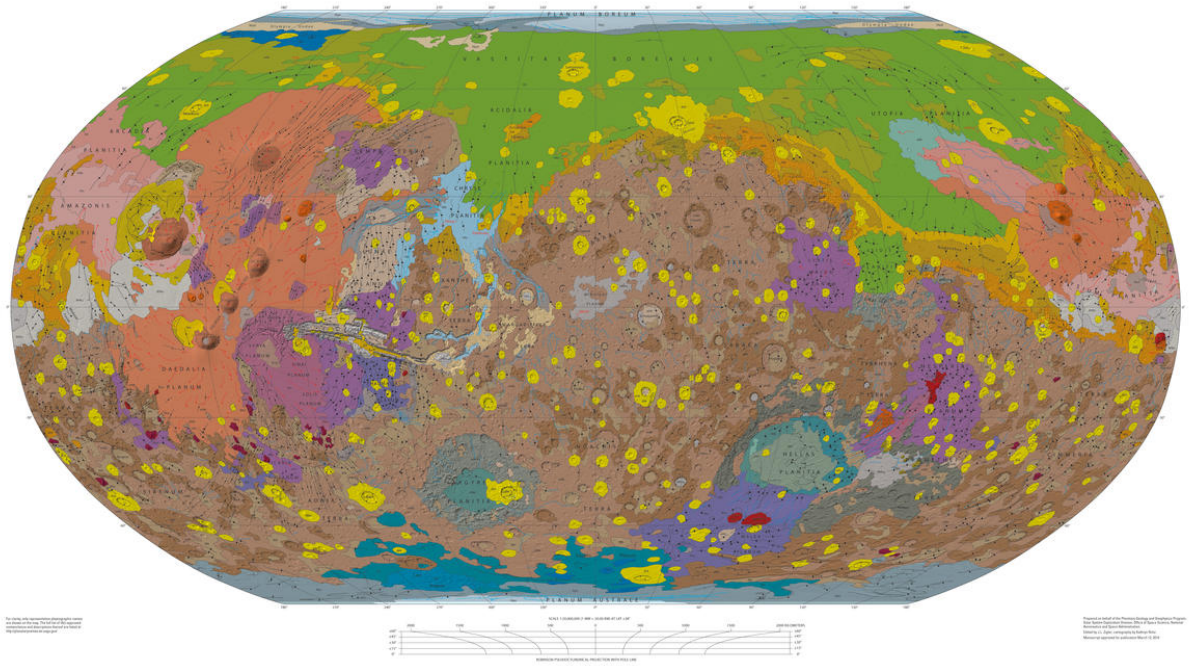

Figure 3.1: Example of a geologic map of Mars recording the distribution of geologic formations and landforms on the planet's surface through time. Note that the labeling is uniform within craters, valleys, plateaus, et cetera. Image courtesy of the United States Geological Survey (USGS) [52].

chical components, visually represented in Figure 3.2:

- The map partitions into k independent, predefined spatial *regions*, labeled as $R : \{r_1, \ldots, r_k\}$, (see left of Figure 3.2).
- The *investigation objectives* estimate abstract properties which are themselves unknowns, labeled as $H : \{h_1, \ldots, h_l\}$. Each region can be explained by an investigation objective with probability $P_R(H)$.
- *Physical properties*, labeled as $Y : \{y_1, \ldots, y_m\}$, have unique associations with the investigation objectives, given by $P_R(Y|H)$.
- The robot collects n sensor *measurements*, labeled as $Z : \{z_1, \ldots, z_n\}$. These measurements could be noisy or dependent on observing conditions, so there is an indirect association between Z and Y given by $P_R(Z|Y)$.

In addition to these components, the science hypothesis map operates under the following assumptions and conditions:

- For mapping and computational convenience (further discussed in Chapter 5), the explored environment is partitioned into disjoint regions, each associated with its own independent conditional distributions relating investigation objectives, physical properties, and mea-

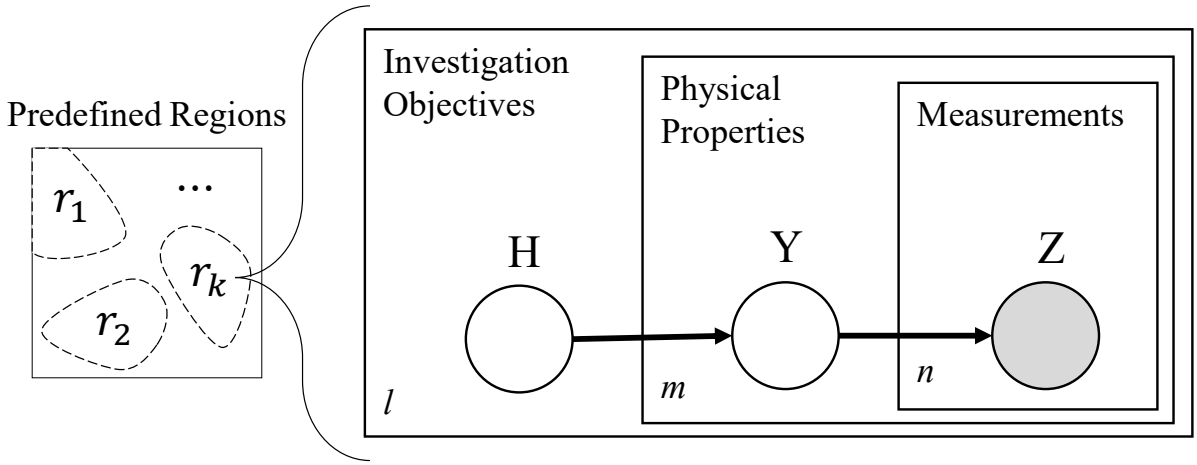

Figure 3.2: Plate notation of the science hypothesis map. Each spatial region $r \in R$ has an independent hierarchical conditional distribution of investigation objectives $h \in H$, physical properties $y \in Y$, and measurements $z \in Z$.

surements (Figure 3.2). This notion is easily traceable and interpretable for the scientist, and it is compatible with the scientist's analysis procedure [43, 61].

- The scientist assigns a prior probability vector $P_R(H)$ to each spatial region. This probability distribution is updated with evidence, and may be observed and modified by the scientist anytime during the mission, granting interpretability and flexibility.
- Initially, the rover only knows the probability vector $P_R(H)$ and the spatial segmentation of the map into regions.
- The rover is only capable of diagnosing the associated physical property Y of a sensor measurement Z at each point-like location.
- Orbital data is only used to define the boundaries of spatial regions, and they are kept fixed throughout the whole mission since there is almost no ambiguity in spatial segmentation (e.g. craters, volcanoes, plateaus, etc. are relatively simple to delimitate) [43, 61].

3.2 Bayesian Inference

The previous probabilistic model can be used to infer the corresponding investigation objective H for each region $r \in R$, given a measurement Z. That is, find a closed-form expression for $P_R(H|Z)$ in terms of the known conditional distributions $P_R(Y|H)$ and $P_R(Z|Y)$. First, the joint probability model of H, Y, and Z for each region in R can be decomposed using the chain

rule:

$$P_R(H, Y, Z) = P_R(H)P_R(Y|H)P_R(Z|Y). \tag{3.1}$$

Integrating over variables:

$$P_R(H, Z) = P_R(H) \sum_Y P_R(Y|H)P_R(Z|Y), \tag{3.2}$$

$$P_R(Z) = \sum_H P_R(H) \sum_Y P_R(Y|H)P_R(Z|Y). \tag{3.3}$$

From the definition of conditional probability:

$$P_R(H|Z) = \frac{P_R(H, Z)}{P_R(Z)}. \tag{3.4}$$

Finally, the direct closed-form expression is:

$$P_R(H|Z) = \frac{P_R(H) \sum_Y P_R(Y|H)P_R(Z|Y)}{\sum_H P_R(H) \sum_Y P_R(Y|H)P_R(Z|Y)}. \tag{3.5}$$

This is equivalent to performing Bayesian inference, with $P_R(H)$ being the prior probability or initial belief of the explaining investigation objective of a region, given by the scientist in advance, and $P_R(H|Z)$ the posterior probability or updated belief.

3.3 Discussion

This approach is constructed in a way that allows to model and accommodate various tasks with a mapping component. Investigation objectives, physical properties, and sensor measurements are defined in a very broad way in this chapter. The next chapter serves as an example to help the reader work through these concepts via a concrete application: a geologic exploration scenario.

Of course, this approach has limitations. One of the fundamental assumptions of the model is the existence of independent spatial regions. Indeed, this notion is very helpful for simplifying the problem and is consistent with the scientist's mental process of analysis and interpretability. It has also been used by other researchers for various tasks such as 3D mapping [13]. Nonetheless, it may oversimplify spatial correlation between regions. A solution might be to use a Gaussian kernel [5] or a spectral similarity criterion with the purpose of propagating the inference process throughout the whole map whenever a new measurement is collected.

Other limitations include assuming that the region boundaries are static, and more importantly, that the rover has only access to these boundaries and the scientist's initial hypotheses

at the beginning of the exploration mission. Incorporating low resolution orbital data into the probabilistic model, as typically available prior to exploration missions, would prove very useful as well for propagating the inference process through other locations of the map.

Chapter 4

Geologic Exploration

Often planetary exploration, including surface exploration of Mars, the Moon, and other planetary bodies, involves mapping surface mineralogy to infer geologic composition, structure, origins, and ages. Geologists use the age of rocks to determine the sequence of events in the history of a planetary body. Geologic and mineralogical composition information helps them estimate what happened over time. Particularly important is the identification of rocks and minerals formed in the presence of water, since water is one of the key elements to sustain life [3]. Therefore, it would seem that planetary exploration reduces to geology.

Consequently, this thesis considers a geologic exploration scenario as proof of concept of the science hypothesis map. The following elements comprise the science hypothesis map for this specific scenario:

- The investigation objectives H are *geologic units* defining distinct units of material with different ages and formation processes. Figure 4.1 shows an example of a hydrothermal formation.
- The physical properties Y are *minerals*, distinctive chemical compositions that diagnose the geologic formation conditions of rocks. Figure 4.1 shows that a hydrothermal formation is characterized by the presence of minerals such as chlorite and serpentine.
- The measurements Z, *reflectance spectra*, representing the fraction of incident light reflected in each wavelength from the visible to shortwave infrared. Many minerals' molecular compounds have distinctive features in this range due to their distinctive chemical structure [17], as may be observed in Figure 4.1.

In order to train and test the science hypothesis map in a geologic mapping scenario, this project focuses on the Cuprite mining district of Nevada, a region with rich mineral content and extensive field study [53]. Figure 4.2 shows maps and images corresponding to the different

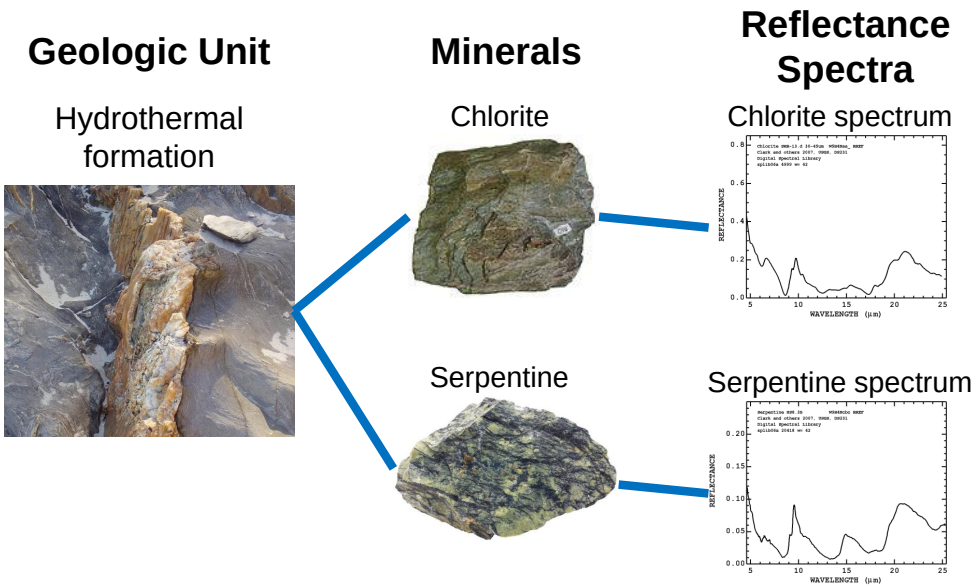

Figure 4.1: Visual example of the relation between geologic units, minerals and reflectance spectrum measurements. Elaborated with information from the United States Geological Survey (USGS) [16].

elements of the science hypothesis map for this site in Cuprite.

The rest of this chapter describes how these elements present in Cuprite: geologic units H, minerals Y, and spectra measurements Z (shown in Figure 4.2) come into play. Specifically, the followed procedure to train the hierarchical probabilistic model in the science hypothesis map using data from Cuprite, Nevada. That is, learn the association between geologic units and minerals, $P(Y|H)$, and the relation between minerals and reflectance spectrum measurements, $P(Z|Y)$.

4.1 Relation Between Geologic Units and Minerals

This subchapter focuses on learning the relation between geologic units H and minerals Y in Cuprite, NV. Maps of these two elements may be observed in Figures 4.2b and 4.2c, respectively. They are based on expert-drawn geologic and mineral maps by Swayze et al [53], which make a rough segmentation of the whole Cuprite site into regions. These maps are adjusted with manual control points so that they align on a per-pixel basis. Each independent geologic unit label and mineral are assigned a different value, illustrated as arbitrary false colors. These maps provide a ground-truth interpretation for each pixel in the scene. Note that each unit generally contains many minerals in different proportions. Scientists typically utilize the presence and abundance of key minerals to determine the involved geologic formation processes within a well-defined region. Therefore, these proportions are used to train a conditional probability table

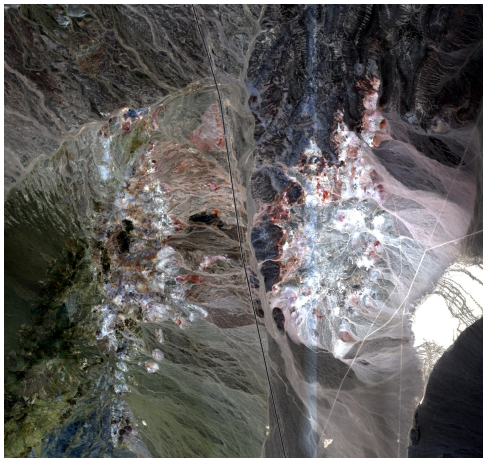

(a) Standard orbital RGB image of Cuprite.

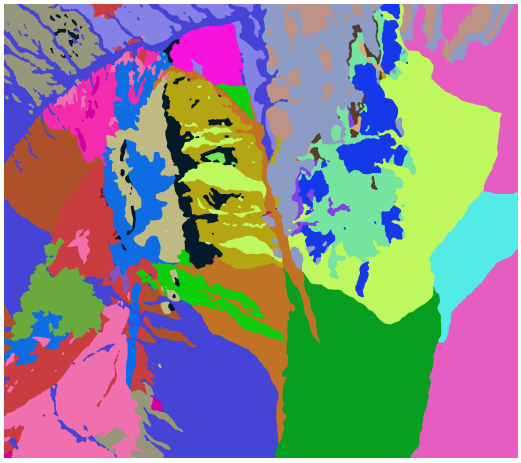

(b) Geologic units H in Cuprite.

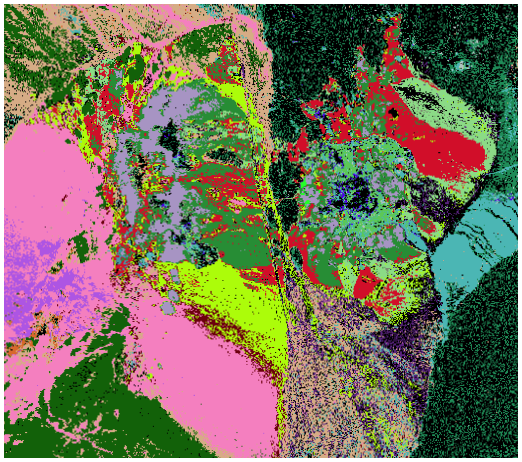

(c) Minerals Y and their abundances in Cuprite.

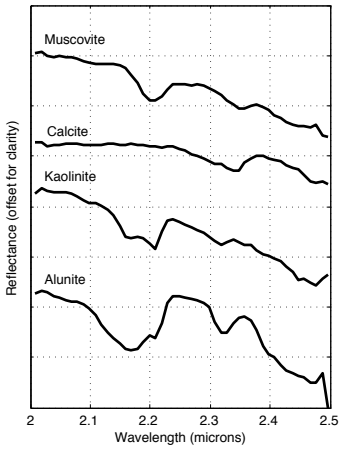

(d) Representative reflectance spectra Z of some key minerals in Cuprite. Each pixel is associated with a full spectrum.

Figure 4.2: Different elements related to the science hypothesis map for a geologic exploration scenario in Cuprite, Nevada. Geologic units and mineral abundances are illustrated as arbitrary false colors, and they are based on the work by Swayze et al [53].

of minerals given geologic classes, $P(Y|H)$, a simple and physically interpretable approach that is compatible with the scientist's notion described earlier. An example of this table is shown on Table 4.1. The actual conditional probability table consists of 32 geologic units and 20 minerals. The training set uses up to 2000 random samples from every predefined region. These datasets are withheld from the analysis and experiments that follow in Chapter 6.

Table 4.1: Example of the conditional probability table of minerals Y given geologic classes H.

$P(Y\|H)$	y_1	y_2	$\dots$	y_m
h_1	$P(y_1\|h_1)$	$P(y_2\|h_1)$	$\dots$	$P(y_m\|h_1)$
h_2	$P(y_1\|h_2)$	$P(y_2\|h_2)$	$\dots$	$P(y_m\|h_2)$
$\vdots$	$\vdots$	$\vdots$	$\ddots$	$\vdots$
h_l	$P(y_1\|h_l)$	$P(y_2\|h_l)$	$\cdots$	$P(y_m\|h_l)$

4.2 Relation Between Minerals and Reflectance Spectra

This subchapter focuses on learning the relation between minerals Y and reflectance spectra Z in Cuprite, NV. Maps and images of these two elements may be observed in Figures 4.2c and 4.2d, respectively. In-situ spectroscopic measurements performed by a robot with a spectrometer are simulated using an airborne instrument. Specifically, data from the Next Generation Airborne Visible Infrared Imaging Spectrometer (AVIRIS-NG) [28]. It assigns a unique reflectance spectrum measurement to every location (pixel) in the scene. AVIRIS-NG mapped the area of Cuprite at high spatial resolution (3.9 m per pixel) with radiance measurements from 380-2510 nm, also with a high spectral resolution (5.0 nm per channel). The data was acquired during overflights in 2014 and converted from measured at-sensor radiance to surface reflectance using the procedure described by [57].

A Gaussian Mixture Model (GMM) is used as a simple but effective probabilistic classifier that predicts the corresponding mineral Y from a spectrum measurement Z. It is a relatively standard approach for this task of spectral classification [40]. Since it uses complete signals rather than just absorption features, it would be considered as a full spectral analysis algorithm. It is simple to train and to implement on an autonomous navigation scenario, and it is relatively robust to noise. It is important to mention that the assumption of normally distributed data may be over-simplistic and yield low training and testing accuracies. However, it has an advantage over most of the full spectra analysis methods: its high physical interpretability. The reason is that it is possible to directly visualize the learned mean and covariance of each class, as opposed to common classifiers such as Random Forests or Support Vector Machines, which operate by

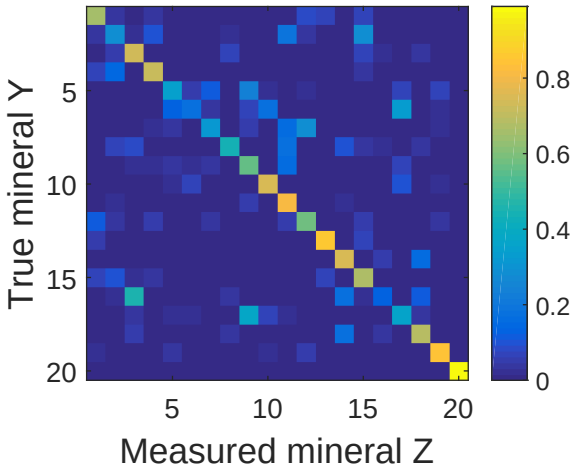

Figure 4.3: Empirical conditional probability table relating actual to estimated minerals for each of 20 possible mineral types in the Cuprite scene.

defining classification boundaries in high-dimensional spaces.

The model is trained using a held-out training set of locations from each mineral class. Since the used AVIRIS dataset has a very high resolution, its dimensionality is reduced to 20 dimensions with Principal Component Analysis (PCA). Then, a mean and covariance matrix is fit to each class using the shrinkage estimator described in [31]. The shrinkage estimator treats the actual covariance as a linear combination of the sample covariance, and a regularized version consisting of the diagonal elements only. A closed form solution to the leave-one-out cross validation (LOOCV) likelihood allows efficient computation of the optimal linear combination. The end result is a probability density for every combination of mineral and spectrum in the image.

Treating spectra measurements Z as a discrete variable simplifies information gain calculations, described with detail in Chapter 5. Consequently, the classification output of the GMM classifier is treated as a categorical answer. Figure 4.3 shows the confusion matrix from a held-out validation set to infer its probability of error on any future measurement. It has an overall accuracy of 57.15%, with a min/max class accuracy of 12.50% and 98.95%, respectively. Although it is a classifier with a relatively low accuracy, its confusion matrix acts as a conditional probability table that compensates the noise in the measurement process which could cause estimated minerals from spectra measurements Z to differ from the true physical variables Y. The used training and validation sets each use up to 1000 samples from every mineral type. As well as with the learned relation between geologic units H and minerals Y, these datasets are withheld from the analysis and experiments that follow in Chapter 6.

4.3 Summary

There are a copuple of key concepts from this chapter that are worth summarizing. First of all, geologic investigations are fundamental for planetary exploration due to the existing correlation between geology and the potential presence of water and life; hence the importance of developing a geologic exploration scenario as proof of concept for the science hypothesis map.

For this geologic scenario, the elements from the science hypothesis map: investigation objectives, physical properties, and measurements; directly traduce to geologic units, minerals and spectra measruments; respectively.

The site of Cuprite, Nevada is chosen for these study and modeling tasks since it is has a well studied geologic and mineralogical diversity.

The probabilistic relation between geologic units and minerals is modeled as a conditional probability table; and the probabilistic relation between geologic units and minerals is modeled as a Gaussian Mixture. These methods are chosen due to their conceptual and implementation simplicity, as well as for their high physical interpretability.

Finally, chapter 6 evaluates the accuracy of the overall Bayes network while used for inferring high-level geologic units from low-level spectra measurements.

Chapter 5

Informative Path Planning

Previous chapters have solely focused on the Bayesian inference process of the Bayesian experimental design loop. This chapter now covers experimental design, the remaining component to close the whole loop.

In our robotic exploration scenario, experimental design directly translates to the class of path planning where a mobile agent must explore an environment by planning an information-optimal path to reach a prescribed end-of-day goal location, optimally balancing navigation and resource costs against meaningful science measurements. During the mission, the robotic explorer collects a sequence of n measurements at spatial locations $X = (x_1, x_2, \ldots, x_n)$, with the corresponding possible combinations of the n measurements as Z^n.

First, this chapter describes with detail the information reward function that needs to be maximized during the mission. Then, it discusses computational challenges for calculating this function, as well as an efficient solution to this particular problem. Finally, it discusses informative path planning and different state-of-the-art planning methods that have been used for similar exploration tasks and that may be integrated with the science hypothesis map.

5.1 Information Gain Objective Function

In probabilistic terms, classical information-driven action selection [6, 41, 55] and Bayesian experimental design [11] have traditionally used information gain as the utility function to optimize. Specifically, information gain is defined as the expected reduction in uncertainty after collecting new information. Uncertainty is measured with Shannon entropy in this context [19]. The Shannon entropy of an independent region in R is:

$$I_R(H) = -\sum_H P_R(H) \log P_R(H). \tag{5.1}$$

The expected entropy of the posterior distribution given n measurements is:

$$I_R(H|Z^n) = -\sum_{Z^n} P_R(Z^n) \sum_H P_R(H|Z^n) \log P_R(H|Z^n). \tag{5.2}$$

From equations 5.1 and 5.2, the information gain for each region can be calculated as:

$$IG_R(H|Z^n) = I_R(H) - I_R(H|Z^n). \tag{5.3}$$

Information gain is additive across independent variables. Then, the objective function for the whole map can be represented as the sum of regions' information gains:

$$IG = \sum_R IG_R(H|Z^n). \tag{5.4}$$

It is important to mention that information gain is not the only existing approach for informative planning. Other authors have used different probabilistic objectives such as variance reduction [7, 8, 33, 50]. They have proven useful in many cases, but they cannot be applied directly to categorical variables, such as minerals and geologic units. The issue is that there is no valid notion of mean or variance with these variables, e.g. it would make no sense to calculate the mean with respect to minerals. But information gain is applicable to categorical variables, so that is the rationale for utilizing information gain as the objective function in this thesis.

5.2 Information Gain Computation

A real time application may face the problem of calculating information gain efficiently. In this case, the sum over all possible sequences of n measurements has exponentially many terms, as given by:

$$|Z^n| = \binom{m+n-1}{n-1}. \tag{5.5}$$

It can be noticed that m is a constant: the total number of possible different minerals. Therefore, the combinatorial explosion when calculating the expectancy over alternatives (equations 5.2 and 5.3) is due to the number of measurements n.

While information gain has been estimated with Markov chain Monte Carlo (MCMC) approximations before [49], here the distributions $P(Z^n)$ and $P(H|Z^n)$ are not known *a priori*.

An alternative is to perform Monte Carlo integration [60], with the possibility of using simple strategies such as uniform or importance sampling. But another problem arises: either the approximation accuracy deteriorates or the computation time increases as n increases linearly.

However, given the problem formulation in this thesis, there is a more robust and efficient approximation for large n. Information gain is a function that as long as Z and H are not independent, and the observations are conditionally independent given H, the following holds:

$$\lim_{n\to\infty} IG_R(H|Z^n) = I_R(H). \tag{5.6}$$

as shown by Haussler and Opper [29]. At first glance, the Bayes network from Figure 3.2 in Chapter 3 suggests that the observations are conditionally independent given Y, and not necessarily given H. However, since information gain (equations 5.2 and 5.3) is calculated by marginalizing Y (equation 3.5), these properties hold true.

This means that it is plausible to fit a curve to the first few data points (i.e. small number of measurements n), which are simpler to calculate or to approximate, specifically, a monotonically increasing curve that converges to a maximum. In this work, a *generalized logistic function* (GLF) is used, also known as Richard's curve, which is a generalization of the sigmoid curve [47]. It was originally developed for biological growth modeling, but it can be applied to other areas as well. It is given by the following expression with six parameters:

$$Y(t) = A + \frac{K - A}{(C + Qe^{-Bt})^{\frac{1}{v}}}. \tag{5.7}$$

To illustrate these concepts, Figure 5.1 demonstrates an example of information gain calculation for a simple probabilistic process with a uniform prior. It uses nonlinear least squares fitting to find a suitable GLF approximation. It also shows a Monte Carlo approach, as well as the true information gain values. It can be easily observed that the accuracy of the Monte Carlo method deteriorates for large n, but not for the GLF method. In this example, only the first six Monte Carlo points are used for the GLF fit. All Monte Carlo approximations are calculated with $10,000$ random points each by performing importance sampling. The proposal distribution is a uniform distribution over all possible combinations of Z^n, where points are sampled with the method by Smith and Tromble [51]. The goodness of fit statistics for $0 \leq n \leq 15$ are: $R^2 = 0.9999$ and $RMSE = 0.0092$, showing that the function can predict IG_R for many samples.

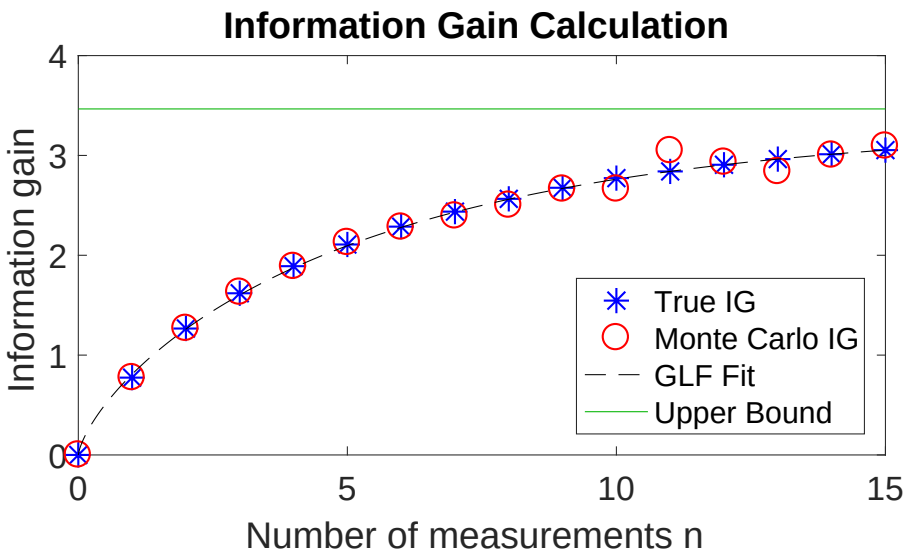

Figure 5.1: Information gain calculation for a uniform prior using three different methods: direct, Monte Carlo integration, and GLF fitting.

5.3 Information Gain-based Path Planning

This chapter formulates the path planning objective as maximization of information gain subject to a fixed resource budget (time, energy, etc.) that the robot can travel before finally reaching a desired end-of-day goal. Implicitly, $x_1 = x_{start}$ and $x_n = x_{end}$. The optimization problem is:

$$\arg\max_{X} \; IG \quad \text{subject to} \quad Cost(X) \leq Budget. \tag{5.8}$$

The science hypothesis map integrated with this formulation in equation 5.8 can accommodate many different path planning strategies, which will be discussed next. However, it is important to first underscore the following statement: information gain, as opposed to some classical objective functions, is a submodular function [23, 34]. Formally, a function $f : 2^{\Omega} \rightarrow \mathbb{R}$, where Ω is a finite set and 2^{Ω} denotes the power set of Ω, is submodular if and only if:

$$\forall S, T \subseteq \Omega \text{ we have that } f(S) + f(T) \geq f(S \cup T) + f(S \cap T). \tag{5.9}$$

In this context, it intuitively means that the reward between samples is not independent, and that sampling the same region over and over again leads to diminishing expected returns. Another way to put it is like this: the marginal utility tends to decrease as more measurements are collected within the same spatial region. The submodularity of the information gain objective function poses an interesting exploitation vs. exploration trade-off that translates to taking many samples from a few highly-rewarding regions, or to visit as many different regions as possible.

Many scenarios assume that the rewards between sampled points are independent. One example is the classical problem of orienteering in graphs [14]. There also exists a straightforward method that uses mixed integer linear programming (MILP) to generate optimal informative paths; it has been successfully applied to simulated adaptive underwater exploration [63]. Although these methods might yield relatively good results, the assumptions obviously do not hold true in our scenario.

Arora et al use a Monte Carlo tree search (MCTS) approach for information-gain based path planning in a more or less similar geologic exploration scenario [5]. Besides, Monte Carlo methods are relatively simple to apply to various scenarios. Nevertheless, their performance is highly dependent on factors such as sampling strategies and number of iterations, and they might yield radically different solutions from one execution to another if they are not properly tuned.

There is a family of significantly more complex algorithms by Krause et al that directly take into account the property of submodularity and produce near-optimal solutions [33, 34, 50]. They integrate multiple methods and perform several approximations during the planning process. They have been mainly applied to sensor placements in Gaussian processes.

Other approaches may be computationally intensive, but potentially closer to optimality, such as using branch and bound (B&B) techniques for both discrete [7] and continuous [32] space representations. They directly consider submodularity, and also have the advantage of using an tunable parameter for the number of lookahead steps, this in order to sacrifice optimality for computation speed.

5.4 Summary

The main ideas that are explained and discussed in this chapter may be reduced to the following. First, information gain is the classical utility function that has been used for information-driven action selection and Bayesian experimental design, and it results especially suitable because of the nature of the inherent categorical variables of our geologic exploration scenario.

Nonetheless, its worth mentioning that using information gain comprises a couple of important implementation challenges. The first one is finding an efficient computation procedure of information gain (which was solved for our specific scenario), and the second one is choosing a suitable path planning algorithm that takes into account its submodularity.

Many existing state-of-the-art informative path planning methods can work together with the science hypothesis map and an information gain objective function, but there might be a substantial trade-off between optimality and complexity.

Now that the whole Bayesian experimental design has been covered, the next chapter evalu-

ates the performance of many of these planners in conjunction with the science hypothesis map in a simulated geologic exploration scenario.

Chapter 6

Experiments and Analysis

In a very broad sense, the main objective of this chapter is to corroborate the validity the following statement: *Bayesian experimental design applied to robotic exploration using the science hypothesis map overperforms conventional science-blind or spatial-coverage methods in terms of prediction accuracy and uncertainty reduction*. It may be useful to remember that within this context, Bayesian experimental design comprises two processes: Bayesian inference and path planning. Therefore, this chapter is organized as follows: the first subchapter focuses on evaluating the Bayesian inference process individually, the second subchapter evaluates the complete loop (Bayesian inference in conjunction with path planning), and the final subchapter contains an overall discussion of these results.

6.1 Bayes Learning

This group of experiments evaluated the model's ability to predict and recover the true geologic unit of a region in Cuprite (with a specific age, formation process, etc.) from high-resolution AVIRIS spectral measurements with the Bayesian update.

6.1.1 Experimental Setup

For each of the 32 regions in Cuprite, 200 random sampling sequences were generated and averaged. Non-training points were sampled without replacement, simulating an exhaustive exploration with no budget constraints.

To evaluate the model's ability to recover the true geologic unit of each region, this experiment tracked two fundamental variables: the posterior probability of the correct geologic unit, and the updated entropy. As more and more measurements are collected, the updated probability

of the true underlying geologic unit should ideally converge to 1, and consequently, the entropy should converge to 0. These convergence rates depend on the prior distribution, therefore, different initial conditions were tested through the following three representative scenarios:

- Starting with an accurate prior: to represent a situation where the scientist has a relatively accurate initial belief, a value of 50% was assigned to the prior probability of the correct geologic unit, while the remaining probability mass is distributed to alternatives. This prior allows ample margin for converging to either a correct or incorrect answer.
- Starting with complete uncertainty: a situation of complete uncertainty was represented with a uniform prior over geologic unit labels.
- Starting with an inaccurate prior: to represent a situation where the scientist's initial beliefs are actually incorrect, a prior probability of 50% was assigned to an incorrect unit: the most similar geologic unit according to the Hellinger distance for probability distributions, a challenging error to correct. The remaining probability mass is distributed uniformly.

6.1.2 Results and Discussion

Starting with an accurate prior: The results for this set of experiments are shown in Figure 6.1. Most beliefs converge to the right answer with just a few samples, while entropies reduce significantly. The three exceptions with poor outcomes correspond to the smallest regions trained with the fewest data points. This means that overall, the model is able to converge by itself to the right geologic unit when starting with a relatively good guess, which is to be expected in this case.

Starting with complete uncertainty: The results for this set of experiments are shown in Figure 6.2. Despite the significantly more challenging situation than earlier, most beliefs still show significant improvement. However, there is more variance and an overall slower convergence rate; but the model is still able to operate with success.

Starting with an inaccurate prior: The results for this set of experiments are shown in Figure 6.3. The plots shows that there is an even slower improvement, as well as a larger variance. But in most cases the model still recovers. It is important to underscore the implications of this notoriously challenging scenario: despite the fact that a scientist makes a mistake, the model is able to infer the correct geologic unit by itself when having the opportunity of collecting sufficient measurements.

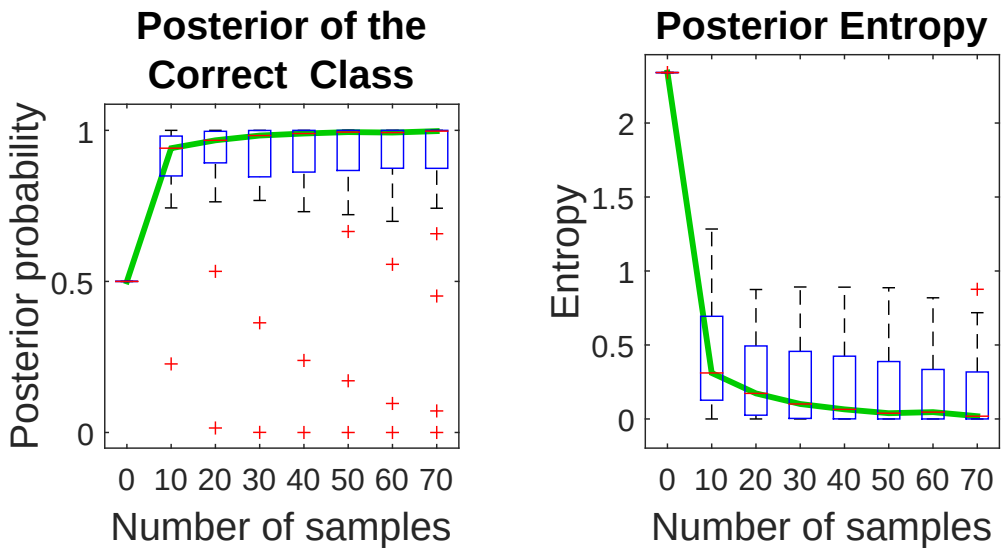

Figure 6.1: Posterior probability (left) and entropy (right) of the correct geologic unit in Cuprite, starting with an accurate prior. Each data point in the box and whisker plots corresponds to one of the 32 different regions, and the green lines denote the medians.

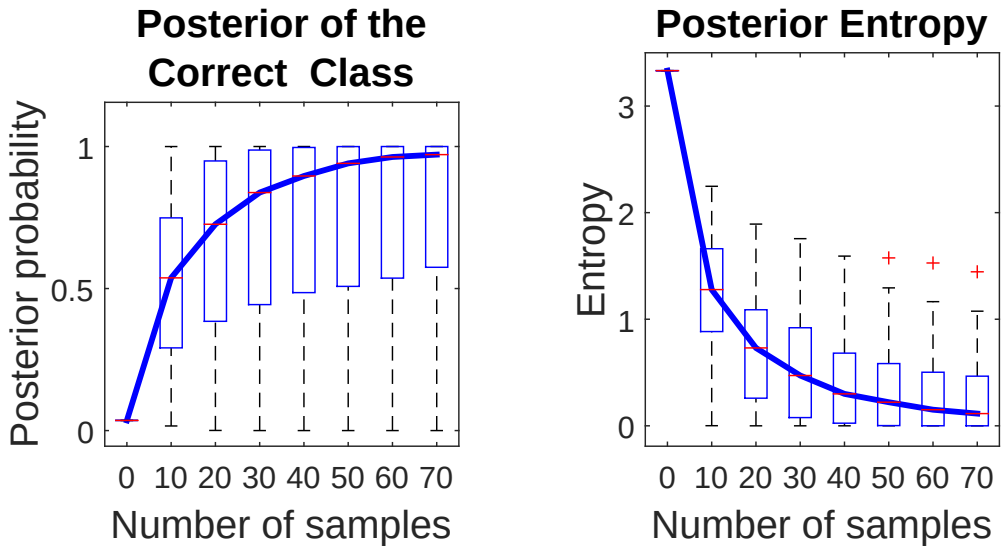

Figure 6.2: Posterior probability (left) and entropy (right) of the correct geologic unit in Cuprite, starting under complete uncertainty. Each data point in the box and whisker plots corresponds to one of the 32 different regions, and the blue lines denote the medians.

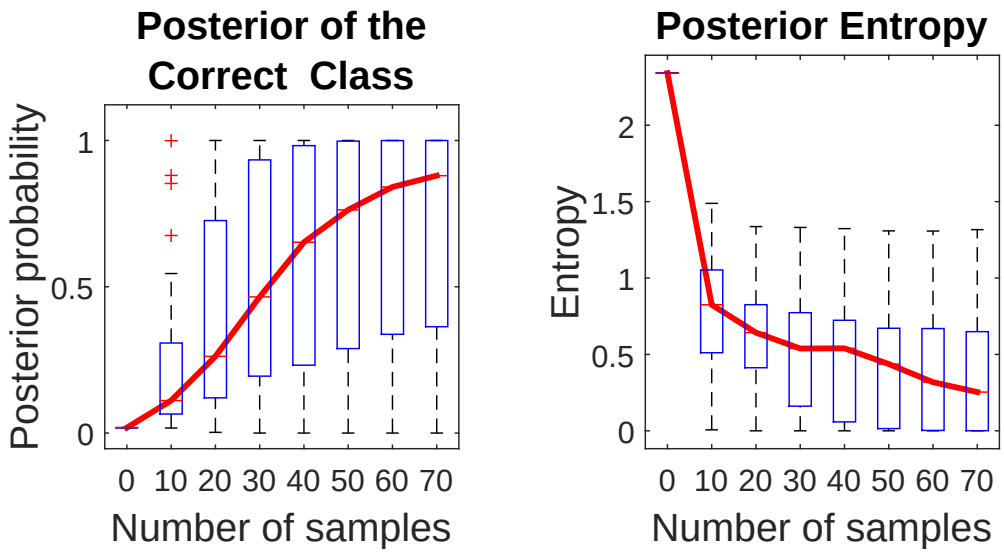

Figure 6.3: Posterior probability (left) and entropy (right) of the correct geologic unit in Cuprite, starting with an inaccurate prior. Each data point in the box and whisker plots corresponds to one of the 32 different regions, and the red lines denote the medians.

6.2 Path Planning

This group of experiments integrates the science hypothesis map with different information gain-driven path planners. It compares their performance in order to evaluate the potential benefits in science yield when using information gain-based planning over conventional science-blind planners.

6.2.1 Experimental Setup

The whole Cuprite region was divided into three sites with the same area: A, B, and C. For simplicity, each site is treated as a 2D discrete graph that is formed by an 8-connected rectangular grid, with locations spaced at 50 pixels (200 m). The full map has a size of 2555×2268 pixels (10.22×9.07 km). For each site, 200 different pairs of random start and end points were generated, located at the East and West edges of the map. The cost function was directly proportional to the path length, and the assigned budget permitted up to 1.5 times the shortest path length, allowing sufficient budget to explore alternative regions. All simulated traverses and regions had the same initial conditions: a prior probability of 40% assigned to both the correct geologic unit and the most similar unit according to the Hellinger distance, while the remaining probability mass was distributed to alternatives. This represented a challenging scenario where the scientist is equally inclined toward two similar options.

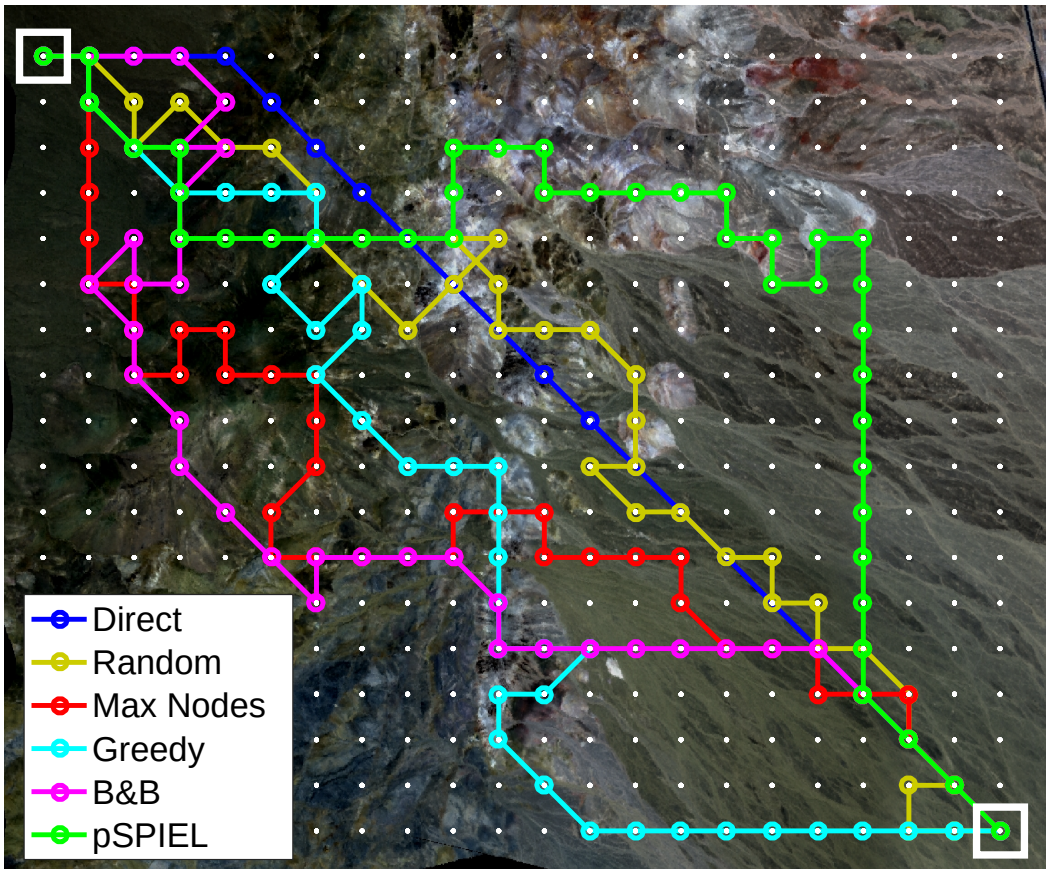

Figure 6.4: Example paths from the 6 path planners. The start is in the upper left, and the end in the lower right.

6.2.2 Tested Path Planners

Figure 6.4 shows an example of paths generated by all the tested planners. As a control case, this experiment evaluates three science-blind planners that ignore information gain:

1. A *Direct* (D) path planner that selects the sequence of waypoints that minimizes the total path length with the standard Dijkstra algorithm.
2. A *Random* (R) path planner that sequentially selects a set of random neighboring waypoints that do not violate the budget on path length.
3. A *Max-N* (M) path planner that maximizes the number of visited nodes using a classic recursive greedy approach for orienteering [14]. Since there may be multiple valid solutions, it performs a random permutation of the waypoints' ids in order to avoid biased paths (e.g. paths that prioritize going to the left).

Against these, this experiment evaluates three science-aware planners that incorporate measurements' information gain:

Table 6.1: Average scores for relevant path planning metrics.

Site	Metric	**Direct**	**Random**	**Max-N**	**Greedy**	**B&B**	**pSPIEL**
A	Path Length (m)	3937	5582	**6129**	5639	5815	5817
	Collected Measurements	15.99	23.24	**27.67**	23.43	24.94	26.74
	Explored Regions	5.32	5.95	6.03	7.71	8.20	**8.77**
B	Path Length (m)	3789	5430	**5892**	5368	5511	5608
	Collected Measurements	13.34	19.88	**23.66**	18.82	20.44	22.51
	Explored Regions	5.62	6.32	6.38	8.05	8.83	**9.65**
C	Path Length (m)	4045	5798	**6285**	5468	5937	5974
	Collected Measurements	12.29	18.77	**21.67**	17.49	20.31	21.32
	Explored Regions	3.94	4.38	4.49	5.24	5.53	**5.88**

1. A *Greedy* (G) path planner that adds at each step the neighboring waypoint that maximizes the objective function, this without exceeding the budget.
2. A *Branch and bound* (B) path planner as a nonmyopic alternative. It selects a sequence of waypoints using the algorithm described by Binney and Sukhatme [7]. In this particular case, implementing a three-step calculation look ahead.
3. A *pSPIEL* (P) planner that combines submodular orienteering algorithms by Singh et al. [50] with this thesis's information gain objective function.

Additionally, these planners updated the hypotheses and paths every time a new measurement was collected, this with the goal to simulate an explorer that adapts its path dynamically with each new observation.

6.2.3 Results and Discussion

First of all, spatial-coverage metrics can be very useful for a preliminary assessment of the performance of the various path planners. Table 6.1 shows the average scores per traverse of three simple but useful metrics: path length, number of collected measurements, and number of explored regions. As expected, the Direct planner minimizes the first two metrics, whereas the Max-N planner maximizes them. The rest of the planners have an intermediate performance, with the Random and Greedy planners getting similar scores, and the pSPIEL planner achieving the second best performance overall. Nonetheless, all the science-aware planners explore more regions than any of the science-blind planners. At first glance, this indicates that exploration tends to be favored over exploitation in this specific scenario. This notion is compatible with the submodular property of information gain.

Now the performance of the planners is measured from an information-theoretic perspective with the following variables: the evolution of the posterior probability and its entropy. The

evaluation of a path across the map is done by simply adding the corresponding updated region's metrics, where a poor performance in a region penalizes the global score, and vice versa. These scores have a high variance because they strongly depend on the assigned path length budget. Therefore, the scores are normalized using feature scaling with respect to the best and worst planners for each given pair of start-end locations. The resulting scaled scores for prediction accuracy and entropy are shown in Figures 6.5 and 6.6. At the same time, Tables 6.2 and 6.3 show a set of paired t-tests comparing the performance among planners with their raw unnormalized scores with respect to the inference and entropy metrics, respectively.

These box plots and hypothesis tests show a couple of things. First of all, most of the times there exists a statistically significant difference in performance among all the planners, validating the notion that some planners indeed yield better results than others. The science-blind planners have a performance more or less proportional to the number of measurements, being Max-N the best. However, that is not necessarily true when these are compared to the science-aware planners: they tend to achieve superior scores in both of the information-theoretic metrics, even when some of them spend less budget or collect less measurements in average (e.g. Greedy vs. Random and Max-N).

This shows that there are more meaningful science measurements than others, and balancing them adequately leads to higher science productivity. For instance, the nonmyopic planners (B&B and pSPIEL) outperform the Greedy method since it only has a one-step look ahead. pSPIEL is the best overall method because it analyzes and weights all reachable regions using an approximation graph, favoring the ones with the highest expected rewards. On the other hand, the B&B method is limited by a three-step calculation horizon.

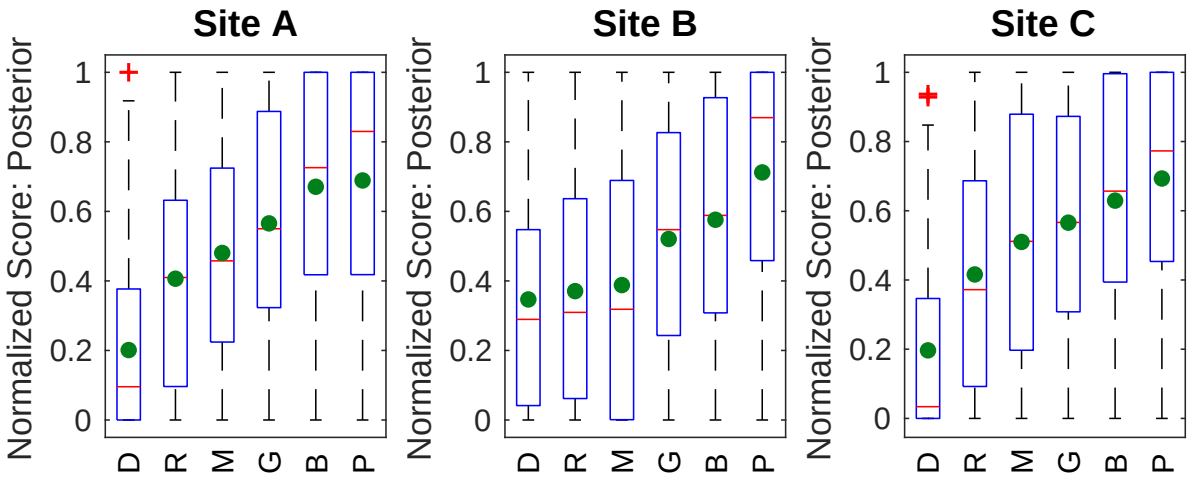

Figure 6.5: Comparison of planners based on their ability to recover correct geologic units per traverse using a normalized score. The green dots are the means.

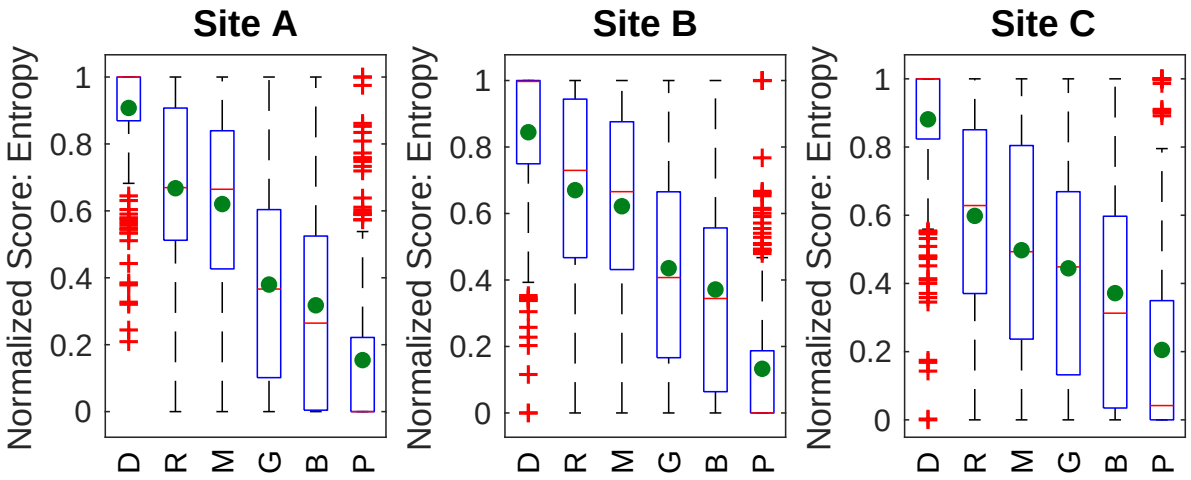

Figure 6.6: Comparison of planners based on their final entropy per traverse using a normalized score. The green dots are the means.

Table 6.2: Set of paired t-tests comparing the performance among planners with respect to the inference metric, using a significance level of 5%. In most cases, there is a statistically significant difference (Y/N).

Site	A						B						C					
Planner	**D**	**R**	**M**	**G**	**B**	**P**	**D**	**R**	**M**	**G**	**B**	**P**	**D**	**R**	**M**	**G**	**B**	**P**
D	-	Y	Y	Y	Y	Y	-	N	N	Y	Y	Y	-	Y	Y	Y	Y	Y
R	-	-	Y	Y	Y	Y	-	-	N	Y	Y	Y	-	-	N	Y	Y	Y
M	-	-	-	Y	Y	Y	-	-	-	Y	Y	Y	-	-	-	N	Y	Y
G	-	-	-	-	Y	Y	-	-	-	-	N	Y	-	-	-	-	Y	Y
B	-	-	-	-	-	N	-	-	-	-	-	Y	-	-	-	-	-	Y
P	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-

Table 6.3: Set of paired t-tests comparing the performance among planners with respect to the entropy metric, using a significance level of 5%. In most cases, there is a statistically significant difference (Y/N).

Site	A						B						C					
Planner	**D**	**R**	**M**	**G**	**B**	**P**	**D**	**R**	**M**	**G**	**B**	**P**	**D**	**R**	**M**	**G**	**B**	**P**
D	-	Y	Y	Y	Y	Y	-	Y	Y	Y	Y	Y	-	Y	Y	Y	Y	Y
R	-	-	Y	Y	Y	Y	-	-	N	Y	Y	Y	-	-	Y	Y	Y	Y
M	-	-	-	Y	Y	Y	-	-	-	Y	Y	Y	-	-	-	N	Y	Y
G	-	-	-	-	Y	Y	-	-	-	-	Y	Y	-	-	-	-	Y	Y
B	-	-	-	-	-	Y	-	-	-	-	-	Y	-	-	-	-	-	Y
P	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-

6.3 Summary

Overall, the experiments seem to support the statement at the beginning of this chapter: *Bayesian experimental design applied to robotic exploration using the science hypothesis map overperforms conventional science-blind or spatial-coverage methods in terms of prediction accuracy and uncertainty reduction.*

First, the experiments demonstrate that the science hypothesis map can infer abstract geologic units with increasing accuracy by repeatedly collecting low-level spectral measurements. The hierarchical structure comprising geologic units, minerals, and spectra measurements turns out to work successfully both in terms of simplicity and accuracy. Even with misleading priors, the model seems to recover. However, the convergence rate and the variance of the results increase significantly the more challenging the initial conditions are. It would be appropriate to begin the mission with a relatively informed initial guess in order to avoid overspending exploration resources.

Second, multiple simulated exploration missions indicate that science-aware planners overperform conventional science-blind planners when having a limited exploration budget. Indeed, the used path-length constraint is over-simplistic and fails to capture more realistic issues such as time constraints, slope traversability and obstacle avoidance. Nonetheless, it proves to be sufficient to show the improvement in performance when using the science hypothesis map together with informative path planning. Especially, planners that have a further lookahead yield better information-theoretic results, such as the pSPIEL and B&B algorithms.

These results from simulations are encouraging to continue developing the science hypothesis map even further. Of course, field experiment results would be desirable as a next step, both on geologic and other exploration scenarios.

Chapter 7

Conclusion and Future Work

7.1 Conclusion

There are several key findings of this project. First, the experiments demonstrate that the science hypothesis map is a sound framework for describing exploration for specific well-defined objectives. Under diverse challenging situations, the probabilistic model can infer abstract geologic units with increasing accuracy by repeatedly collecting low-level spectral measurements.

In the case of robotic path planning, path planners that exploit information gain, efficiently calculated with the GLF approximation, consistently better reduce uncertainty over the investigation objectives when compared to typical spatial-coverage planners. Although the Max-N planner tends to collect the highest number of measurements, selecting sampling locations intelligently turns out to be critical. This is especially useful when there exists a limited exploration budget. Not surprisingly, the further the lookahead of the planner, the better the results. That is why the pSPIEL planner proves to be the best algorithm, followed by the B&B planner with a three step lookahead, and then by the Greedy approach with just one step lookahead.

In general, the complete Bayesian experimental design loop, i.e. the cyclic process of Bayesian inference in conjunction with path planning, proves to be successful in simulations. The experiments show that there is a statistically significant improvement in performance during the multiple simulated exploration missions, both in terms of uncertainty reduction and prediction accuracy. These results show promise and support the idea of applying these integrated methods to planetary remote exploration for geologic mapping.

A more surprising finding is the strength of the statistical link between measured spectra, the corresponding mineral types, and the geologic unit classes. This is not a foregone conclusion because geologists form this classification from countless other features such as local morphology, elevation, the wide-area geographic three-dimensional structure of different strata, and domain

knowledge. It is striking that AVIRIS-NG reflectance spectra, drawn independently and randomly from each unit, so strongly predict the units unique type out of over thirty alternatives.

7.2 Future Work

In general, future work will expand or improve two main areas. One is the fidelity and realism of the models and experiments, and the other one is the application of these methods and techniques to other problems besides geologic exploration.

In terms of the model itself, there are a couple of ideas and concepts that may be included. A straightforward addition would be to model sensor error as a function of spatial regions or investigation objectives, i.e. situations where measurements are unfeasible or increasingly difficult due to the inherent properties of a location. This would incorporate a new trade-off between exploring highly unknown locations and obtaining useful measurements.

One of the fundamental assumptions of the model is the existence of independent spatial regions. Although this helps to simplify the problem and is consistent with the scientist's mental process of analysis and interpretability, especially in the geologic exploration scenario, it would be interesting to incorporate some sort of spatial correlation between regions. For example, using a Gaussian kernel or a spectral similarity criterion with the purpose of propagating the inference process throughout the whole map whenever a new measurement is collected. This may be done by incorporating low resolution orbital data into the probabilistic model, as they are typically available prior to exploration missions.

Despite the fact that diverse and challenging initial hypotheses were tested, an appropriate next step would be to incorporate initial and revised hypotheses by a scientist on multiple sequential command cycles, as well as to evaluate and characterize the evolution of this process.

In terms of path planning, the realism of the constraints can be improved. The experiments were carried out with straightforward path-length planning, but realistic time and energy constraints in conjunction with digital elevation models (DEM) of the site may be included. Additionally, a Monte Carlo planner such as the one by Arora et al [5] may be implemented and compared with the rest of the algorithms. Although it is driven by a stochastic process that does not guarantee optimality nor convergence if the number of iterations is not large enough, it is an anytime, scalable, and flexible algorithm, especially when attempting to investigate various scenarios and constraints.

After most of these ideas have been implemented and tested in simulations, the next logical step would be to validate them in a real field experiment. Most likely, these tests will involve taking Zoë (the autonomous rover from the Field Robotics Center at Carnegie Mellon University)

to Cuprite, Nevada in order to perform an in situ geologic mapping of the site.

And finally, geologic exploration is just one area of application of the science hypothesis map. Another interesting and feasible proof of concept would be a coral reef exploration scenario using unmanned underwater vehicles (UUV), which could also be assisted or simulated with existing remote sensing data products [37].

Bibliography

[1] Radhakrishna Achanta, Appu Shaji, Kevin Smith, Aurelien Lucchi, Pascal Fua, and Sabine Süsstrunk. SLIC superpixels compared to state-of-the-art superpixel methods. *IEEE Transactions on Pattern Analysis Machine Intelligence*, 34(11):2274–2282, 2012. 2.2

[2] National Aeronautics and Space Agency. Curiosity at work on mars (artist's concept). `https://mars.jpl.nasa.gov/msl/multimedia/images/?imageid=3710`, 2017. Accessed: 2017-08-06. (document), 1.1

[3] National Aeronautics and Space Agency. Mars exploration: Program & missions, science, geology. `https://mars.nasa.gov/programmissions/science/goal3/`, 2017. Accessed: 2017-08-06. 4

[4] Environmental Protection Agency: AirNow. Air quality index (AQI) basics. `https://airnow.gov/index.cfm?action=aqibasics.aqi`, 2017. Accessed: 2017-08-05. 1

[5] Akash Arora, Robert Fitch, and Salah Sukkarieh. Extending autonomy of planetary rovers by encoding geological knowledge in a bayesian framework. In *10th International Cognitive Robotics Workshop*, 2016. 2.1, 3.3, 5.3, 7.2

[6] J.M. Bernardo. Expected information as expected utility. *The Ann. of Stat.*, 7(3):686–690, 1979. 5.1

[7] J. Binney and G. S. Sukhatme. Branch and bound for informative path planning. In *2012 IEEE International Conference on Robotics and Automation*, pages 2147–2154, May 2012. doi: 10.1109/ICRA.2012.6224902. 5.1, 5.3, 2

[8] J. Binney, A. Krause, and G. S. Sukhatme. Informative path planning for an autonomous underwater vehicle. In *2010 IEEE International Conference on Robotics and Automation*, pages 4791–4796, May 2010. doi: 10.1109/ROBOT.2010.5509714. 1, 2.1, 5.1

[9] A. J. Brown. Spectral curve fitting for automatic hyperspectral data analysis. *IEEE Transactions on Geoscience and Remote Sensing*, 44(6):1601–1608, June 2006. ISSN 0196-2892.

doi: 10.1109/TGRS.2006.870435. 2.2

[10] A. Castano, A. Fukunaga, J. Biesiadecki, L. Neakrase, P. Whelley, R. Greeley, M. Lemmon, R. Castano, and S. Chien. Autonomous detection of dust devils and clouds on mars. In *2006 International Conference on Image Processing*, pages 2765–2768, Oct 2006. doi: 10.1109/ICIP.2006.313120. 2.1

[11] Kathryn Chaloner and Isabella Verdinelli. Bayesian experimental design: a review. *Statistical Science*, 10(3):273–304, 1995. 1, 5.1

[12] Chein-I Chang. Spectral information divergence for hyperspectral image analysis. In *Geoscience and Remote Sensing Symposium, 1999. IGARSS '99 Proceedings. IEEE 1999 International*, volume 1, pages 509–511 vol.1, 1999. doi: 10.1109/IGARSS.1999.773549. 2.2

[13] B. Charrow, S. Liu, V. Kumar, and N. Michael. Information-theoretic mapping using cauchy-schwarz quadratic mutual information. In *2015 IEEE International Conference on Robotics and Automation (ICRA)*, pages 4791–4798, May 2015. doi: 10.1109/ICRA.2015.7139865. 3.3

[14] Chandra Chekuri and Martin Pal. A recursive greedy algorithm for walks in directed graphs. In *Proceedings of the 46th Annual IEEE Symposium on Foundations of Computer Science*, FOCS '05, pages 245–253, Washington, DC, USA, 2005. IEEE Computer Society. ISBN 0-7695-2468-0. doi: 10.1109/SFCS.2005.9. URL `http://dx.doi.org/10.1109/SFCS.2005.9`. 5.3, 3

[15] Steve Chien, Rob Sherwood, Daniel Tran, Benjamin Cichy, Gregg Rabideau, Rebecca Castano, Ashley Davis, and Darrell Boyer. Using autonomy flight software to improve science return on earth observing one. *Journal of Aerospace Computing, Information, and Communication*, 2:196–216, 2005. 2.1

[16] R.N. Clark et al. Usgs digital spectral library splib06a. *U.S. Geological Survey, Digital Data Series 231*, 2007. (document), 4.1

[17] Roger N Clark. Spectroscopy of rocks and minerals, and principles of spectroscopy. *Manual of Remote Sensing*, 3:3–58, 1999. 2.2, 2.2, 4

[18] Roger N. Clark, Gregg A. Swayze, K. Eric Livo, Raymond F. Kokaly, Steve J. Sutley, J. Brad Dalton, Robert R. McDougal, and Carol A. Gent. Imaging spectroscopy: Earth and planetary remote sensing with the usgs tetracorder and expert systems. *Journal of Geophysical Research: Planets*, 108(E12):n/a–n/a, 2003. ISSN 2156-2202. doi: 10.1029/2002JE001847. URL `http://dx.doi.org/10.1029/2002JE001847`. 5131. 2.2

[19] Thomas M. Cover and Joy A. Thomas. *Elements of Information Theory*. Wiley-Interscience, Hoboken, New Jersey, 2nd edition, 2006. 5.1

[20] Wuhui Duan, Shutao Li, and Leyuan Fang. *Spectral-Spatial Hyperspectral Image Classification Using Superpixel and Extreme Learning Machines*, pages 159–167. Springer Berlin Heidelberg, Berlin, Heidelberg, 2014. ISBN 978-3-662-45646-0. doi: 10.1007/978-3-662-45646-0_17. URL https://doi.org/10.1007/978-3-662-45646-0_17. 2.2

[21] Tara A. Estlin, Benjamin J. Bornstein, Daniel M. Gaines, Robert C. Anderson, David R. Thompson, Michael Burl, Rebecca Castaño, and Michele Judd. Aegis automated science targeting for the mer opportunity rover. *ACM Trans. Intell. Syst. Technol.*, 3(3):50:1–50:19, May 2012. ISSN 2157-6904. doi: 10.1145/2168752.2168764. URL http://doi.acm.org/10.1145/2168752.2168764. 1, 2.1, 2.1

[22] Fan Fan, Yong Ma, Chang Li, Xiaoguang Mei, Jun Huang, and Jiayi Ma. Hyperspectral image denoising with superpixel segmentation and low-rank representation. *Information Sciences*, 397:48 – 68, 2017. ISSN 0020-0255. doi: http://dx.doi.org/10.1016/j.ins.2017.02.044. URL http://www.sciencedirect.com/science/article/pii/S0020025517305443. 2.2

[23] M. L. Fisher, G. L. Nemhauser, and L. A. Wolsey. *An analysis of approximations for maximizing submodular set functions—II*, pages 73–87. Springer Berlin Heidelberg, Berlin, Heidelberg, 1978. ISBN 978-3-642-00790-3. doi: 10.1007/BFb0121195. URL https://doi.org/10.1007/BFb0121195. 5.3

[24] Greydon Foil. *Efficiently Sampling from Underlying Physical Models*. PhD thesis, Robotics Institute, Carnegie Mellon University, Pittsburgh, PA, October 2016. 2.1

[25] Greydon Foil, David R Thompson, William Abbey, and David S Wettergreen. Probabilistic surface classification for rover instrument targeting. In *Intelligent Robots and Systems (IROS), 2013 IEEE/RSJ International Conference on*, pages 775–782. IEEE, 2013. 2.1

[26] M. Gallant, A. Ellery, and J. Marshall. Science-influenced mobile robot guidance using bayesian networks. In *2011 24th Canadian Conference on Electrical and Computer Engineering(CCECE)*, pages 001135–001139, May 2011. doi: 10.1109/CCECE.2011.6030639. 2.1

[27] Y. Girdhar and G. Dudek. Modeling curiosity in a mobile robot for long-term autonomous exploration and monitoring. *Autonomous Robots*, 40(7):1267–1278, 2016. 2.1

[28] L. Hamlin, R.O. Green, P. Mouroulis, M. Eastwood, D. Wilson, M. Dudik, and C. Paine.

Imaging spectrometer science measurements for terrestrial ecology: Aviris and new developments. In *Aerospace Conference, 2011 IEEE*, pages 1–7. IEEE, 2011. 4.2

[29] David Haussler and Manfred Opper. General bounds on the mutual information between a parameter and n conditionally independent observations. In *Proceedings of the Eighth Annual Conference on Computational Learning Theory*, COLT '95, pages 402–411, New York, NY, USA, 1995. ACM. ISBN 0-89791-723-5. doi: 10.1145/225298.225347. URL http://doi.acm.org/10.1145/225298.225347. 5.2

[30] Andrew N. Hock, Nathalie A. Cabrol, James M. Dohm, Jennifer Piatek, Kim Warren-Rhodes, Shmuel Weinstein, David S. Wettergreen, Edmond A. Grin, Jeffrey Moersch, Charles S. Cockell, Peter Coppin, Lauren Ernst, Gregory Fisher, Craig Hardgrove, Lucia Marinangeli, Edwin Minkley, Gian Gabriele Ori, Alan Waggoner, Mike Wyatt, Trey Smith, David Thompson, Michael Wagner, Dominic Jonak, Kristen Stubbs, Geb Thomas, Erin Pudenz, and Justin Glasgow. Life in the atacama: A scoring system for habitability and the robotic exploration for life. *Journal of Geophysical Research: Biogeosciences*, 112(G4):n/a–n/a, 2007. ISSN 2156-2202. doi: 10.1029/2006JG000321. URL http://dx.doi.org/10.1029/2006JG000321. G04S08. 1, 2.1

[31] J. P. Hoffbeck and D. A. Landgrebe. Covariance matrix estimation and classification with limited training data. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 18 (7):763–767, Jul 1996. ISSN 0162-8828. doi: 10.1109/34.506799. 4.2

[32] Geoffrey A. Hollinger and Gaurav S. Sukhatme. Sampling-based robotic information gathering algorithms. *The International Journal of Robotics Research*, 33(9):1271–1287, 2014. doi: 10.1177/0278364914533443. URL http://dx.doi.org/10.1177/0278364914533443. 5.3

[33] Andreas Krause, Carlos Guestrin, Anupam Gupta, and Jon Kleinberg. Near-optimal sensor placements: Maximizing information while minimizing communication cost. In *Proceedings of the 5th International Conference on Information Processing in Sensor Networks*, IPSN '06, pages 2–10, New York, NY, USA, 2006. ACM. ISBN 1-59593-334-4. doi: 10.1145/1127777.1127782. URL http://doi.acm.org/10.1145/1127777.1127782. 5.1, 5.3

[34] Andreas Krause, Ajit Singh, and Carlos Guestrin. Near-optimal sensor placements in gaussian processes: Theory, efficient algorithms and empirical studies. *J. Mach. Learn. Res.*, 9:235–284, June 2008. ISSN 1532-4435. URL http://dl.acm.org/citation.cfm?id=1390681.1390689. 5.3, 5.3

[35] F.A. Kruse, A.B. Lefkoff, J.W. Boardman, K.B. Heidebrecht, A.T. Shapiro, P.J. Barloon,

and A.F.H. Goetz. The spectral image processing system (sips): interactive visualization and analysis of imaging spectrometer data. *Remote Sensing of Environment*, 44 (2):145 – 163, 1993. ISSN 0034-4257. doi: http://dx.doi.org/10.1016/0034-4257(93)90013-N. URL http://www.sciencedirect.com/science/article/pii/003442579390013N. Airborne Imaging Spectrometry. (document), 2.2, 2.2, 2.4

[36] Fred A. Kruse. Expert system analysis of hyperspectral data. In *Proc. SPIE*, volume 6966, pages 69660Q–69660Q–12, 2008. doi: 10.1117/12.767554. URL http://dx.doi.org/10.1117/12.767554. 2.2

[37] NASA Jet Propulsion Laboratory. COral Reef Airborne Laboratory (CORAL). https://coral.jpl.nasa.gov/about-coral, 2017. Accessed: 2017-08-05. 7.2

[38] David A Landgrebe. *Signal Theory Methods in Multispectral Remote Sensing*. Wiley, 2003. 2.2, 2.2

[39] J.K. Lein. *Environmental Sensing: Analytical Techniques for Earth Observation*. Springer, 1st edition, 2012. 2.2, 2.2, 2.2, 2.2

[40] T.M. Lillesand and R. W. Kiefer. *Remote Sensing and Image Interpretation*. John Wiley & Sons, 4th edition, 2002. ISBN 9971-51-427-3. 4.2

[41] D.V. Lindley. On the measure of information provided by an experiment. *The Annals of Mathematical Statistics*, 27(4):986–1005, 1956. 5.1

[42] Alan Mazer, Miki Martin, Meemong Lee, and Jerry E. Solomon. Image processing software for imaging spectrometry data analysis. *Remote Sensing of Environment*, 24, 03 1988. 2.2

[43] Eldar Noe Dobrea. personal communication, 2015. 3.1

[44] L. Pedersen, M. Wagner, D. Apostolopoulos, and W. R. Whittaker. Autonomous robotic meteorite identification in antarctica. In *Proceedings 2001 ICRA. IEEE International Conference on Robotics and Automation (Cat. No.01CH37164)*, volume 4, pages 4158–4165 vol.4, 2001. doi: 10.1109/ROBOT.2001.933268. 2.1, 2.1

[45] H. Petersson, D. Gustafsson, and D. Bergstrom. Hyperspectral image analysis using deep learning - a review. In *2016 Sixth International Conference on Image Processing Theory, Tools and Applications (IPTA)*, pages 1–6, Dec 2016. doi: 10.1109/IPTA.2016.7820963. 2.2

[46] Andrew N. Rencz and Robert A. Ryerson. *Manual of Remote Sensing, Volume 3, Remote Sensing for the Earth Sciences*. Wiley, 3rd edition, March 1999. ISBN 978-0-471-29405-4. (document), 2.3

[47] F.J. Richards. A flexible growth function for empirical use. *Journal of Experimental*

Botany, 10(2):290–301, 1959. 5.2

[48] F. C. Robey, D. R. Fuhrmann, E. J. Kelly, and R. Nitzberg. A cfar adaptive matched filter detector. *IEEE Transactions on Aerospace and Electronic Systems*, 28(1):208–216, Jan 1992. ISSN 0018-9251. doi: 10.1109/7.135446. 2.2

[49] K. Ryan. Estimating expected information gains for experimental designs with application to the random fatigue-limit model. *Journal of Computational and Graphical Statistics*, 12 (3):585–603, 2003. 5.2

[50] Amarjeet Singh, Andreas Krause, and William J. Kaiser. Nonmyopic adaptive informative path planning for multiple robots. In *Proceedings of the 21st International Joint Conference on Artificial Intelligence*, IJCAI'09, pages 1843–1850, San Francisco, CA, USA, 2009. Morgan Kaufmann Publishers Inc. URL `http://dl.acm.org/citation.cfm?id=1661445.1661741`. 5.1, 5.3, 3

[51] Noah A. Smith and Roy W. Tromble. Sampling uniformly from the unit simplex. Technical Report 29, Johns Hopkins Univ., Baltimore, MD, 2004. 5.2

[52] United States Geological Survey. Geologic map of mars. `https://www.usgs.gov/media/images/geologic-map-mars`, 2017. Accessed: 2017-08-06. (document), 3.1

[53] Gregg A. Swayze, Roger N. Clark, Alexander F.H. Goetz, K. Eric Livo, George N. Breit, Fred A. Kruse, Stephen J. Sutley, Lawrence W. Snee, Heather A. Lowers, James L. Post, et al. Mapping advanced argillic alteration at cuprite, nevada, using imaging spectroscopy. *Economic Geology*, 109(5):1179–1221, 2014. (document), 4, 4.1, 4.2

[54] D. R. Thompson, L. Mandrake, M. S. Gilmore, and R. Castano. Superpixel endmember detection. *IEEE Transactions on Geoscience and Remote Sensing*, 48(11):4023–4033, Nov 2010. ISSN 0196-2892. doi: 10.1109/TGRS.2010.2070802. 2.2

[55] David R. Thompson. *Intelligent Mapping for Autonomous Robotic Survey*. PhD thesis, Robotics Institute, Carnegie Mellon University, Pittsburgh, PA, August 2008. 2.1, 3.1, 5.1

[56] David R. Thompson, David S. Wettergreen, and Francisco J. Calderón Peralta. Autonomous science during large-scale robotic survey. *Journal of Field Robotics*, 28(4):542–564, 2011. ISSN 1556-4967. doi: 10.1002/rob.20391. URL `http://dx.doi.org/10.1002/rob.20391`. 1, 1

[57] David R. Thompson, Bo-Cai Gao, Robert O. Green, Dar A. Roberts, Philip E. Dennison, and Sarah R. Lundeen. Atmospheric correction for global mapping spectroscopy: Atrem advances for the hyspiri preparatory campaign. *Remote Sensing of Environment*, 167:64–

77, 2015. 4.2

[58] David Ray Thompson, David Wettergreen, Greydon T Foil, P Michael Furlong, and Anatha Ravi Kiran. Spatio-spectral exploration combining in situ and remote measurements. In *AAAI*, pages 3679–3685, 2015. 2.1

[59] T.N. Tran, R. Wehrens, and L.M.C Buydens. Clustering multispectral images: a tutorial. *Chemometrics and Intelligent Laboratory Systems*, 77:3–17, 2005. 2.2

[60] Jojanneke Van Den Berg, Andrew Curtis, and Jeannot Trampert. Optimal nonlinear bayesian experimental design: an application to amplitude versus offset experiments. *Geophysical Journal International*, 155(2):411–421, 2003. ISSN 1365-246X. doi: 10.1046/j.1365-246X.2003.02048.x. URL `http://dx.doi.org/10.1046/j.1365-246X.2003.02048.x`. 5.2

[61] David Wettergreen. personal communication, 2016. 3.1

[62] Mark Woods, Andy Shaw, Dave Barnes, Dave Price, Derek Long, and Derek Pullan. Autonomous science for an exomars rover-like mission. *Journal of Field Robotics*, 26(4): 358–390, 2009. ISSN 1556-4967. doi: 10.1002/rob.20289. URL `http://dx.doi.org/10.1002/rob.20289`. 2.1

[63] N. K. Yilmaz, C. Evangelinos, P. F. J. Lermusiaux, and N. M. Patrikalakis. Path planning of autonomous underwater vehicles for adaptive sampling using mixed integer linear programming. *IEEE Journal of Oceanic Engineering*, 33(4):522–537, Oct 2008. ISSN 0364-9059. doi: 10.1109/JOE.2008.2002105. 5.3