

Bayesian Models for Science-Driven Robotic Exploration

Alberto Candela Garza

CMU-RI-TR-21-68

September 2021

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

Thesis Committee:

David Wettergreen, Chair

Nathan Michael

Jeff Schneider

David R. Thompson, Jet Propulsion Laboratory

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

Copyright © 2021 Alberto Candela Garza

Abstract

Planetary rovers allow for science investigations in remote environments. They have traversed many kilometers and made major scientific discoveries. However, rovers spend a considerable amount of time awaiting instructions from mission control. The reason is that they are designed for highly supervised data collection, not for autonomous exploration. The exploration of farther worlds will face increasing challenges and constraints. Such missions will demand a new approach.

This work advocates Bayesian models as powerful tools for a new paradigm for robotic explorers. In this approach, the explorer's description of where to go is not prescribed by a fixed set of instructions, but instead by a model of what the explorer believes. This formulation has several benefits. Bayesian models provide a mathematically grounded framework to reason about uncertainty. They can allow robots to gain a deeper understanding of the evolving scientific goals guiding the mission. Furthermore, they can empower scientists by providing explainable results.

To this end, this research develops models that allow for data interpretation by learning and exploiting structure in the data and the environment. It shows how these models enable robotic explorers to make intelligent decisions based on instantaneous information. Ultimately, it demonstrates how science productivity is improved by measuring science value with information-theoretic variables and by formulating the exploration problem in terms of Bayesian experimental design.

This work makes several contributions to the field of science-driven robotic exploration. First, it introduces three different deep generative models for the analysis of data that enable scientists to quantify and interpret learned statistical dependencies. Then, it presents an adaptive exploration model that leverages contextual information from remote data to efficiently extrapolate features from *in situ* observations, as well as a corresponding strategy for improving science productivity. Afterward, it establishes a hierarchical probabilistic structure in which scientists initially describe their abstract beliefs and hypotheses, and then this belief evolves as the robot makes raw measurements; additionally, science information gain is efficiently computed and maximized. Finally, it proposes a comprehensive model for planetary rover exploration that considers both science productivity and risk.

The presented Bayesian models are validated and evaluated in various science investigation scenarios that can provably benefit from autonomous robotic exploration. Such scenarios include terrestrial, airless, and Martian surface surveys, as well as marine biology studies. Emphasis is placed on spectroscopic data, which is widely used in the natural sciences for composition analysis. Promising results are shown in simulations and field experiments using the autonomous rover Zoë.

Acknowledgments

This thesis was only possible through the help of many others, so I would like to dedicate a few words to acknowledge their support.

First and foremost, I owe a great debt of gratitude to my advisor David Wettergreen. From the first time we met, he knew right away that this dissertation topic would suit my skills and interests. It has been a tremendous honor and pleasure to work under his guidance. I will be forever grateful for his advice and support over the years, not only in terms of research but also in life.

I would like to thank David R. Thompson since working under his mentorship has been an incomparable experience. I have learned such a great deal by his example. Furthermore, he has granted me the opportunity to work at the Jet Propulsion Laboratory for several occasions, for which I am very thankful.

Thanks to Nathan Michael and Jeff Schneider, my committee members, for taking the time to read this thesis, as well as for providing valuable comments that ultimately led to improvements to this work.

I would like to extend my gratitude to Eldar Noe Dobrea, Roger Clark, Tom Prettyman, and Neil Pearson from the Planetary Science Institute for their invaluable help regarding geology and planetary science.

I also want to thank Michelle Gierach Gail Woodward, and David R. Thompson (again) from the Jet Propulsion Laboratory. This dissertation was originally going to focus exclusively on planetary science and rovers, but thanks to them we were able to broaden its applications to coral reef studies.

Thanks to my Carnegie Mellon lab mates: Greydon Foil, Shivam Gautam, Bishwamoy Roy, Suhit Kodgule, Himanshi Yadav, Kevin Edelson, Srini Vijayarangan, Abby Breitfeld, and Maggie Hansen. They have been excellent collaborators, officemates, and adventure companions.

I would like to give a very special thanks to Sharon for her patience, encouragement, and endless love. She has been my greatest source of motivation, as well as my biggest supporter through the best and worst episodes.

Finally, this work was funded by a CONACYT scholarship, a Fulbright-García Robles scholarship, a SEP-DGRI scholarship, a Lewis and Clark Grant Fund for Exploration and Field Research in Astrobiology, a Jet Propulsion Laboratory Strategic University Research Partnership Grant, the National Science Foundation National Robotics Initiative Grant #IIS-1526667, the NASA Solar System Exploration Virtual Institute and the Planetary Science Institute project TREX: Toolbox for Research and Exploration (award 80ARC017M0005), and the NASA Astrobiology Institute Cycle 7 (SETI Team). Their support is gratefully acknowledged.

Contents

1	Introduction	1
1.1	Motivation	1
1.2	Bayesian Model of Exploration	3
1.3	Thesis Statement	5
1.4	Contributions	5
1.5	Thesis Overview	7
2	Methods of Bayesian Modeling for Analysis, Inference, and Decision-Making	9
2.1	Bayesian Inference	9
2.2	Uncertainty Quantification	10
2.3	Information Theory	12
2.4	Bayesian Experimental Design	14
3	Spectroscopic Data and Analysis Techniques	17
3.1	Spectroscopic Data	17
3.2	Analysis Techniques	18
4	Scenarios of Science-Driven Robotic Exploration	23
4.1	Robotic Platform Zoë	23
4.2	Geologic Surveys	24
4.2.1	Terrestrial Surface Investigations	25
4.2.2	Airless Surface Investigations	26
4.2.3	Martian Surface Investigations	29
4.3	Coral Reef Studies	31
4.4	Summary	33
5	Deep Generative Models for Spectroscopic Data Analysis	35
5.1	Introduction	35
5.2	Related Work	36
5.3	Variational Autoencoder	38

5.3.1	Model and Architecture	38
5.3.2	Experiments and Results	40
5.4	Deep Conditional Gaussian Model	41
5.4.1	Model and Architecture	41
5.4.2	Experiments and Results	46
5.5	Deep Conditional Dirichlet Model	55
5.5.1	Model and Architecture	55
5.5.2	Experiments and Results	57
5.6	Discussion and Conclusion	65
6	Combining Remote and In Situ Measurements for Adaptive Exploration	67
6.1	Introduction	67
6.2	Related Work	68
6.3	Probabilistic Mapping Model	69
6.3.1	Spectral Feature Extraction	70
6.3.2	Gaussian Processes for Spatio-Spectral Regression	71
6.3.3	Learning-Based Spectral Composition Analysis	72
6.3.4	Integrated Model	73
6.4	Informative Exploration	73
6.5	Cuprite Scenario	76
6.5.1	Experimental Setting	76
6.5.2	Simulation Experiments and Results	77
6.5.3	Field Experiments and Results	78
6.6	Mars Scenario	83
6.6.1	Experiments	83
6.6.2	Results	85
6.7	Coral Scenario	88
6.7.1	Experiments	88
6.7.2	Results	93
6.8	Discussion and Conclusion	95
7	The Science Hypothesis Map for Co-Exploration	99
7.1	Introduction	99
7.2	Related Work	100
7.3	Probabilistic Mapping Model	101
7.3.1	Science Hypothesis Map	101
7.3.2	Bayesian Inference	102
7.3.3	Geologic Model	102
7.4	Informative Exploration	103

7.5	Cuprite Scenario	106
7.5.1	Geologic Model	106
7.5.2	Experiments and Results	107
7.6	TREX Scenario	113
7.6.1	Geologic Model	113
7.6.2	Experiments and Results	114
7.7	Discussion and Conclusion	120
8	Science and Risk-Aware Exploration	123
8.1	Introduction	123
8.2	Related Work	124
8.2.1	Science-Aware Planning	124
8.2.2	Risk-Aware Planning	124
8.3	Approach	125
8.3.1	Science Model	126
8.3.2	Risk Model	126
8.3.3	Path Planning Formulation	131
8.4	Experiments and Results	133
8.4.1	Experiments	133
8.4.2	Results	134
8.5	Discussion and Conclusion	140
9	Conclusion	141
9.1	Contributions	141
9.2	Future Work	143
9.3	Closing Remarks	145
	Bibliography	147

Chapter 1

Introduction

1.1 Motivation

Rovers are the means to better understanding Mars (Figure 1.1). Over the past two decades, robots have provided transformative information about the Martian geological present and its ancient past. With their ability to move around and investigate multiple locations, rovers have tremendous advantages over static landers. By investigating many sites and correlating observations, Mars rovers have provided compelling evidence that the planet was once more habitable than it is today [66]. More missions are scheduled for the near future and they will carry out new exciting tasks. For instance, NASA’s Perseverance rover is currently caching samples that will one day be brought back to Earth.

In the next few years, rovers will be critical for studying the Moon. Whether ice exists in useful amounts is one of the most pressing questions today since ice could eventually be harvested to sustain human exploration. MoonRanger, a small robot being developed by Carnegie Mellon University and Astrobotic Technologies, will be the first rover to search for signs of water at the South Pole of the Moon in 2023 [23]. NASA will follow up these efforts with its more capable Volatiles Investigating Polar Exploration Rover (VIPER), which will collect about 100 days of data using four different science instruments. Its goal will be to perform more rigorous and sustained scientific characterization of the ice at the South Pole [130].

Modern rovers may be described as semi-autonomous agents that are designed for data collection [55]. These rovers are guided by scientists directly selecting a strategy that they believe will lead to greatest discovery and success. The scientists’ instructions are formed from their expert knowledge of the area and expectations as to where and how to best gather information. Human scientists continually reinterpret their measurements with a growing contextual knowledge, contrary to existing experimental design formalisms, real missions are characterized by a frequent reformulation and replanning throughout the investigation [79].

Achieving high levels of productivity for planetary rover operations can be difficult [51]. These missions are limited in bandwidth, delayed due to distance and restricted to just a few




Figure 1.1: Artist's concept of Curiosity, the NASA rover deployed on Mars. It first travels to locations of interest. It then uses its ChemCam instrument together with other devices to investigate the composition of rocks as scientists try to understand the presence of water on ancient Mars. Image courtesy of NASA [2].

communication cycles per day. Rovers spend a considerable amount of time awaiting instructions from ground operators that require substantial effort when interpreting data and planning commands. For example, a study using data from the Mars Science Laboratory mission revealed that activity plans were conservative since they completed on average 28% faster than predicted [50].

A few automated science data collection procedures have been developed as these are a way to alleviate the aforementioned communications and operations bottleneck [43, 180]. The Autonomous Exploration for Gathering Increased Science (AEGIS) system analyzes rover images onboard Mars rovers to detect pre-defined science features of interest, enabling targeted instrument data to be acquired immediately with no delays for ground communication [43]. However, AEGIS pursues static objectives that are fixed at the outset, providing little adaptability as the knowledge of the field area evolves with incoming data.

The exploration of farther celestial bodies will require more adaptive and sustained autonomy. Ocean Worlds such as Europa or Enceladus have become key astrobiology targets because scientists believe they possess subsurface oceans covered in ice. However, exploring these worlds will be extremely challenging. It will involve building aquatic robots capable of navigating with virtually no communication under ice sheets that could be up to 20 kilometers thick [92].

Upcoming deep space missions will demand a new paradigm. Robotic agents designed for automated data collection will need to evolve into true robotic explorers that are capable of modeling scientific endeavor itself [55]. Envisioning and addressing the next generation of robotic

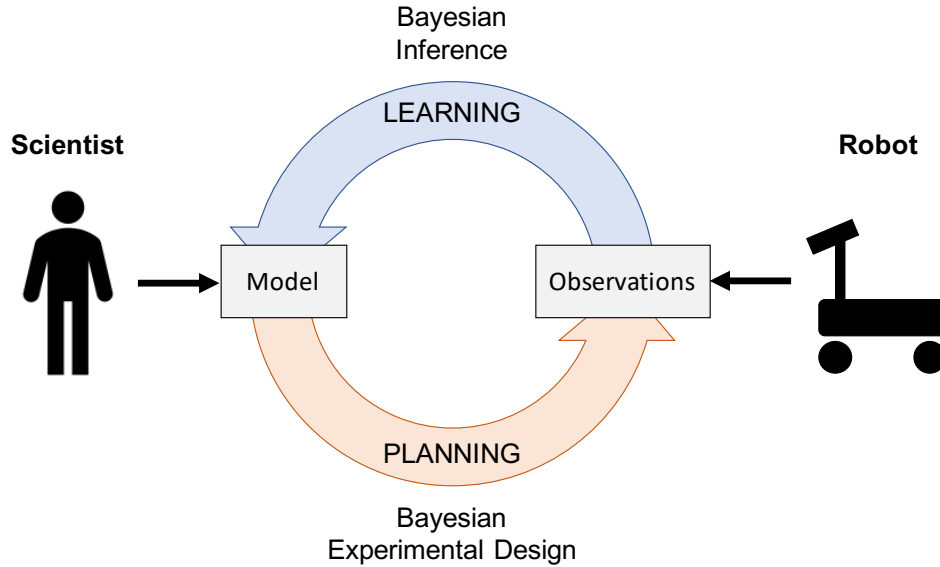


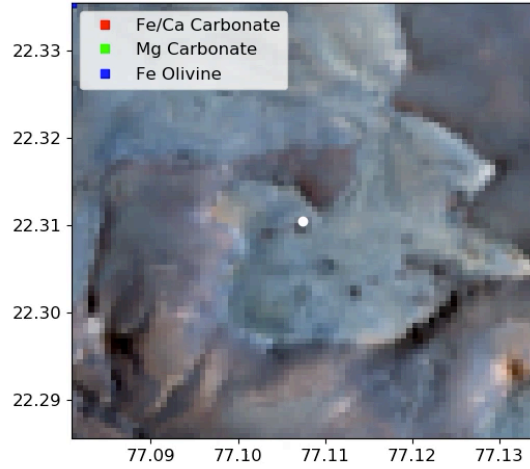
Figure 1.2: Science-driven robotic exploration as a cycle of two processes: Bayesian inference (learning) and Bayesian experimental design (planning). The learning step updates a model of what the explorer believes as new observations occur. The planning step selects new observations that yield the greatest amount of information according to the current state of the model. Scientists can supervise the overall process by defining hypotheses, priors, revisions, and constraints. A robot can continually calculate navigation plans that optimize scientific discovery.

explorers is what motivates this research.

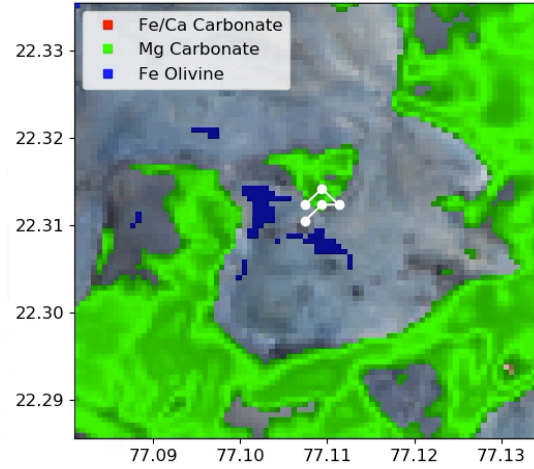
1.2 Bayesian Model of Exploration

This work advocates Bayesian models as powerful tools for creating the next generation of science-driven robotic explorers. In this research, the explorer’s description of where to go is not prescribed by a fixed set of instructions, but instead by a model of what the explorer believes and that is able to adapt to real-time information. This formulation has several benefits. It allows a robot to gain a deeper understanding of the evolving scientific goals guiding the missions. Moreover, it can empower scientists by providing explainable results.


This research formulates adaptive robotic exploration as a cycle of two processes: Bayesian inference and Bayesian experimental design (Figure 1.2). The Bayesian inference step is concerned with how to update the model of what the explorer believes given new observations, which is learning. The Bayesian experimental design step is focused on selecting new observations that yield the greatest amount of quantified information given the current state of the model, which is planning. Scientists can supervise the overall process by defining hypotheses, priors, revisions, and constraints. In real time, a robot can continually calculate navigation plans that maximize scientific discovery.




(a) In this case the robot begins with little or no prior knowledge of the explored environment.



(b) The robot identifies, extrapolates, and maps two mineral signatures (corresponding to Mg carbonate and Fe olivine) using a Bayesian spatial model.



(c) The robot collects more samples in order to update and refine the mineral map, especially when mapping Fe olivine.



(d) The rover discovers and maps a third mineral (Fe/Ca carbonate) in the scene.

Figure 1.3: Conceptual example of an autonomous robotic explorer performing a geologic survey near Nili Fossae on Mars.

An example that illustrates how a robotic exploration mission can be formulated in terms of Bayesian inference and Bayesian experimental design serves as a preview of this research. Imagine a mission where a rover is characterizing and mapping the surface mineralogy of a location near Nili Fossae on Mars (Figure 1.3). Minerals such as carbonates and olivines have been found in Nili Fossae, making it a location of high astrobiological interest [173]. Carbonates are minerals that contain a carbonate ion, and are of interest because they are typically formed

through the interaction of carbon dioxide and liquid water. This makes carbonate deposits of interest to astrobiologists who are trying to determine aspects of the ancient environment of Mars. Olivines are a group of minerals that are associated with volcanic activity. Olivines are silicates containing varying amounts of magnesium (Mg) and iron (Fe). A process known as hydrothermal serpentinization produces carbonates in the presence of olivines [16]. In order to identify such minerals, a Martian rover carries special instruments that are used to analyze the composition of samples. In adaptive science-driven robotic exploration, the rover not only is able to identify key minerals in the samples, but its model of the explored environment is updated accordingly. Spatial models can be especially powerful for exploration since maps are an inherent aspect of many scientific endeavors [162]. This step corresponds to the Bayesian inference process. After the model is updated, the rover collects new observations that are linked to minerals that are yet to be found in the scene. This step corresponds to the Bayesian experimental design process, where the designed experiment is the path or series of mineral observations that maximize diversity, and hence information from the explored scene. In summary, these Bayesian steps guide the robot and allow it to make intelligent decision based on an evolving knowledge of the area throughout the mission.

1.3 Thesis Statement

Formally, this research advocates the following conclusion:

Bayesian models can enable and improve robotic exploration for science investigations.

Specifically, this work focuses on demonstrating the following statements:

- Bayesian models can enable data interpretation by learning and exploiting structure in the data and the environment.
- Bayesian models can enable decision making by helping plan intelligent actions based on current information.
- Bayesian models can improve science productivity as measured by formal principles of information theory and Bayesian experimental design.

1.4 Contributions

This research contains significant contributions to the field of science-driven robotic exploration. In general, it demonstrates how Bayesian models can be used for science data interpretation and intelligent decision-making; and ultimately for improving science return.

To this end, this work constructs and studies a series of relevant science investigation scenarios that can provably benefit from autonomous robotic exploration. These scenarios focus on terrestrial, airless, and Martian surface surveys with rovers. Additionally, this work addresses

marine biology studies as they can too be assisted by autonomous robots. Herein we pay special attention to the construction of spatial models for exploration since maps are an inherent aspect of many scientific investigations [162]. Emphasis is placed on observations from spectrometers as they are powerful instruments for estimating the chemical composition of many materials [32].

This research describes three different deep generative models, which combine ideas from Bayesian models and deep neural networks. The objective is to allow scientists to quantify, visualize and interpret learned statistical dependencies while performing data analysis. We focus on spectroscopic data that is relevant for the aforementioned scientific scenarios. We start with the Variational Autoencoder, which is used to learn, extract, and normalize science features. We then describe the Deep Conditional Gaussian Model, a tool for probabilistic regression that performs Gaussian parameter estimation for each different input. We apply this model to the problem of enhancing low-resolution data, also known as super resolution. We then present the Deep Conditional Dirichlet Model, a similar architecture that instead performs Dirichlet parameter estimation. It can be used for classification tasks, and more generally, for predicting proportional abundances of mixed classes. Results are shown when applied to the problem of spectral unmixing.

This work presents a spatial model that leverages contextual information from remote data to efficiently extrapolate features from robot measurements throughout an explored environment. It is based on Gaussian process regression, but enhanced in such a way that is able to map and extrapolate multivariate data, which is latter used to map composition throughout wide areas. Additionally, an objective function is derived for improving science productivity in exploration. Many different sampling strategies are compared. These methods are evaluated both in extensive simulations at diverse locations and field experiments with a rover in Nevada.

This research develops a Bayesian framework in which scientists initially describe their abstract beliefs and hypotheses, then the state of this belief evolves as the robot makes raw measurements. This is possible by constructing a hierarchical probabilistic structure that relates high-level science representations to measurable data. We call this model the *science hypothesis map*. We also derive, efficiently compute, and maximize scientific information gain in the model. This system is tested in simulations where various science-blind and science-aware path planners are compared. Results with a rover are also shown.

Finally, this work addresses a comprehensive strategy for planetary rovers that accounts for science and risk. Most works related to motion planning for rovers focus on either increasing science productivity or reducing traversability risk; but not both at the same time. We present a Bayesian framework for the integration of these two tasks. Specifically, one that is able to quantify information and risk, provide probabilistic guarantees, and design paths accordingly. This approach is evaluated in a Mars surface simulation study using both science and terrain-relevant data.

1.5 Thesis Overview

The organization of this document is as follows.

Chapter 2 Methods of Bayesian Modeling for Analysis, Inference, and Decision-Making covers concepts that are relevant for this research. It explains definitions and well-established principles of Bayesian inference, uncertainty quantification, information theory, and Bayesian experimental design.

Chapter 3 Spectroscopic Data and Analysis Techniques includes a description on spectroscopic data, as well as a discussion on corresponding analysis techniques. The objective of this chapter is to understand how to enable robots to analyze spectroscopic data.

Chapter 4 Scenarios of Science-Driven Robotic Exploration addresses the various science scenarios, instruments, data sets, and tools that are used in this research. Specifically, this chapter describes scientific tasks involving diverse geologic surveys, as well as coral reef studies.

Chapter 5 Deep Generative Models for Spectroscopic Data Analysis starts by discussing the Variational Autoencoder, which is used to learn and normalize science features. It then describes the Deep Conditional Gaussian Model and shows results when applied to the problem of spectral super-resolution. It then presents the Deep Conditional Dirichlet Model and demonstrates results when applied to the problem of spectral unmixing.

Chapter 6 Combining Remote and In Situ Measurements for Adaptive Exploration describes an active model for efficient wide-area mapping and robotic exploration. To this end, this model integrates ideas from deep generative models, Gaussian processes, and learning-based composition analysis.

Chapter 7 The Science Hypothesis Map for Co-Exploration addresses a Bayesian framework and hierarchical probabilistic structure. This model relates high-level science representations to measurable data, allowing scientists to initially describe their abstract beliefs and hypotheses, then enabling the robot to update the state of this belief with new observations.

Chapter 8 Science and Risk-Aware Exploration introduces a framework for planetary rovers that improves science return while accounting for traversability risk. It describes a probabilistic risk model that combines geometric and semantic information.

Chapter 9 Conclusions closes this thesis with a summary of this work and a discussion on its contributions. It also discusses relevant avenues for future research.

Chapter 2

Methods of Bayesian Modeling for Analysis, Inference, and Decision-Making

This research focuses on Bayesian models for one important reason: Bayesian models provide a mathematically grounded framework to reason about knowledge and uncertainty. Consequently, they can allow robots to gain a deeper understanding of the evolving scientific goals guiding their missions. This chapter describes useful concepts that serve as the basis for science analysis and robotic exploration in this research. Specifically, it explains concepts from Bayesian inference, uncertainty quantification, information theory, and Bayesian experimental design.

2.1 Bayesian Inference

Bayesian inference is a method of statistical inference in which the probability for a hypothesis is updated as more information from observations becomes available. In this research we are interested in *science hypotheses*. An example may be the spatial distribution of minerals in a region of Mars, as described earlier in Chapter 1. Such hypotheses need to be updated with *scientific observations* from instruments. These observations may come in the form of images, spectroscopic measurements, thermal readings, etc. Ultimately, the goal of this work is to define the variables and to construct the probability distributions required to formally relate low-level instrument observations to high-level science hypotheses so the latter can be updated with the former. To this end, let us cover the following basic definitions in Bayesian inference:

- $\theta \in \Theta$ is the *hypothesis* or parameter to estimate.
- $x \in X$ is the *observation* or new data.
- $p(\theta)$ is known as the *prior distribution* or prior knowledge. It is the distribution of the hypothesis θ before any data is observed.
- $p(x|\theta)$ is known as the *likelihood* or sampling distribution. It indicates the compatibility of an observation x given the hypothesis θ .

- $p(x)$ is known as the *marginal likelihood* or *evidence*. It is the distribution of the observation x , often marginalized over the possible hypotheses, i.e. $p(x) = \int_{\Theta} p(x|\theta)p(\theta)d\theta$.
- $p(\theta|x)$ is known as the *posterior distribution* or posterior knowledge. It is the updated probability of the hypothesis θ given observation x .

All of these concepts are related to each other in Bayes' theorem, which is given by:

$$p(\theta|x) = \frac{p(x|\theta)p(\theta)}{p(x)}. \quad (2.1)$$

2.2 Uncertainty Quantification

Science exploration seeks to discover the unknown. It requires formal procedures to quantify uncertainty so as to make informed analyses and decisions. In this research we advocate Bayesian models since they provide a mathematically grounded framework to reason about uncertainty. But first, it is important to understand what uncertainty is and how to quantify it.

Uncertainty may be classified into two categories: *aleatoric* and *epistemic* [95, 98]. Aleatoric uncertainty is the intrinsic randomness or noise in the data. This could be for example instrument noise, usually resulting in uncertainty that cannot be reduced even if more data were to be collected. In contrast, epistemic uncertainty accounts for uncertainty in the model and its parameters. This uncertainty can be explained away given enough data; for example, by gathering information from scientific instruments that complement each other, or by getting data from unsampled locations. The Bayesian models that we present in this research will address these two kinds of uncertainty. For instance, our deep generative models (Chapter 3) focus on learning aleatoric uncertainties in the data, while Gaussian processes (Chapter 4) are typically used to model and reduce epistemic uncertainty in certain cases. An example of aleatoric and epistemic uncertainties is shown in Figure 2.1. Note that some data might be redundant and may provide no additional information. A scientifically relevant example would be to take multiple pictures of the same rock at the same location under the same conditions, essentially just repeating the same information. Hence, we point out the notion that some samples are more informative than others; a robotic explorer would like to select them accordingly. Intuitively, they robot should prioritize samples that are different to what it has already seen. This idea will be further explained and formalized when we later discuss information theory and Bayesian experimental design.

Furthermore, there is a distinction between two types of aleatoric uncertainty or noise: *homoscedastic* and *heteroscedastic* (Figure 2.2). Essentially, homoscedastic uncertainty refers to a constant noise function for all inputs, whereas heteroscedastic uncertainty involves variable noise functions. This distinction is important in this research as some of our Bayesian models assume homoscedastic uncertainties, while others are able to represent heteroscedastic noise. It is worth mentioning that both independent and dependent variables, or inputs and outputs, can have aleatoric uncertainties. *Forward uncertainty propagation* is the quantification of uncertainties in

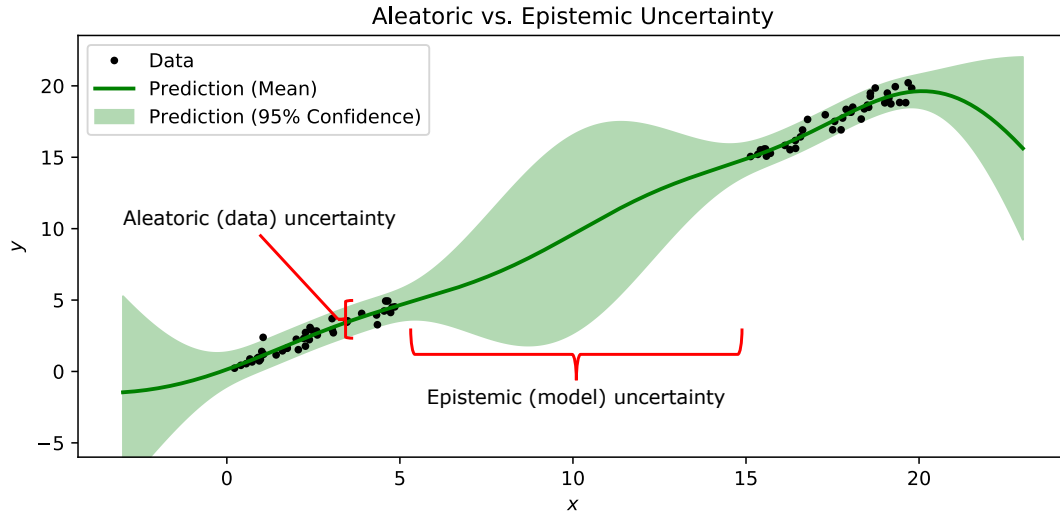


Figure 2.1: Example of aleatoric (data) and epistemic (model) uncertainties. In this example, epistemic uncertainty could be reduced by collecting data in the interval $x = (5, 15)$.

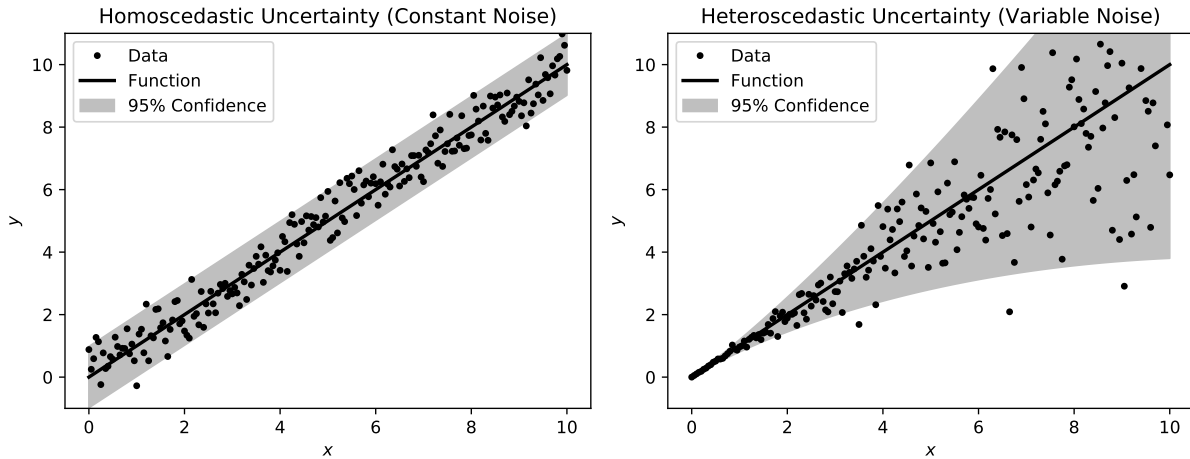


Figure 2.2: Homoscedastic and heteroscedastic uncertainties. Homoscedasticity refers to a constant noise function (left), whereas heteroscedasticity means that noise changes for different values of x and y (right).

system outputs propagated from uncertain inputs. A key insight is that the overall uncertainty in the system response tends to increase where there is a high derivative in the underlying function (Figure 2.3). Intuitively, this means that even a small amount of noise in the input can considerably change the value of the output, and thus increase the associated uncertainty. This insight will prove useful we demonstrate how our work on deep generative models (Chapter 5) and Gaussian processes (Chapter 6) is able to capture this propagation of uncertainties.




Figure 2.3: Forward uncertainty propagation. This example involves a sigmoid function with constant Gaussian noise models for x (input) and y (output). The overall uncertainty in the system response increases at points where the function has a large derivative since small changes in the input can produce large variations in the output.

2.3 Information Theory

Information theory studies the quantification, storage, and communication of information [35]. It is based on probability theory and statistics. In this work we are especially interested in measures of information and uncertainty for random variables describing scientific hypotheses and observations. These measures and properties are powerful tools when applied to our Bayesian models for science exploration. They allow a robotic explorer to quantify information and uncertainty in science hypotheses and observations; they also permit the robot to understand how new observations will affect the science hypotheses guiding its mission. Next we describe the ones that are used in this research.

Shannon entropy is the fundamental unit of measure in information theory. It quantifies the average amount of information I of a probability distribution. It is defined as follows:

$$H(X) = \mathbb{E}_x [I(x)] = - \sum_{x \in X} P(x) \log_b P(x), \quad (2.2)$$

where b is the logarithm base. When $b = 2$, entropy (and information) is measured in bits. When $b = e$, entropy is measured in natural units (nats). We employ the natural logarithm throughout this work when calculating Shannon entropy. Note that $H(X) \geq 0$ as long as $b > 1$. A small entropy value indicates that there is little uncertainty in the distribution; and vice versa, a large entropy value means that there is high uncertainty. An advantage of using Shannon entropy for measuring uncertainty is that its value is always a scalar, even for multivariate distributions.

The *conditional entropy* is defined as the expected entropy of X given Y , that is:

$$H(X|Y) = \mathbb{E}_Y [H(X|y)] = - \sum_{y \in Y} P(y) \sum_{x \in X} P(x|y) \log P(x|y). \quad (2.3)$$

Mutual information quantifies how much information is shared between two random variables. It is defined as follows

$$MI(X; Y) = \sum_{x \in X} \sum_{y \in Y} P(x, y) \log \frac{P(x, y)}{P(x)P(y)}. \quad (2.4)$$

Information gain measures the amount of information that can be obtained about one random variable after observing another. It is defined as the expected reduction in the entropy of X given Y :

$$IG(X|Y) = H(X) - H(X|Y). \quad (2.5)$$

It has been shown that information gain and mutual information are equivalent [35]:

$$IG(X|Y) = MI(X; Y). \quad (2.6)$$

Additionally, we observe that both measures are symmetric and non-negative

$$IG(X|Y) = IG(Y|X) = MI(X; Y) = MI(Y; X) \geq 0. \quad (2.7)$$

The *Kullback-Leibler divergence* (KLD), also known as information divergence or relative entropy, quantifies the difference between two probability distributions $P(X)$ and $Q(X)$. It is defined as follows:

$$D_{KL}(P(X)||Q(X)) = \sum_{x \in X} P(x) \log \frac{P(x)}{Q(x)}. \quad (2.8)$$

Note that this measure is not symmetric, i.e., $D_{KL}(P(X)||Q(X)) \neq D_{KL}(Q(X)||P(X))$. The Kullback-Leibler divergence is often used to evaluate the difference between the “true” or posterior distribution $P(X)$ and an “arbitrary” or prior distribution $Q(X)$. Assuming that $P(X|Y)$ is the posterior and $P(X)$ is the prior, we find that information gain is equivalent to the expected Kullback-Leibler divergence between all possible outcomes for the posterior and the prior

$$IG(X|Y) = \mathbb{E}_Y [D_{KL}(P(X|y)||P(X))]. \quad (2.9)$$

Cross-entropy also measures the difference between two probability distributions $P(X)$ and $Q(X)$. It is given by:

$$H(p, q) = - \sum_{x \in X} P(x) \log Q(x). \quad (2.10)$$

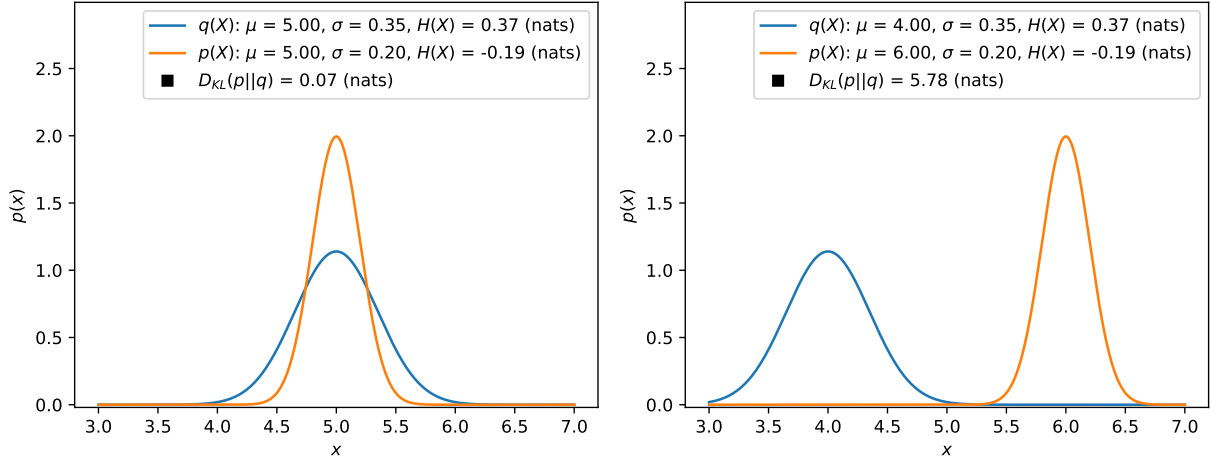


Figure 2.4: Differential entropy and Kullback-Leibler divergence for different Gaussian distributions. Large variances yield large entropies (e.g., $q(X)$). Note that differential entropy can be negative for small variances (e.g., $p(X)$). Similar distributions have a small Kullback-Leibler divergence (left); and vice versa, dissimilar distributions have a large Kullback-Leibler divergence (right).

Cross-entropy is a smooth function that is commonly used when training neural networks for classification tasks. It may also be formulated in terms of entropy and the Kullback-Leibler divergence:

$$H(p, q) = H(P(X)) + D_{KL}(P(X)||Q(X)). \quad (2.11)$$

For continuous distributions, Shannon entropy is called *differential entropy* and replaces summations with integrations

$$H(X) = \mathbb{E}_x [I(x)] = - \int_X p(x) \log p(x) dx. \quad (2.12)$$

The same idea regarding integration applies to all the information measures previously mentioned. Note that differential entropy can have negative values. However, information gain, mutual information, and the Kullback-Leibler divergence are always non-negative for continuous distributions. A relevant example involving Gaussian distributions is shown in Figure 2.4.

Additionally, entropy, mutual information, information gain, and the Kullback-Leibler divergence are additive for independent random variables (discrete and continuous). For instance,

$$X_1 \perp\!\!\!\perp X_2 \quad \Rightarrow \quad H(X_1, X_2) = H(X_1) + H(X_2). \quad (2.13)$$

2.4 Bayesian Experimental Design

Bayesian experimental design provides a probabilistic framework for making decisions under uncertainty [25]. Specifically, it is used to find, or design, the experiment that will be most use-

ful to explain a hypothesis $\theta \in \Theta$. This research focuses on *science experiments* performed by robotic explorers. Particularly, experimental designs where a robot needs to decide which are the best new observations to be collected, that is, the ones that will most significantly reduce the uncertainty of the science hypotheses in the exploration model. The formalism provided by Bayesian experimental design allows robotic explorers to follow principled methods to select such science observations that yield the greatest amount of information gain. To this end, we follow the definitions from the seminal work by Lindley [113]. We start by noting that observations $x \in X$ depend on the chosen experiment ε , i.e., $p(x) = f(\varepsilon)$.

The amount of information of the prior knowledge $p(\theta)$ is quantified using entropy:

$$H(\Theta) = - \int_{\Theta} p(\theta) \log p(\theta) d\theta. \quad (2.14)$$

After the experiment has been performed and a value x has been observed, the amount of information of the posterior knowledge is quantified using conditional entropy:

$$H(\Theta|x) = - \int_{\Theta} p(\theta|x) \log p(\theta|x) d\theta. \quad (2.15)$$

The amount of information provided by the experiment ε when the observation is x , with prior knowledge $p(\theta)$, is given by the reduction in entropy:

$$g(\varepsilon, p(\theta), x) = H(\Theta) - H(\Theta|x). \quad (2.16)$$

There are multiple possible outcomes for the observation x . Therefore, the average amount of information provided by the experiment ε , with prior knowledge $P(\theta)$, is given by the expected reduction in entropy:

$$g(\varepsilon, p(\theta)) = \mathbb{E}_x [H(\Theta) - H(\Theta|x)]. \quad (2.17)$$

Note that this expression is equivalent to information gain (see Equation 2.5):

$$g(\varepsilon, p(\theta)) = H(\Theta) - H(\Theta|X). \quad (2.18)$$

The objective of Bayesian experimental design is to find the optimal experimental design ε^* , which is the one that maximizes information gain:

$$\varepsilon^* = \arg \max_{\varepsilon} g(\varepsilon, p(\theta)). \quad (2.19)$$

Note that the prior knowledge $H(\Theta)$ is affected by neither the experiment nor the observations, so maximizing information gain is the same as minimizing the expected posterior entropy:

$$\varepsilon^* = \arg \min_{\varepsilon} \mathbb{E}_x [H(\Theta|x)]. \quad (2.20)$$

Chapter 3

Spectroscopic Data and Analysis Techniques

This work pays special attention to spectroscopic data; a critical element for Earth and planetary studies. This chapter first explains what spectroscopic data is. It then includes a discussion on common analysis techniques for spectroscopic data. Understanding spectroscopic data analysis techniques, as well as their benefits and drawbacks, is important for our research since we want to enable robotic explorers to process spectroscopic data in a principled and automated manner.

3.1 Spectroscopic Data

Earth and planetary sciences often rely upon the analysis of spectroscopic data. Instruments such as the Compact Reconnaissance Imaging Spectrometer for Mars (CRISM) [128] and the Thermal Emission Imaging System (THEMIS) [30] have been vital for understanding the geology of Mars, as well as for studying climate and habitability implications [12]. Another example is the Moon Mineralogy Mapper (M3) instrument, which produced the first mineralogical map of the Moon and also found water ice in the polar regions [112]. Therefore, herein we place emphasis on spectroscopic data.

In general, spectroscopy is a discipline that studies the interaction between matter and electromagnetic radiation. Spectrometers are instruments that measure different wavelengths of the electromagnetic spectrum and thus capture more information than can be seen with the eye. By observing unique patterns in the reflectance and absorption of light throughout the wavelengths of the electromagnetic spectrum, the chemical composition of an object may be derived, such as the presence of water or certain minerals [32, 111]. The measured signals are called *spectra* and the aforementioned patterns are known as *spectral features*. Herein we focus on *reflectance spectra*, which is the fraction of incident electromagnetic power that is reflected from a surface.

There are many different spectroscopic sensors and techniques, both for remote and *in situ*

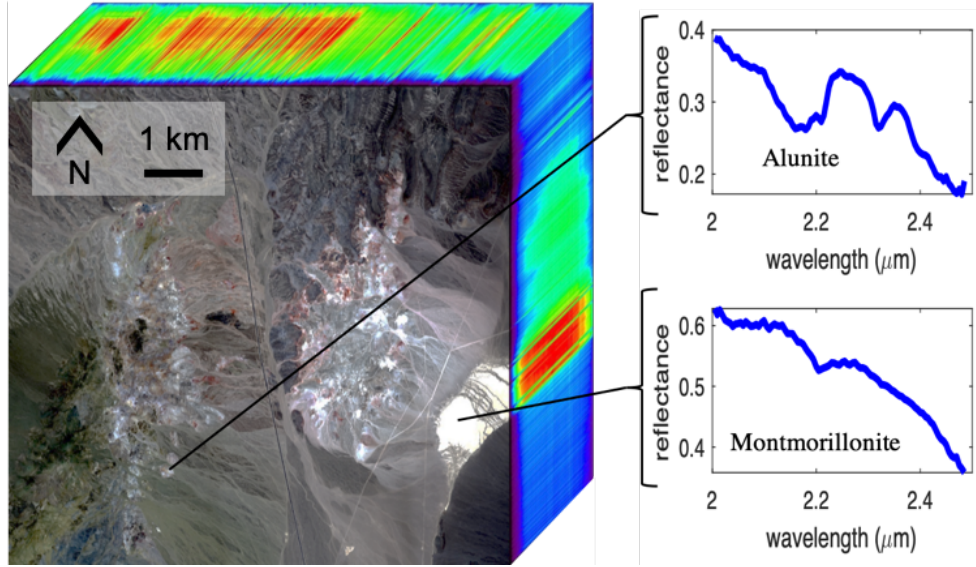


Figure 3.1: Example of a spectroscopic image. A 3-D representation is shown on the left, while two reflectance spectra corresponding to two pixels in the image appear on the right.

sensing. We are especially interested in imaging spectroscopy, which is the acquisition of images where each pixel stores information from many contiguous wavelengths of the electromagnetic spectrum, instead of just the three bands that comprise the standard RGB color model. Spectroscopic images are also known as hyperspectral images or hyperspectral cubes. They contain much more information than a conventional RGB image, yet a similar spatial structure (Figure 3.1).

3.2 Analysis Techniques

The analysis of spectroscopic data may be 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. This analysis becomes increasingly hard on new planets or environments. Therefore, adequate processing 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. Understanding these spectroscopic analysis techniques, as well as their benefits and drawbacks, is important for our research since we want to enable robotic explorers to process spectroscopic data in a principled and automated manner.

Spectral analysis can be performed by focusing on relevant features (Figure 3.2). In this context, spectroscopists tend to use the absorption bands of spectra as characteristic features for material identification [32, 104]. Spectral absorption features have basic elements such as position, depth, and width. These elements have been used to construct expert systems for simple shape matching [105]. More sophisticated methods perform a complete shape matching of the

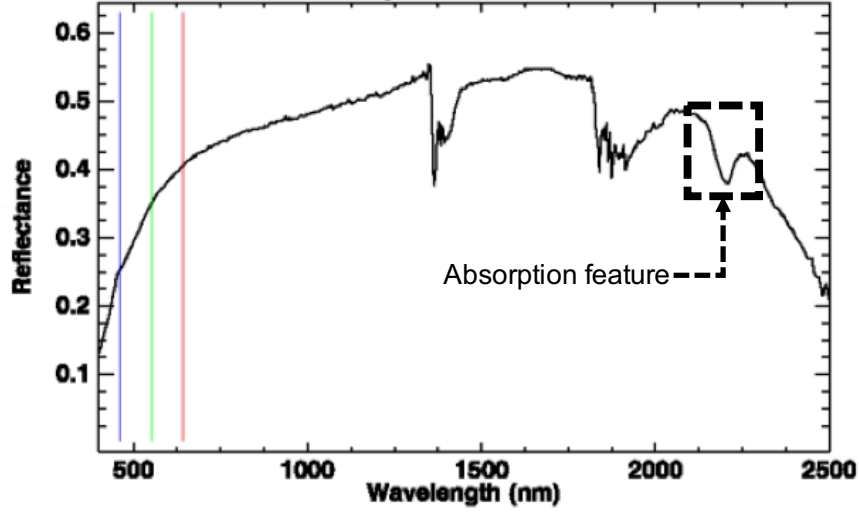


Figure 3.2: Example of a spectral absorption feature in the infrared wavelengths.

absorption features using a least squares criterion in conjunction with encoded expert knowledge, most notably the Tetracorder system [33]. Finally, some techniques attempt to model the features by fitting functions such as Gaussian, Lorentzian, or Voigt curves to spectral signals [15]. In general, these methods usually require more effort to implement and are highly sensitive to noise, but their results and way of operation are better in terms of scientific interpretability.

Another approach to spectral analysis is to look at the whole spectrum and compare it to spectra of known materials. A very popular method involves comparing pairs of spectra (\mathbf{x}, \mathbf{y}) using distance metrics. Probably the most common distance is the Spectral Angle Mapper (SAM), which is a direct analog of the cosine distance function or the normalized cross-correlation operation [104]:

$$SAM(\mathbf{x}, \mathbf{y}) = \cos^{-1} \left(\frac{\mathbf{x} \cdot \mathbf{y}}{\|\mathbf{x}\|_2 \|\mathbf{y}\|_2} \right). \quad (3.1)$$

Other examples of distance metrics within this context are the Spectral Information Divergence (SID), a symmetric variation of the Kullback-Leibler divergence that treats spectra as probability vectors [26]:

$$SID(\mathbf{x}, \mathbf{y}) = D_{KL}(\mathbf{x} \parallel \mathbf{y}) + D_{KL}(\mathbf{y} \parallel \mathbf{x}); \quad (3.2)$$

as well as the Hamming distance, which involves encoding spectra as binary signals [117]. Full spectral analysis methods are relatively simple to automate and are more robust to the presence of noise, but they tend to lose physical interpretability.

The aforementioned strategies may not be very useful when dealing with mixtures of different materials, which is a common problem as it is often difficult to find pure materials in the environment. These pure constituents are called *endmembers*. The computation of the fractional abundance of each endmember is known as *spectral unmixing*. In general, there are two different

models: linear and nonlinear. Linear models assume that a spectrum \mathbf{x} can be represented as a linear combination of its endmembers:

$$\mathbf{x} = \mathbf{E}\mathbf{r}, \quad (3.3)$$

where $\mathbf{E} = [\mathbf{e}^1, \mathbf{e}^2, \dots, \mathbf{e}^n]$ is a spectral library consisting of n endmembers, and $\mathbf{r} = [r_1, r_2, \dots, r_n]$ is the vector of mixing ratios, or fractional abundances, and the variable to solve in the linear equation. This linear equation can be solved by using the least squares method, which minimizes the following expression:

$$\min_{\mathbf{r}} \|\mathbf{E}\mathbf{r} - \mathbf{x}\|_2^2 \quad (3.4)$$

There are three main techniques for linear unmixing using the least squares method:


- Unconstrained least squares: solves the equation for \mathbf{r} using normal least squares without any constraints on \mathbf{r} .
- Non-negative least squares: solves the equation for \mathbf{r} with $r_i \geq 0, \forall r_i$
- Fully-constrained least squares: solves the equation for \mathbf{r} with $r_i \geq 0, \forall r_i$ and $\sum_{i=1}^n r_i = 1$.

However, linear models may be too simplistic in some situations. There exist more complex nonlinear models for spectral unmixing [32, 109, 111]. These assume that a spectrum \mathbf{x} can be represented as nonlinear function of the endmembers and their fractional abundances:


$$\mathbf{x} = f(\mathbf{E}, \mathbf{r}) \quad (3.5)$$

Nonlinear mixing models may consist of operations ranging from polynomial functions to sophisticated neural networks. They may also be based on radiative transfer and light scattering models that are more physically grounded [69, 70]. Nonlinear models may be capable of fitting the data with more accuracy, but have the drawback of potential overfitting, even when using a small number of model parameters.

In principle, many image processing, machine learning, and computer vision tools can be applied to spectroscopic images. Algorithms ranging from Gaussian Classifiers and K-means [171] to Deep Neural Networks [139] have been applied to these. An important thing to underscore is that spectroscopic images may be very large files, therefore dimensionality reduction can be used 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 [109, 111]. Another 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. They have also been used for various spectroscopic data analysis tasks [40, 45, 163]. Figure 3.3 shows an example of the Simple Linear Iterative Clustering (SLIC) superpixel algorithm [1] when applied to a spectroscopic image.



(a) RGB channels.



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

Figure 3.3: Example of the superpixel segmentation of a spectroscopic image using SLIC [1].

Table 3.1: Summarized comparison of the spectral analysis techniques.

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

In conclusion, there are multiple strategies for the analysis of spectroscopic data. A simple and compact taxonomy can be described as having full spectral vs. spectral feature analysis techniques, as well as having linear vs. non linear spectral unmixing methods (Table 3.1). Among them, there is an overall trade-off between computational simplicity (both in terms of implementation and data processing), physical interpretability, and accuracy. Additionally, given the high dimensionality of spectroscopic data, preprocessing and dimensionality reduction techniques can be quite useful. Our presented Bayesian models, especially the ones described in Chapter 5, address many of the aforementioned aspects of spectroscopic data analysis.

Chapter 4

Scenarios of Science-Driven Robotic Exploration

The thesis statement (Chapter 1) is supported with evidence from scientifically relevant scenarios that can provably benefit from autonomous robotic exploration. This is shown in extensive simulations, and in certain instances, in field experiments with the robotic platform Zoë. Overall, we focus on geologic investigations as they currently are a major component of planetary robotic exploration. We present three different geologic scenarios: terrestrial exploration in Nevada, studies of airless planetary bodies, and Martian surveys. In addition to our geologic exploration scenarios, we recognize that other science endeavors can benefit from autonomous robotic exploration. Hence, we address marine biology studies via simulations of coral reef investigations, which may be performed with autonomous underwater vehicles (AUV).

This chapter starts by describing the robotic platform Zoë. Then, it explains the four aforementioned geologic exploration scenarios. Afterward, it discusses the coral reef study simulation. Finally, it includes a summary relating common elements among these science scenarios.

4.1 Robotic Platform Zoë

Zoë is an autonomous astrobiology rover developed by Carnegie Mellon University (Figure 4.1). It is named after the Greek word for "life". It was originally developed to carry out a series of astrobiology investigations in the Atacama Desert, Chile [175]. Specifically, Zoë was designed to provide several improvements over its previous generation, the rover Hyperion [176].

The rover has a width of 1.63 m, a length of 2.0 m, and a mass of 198 kg. Zoë has 4 independently driven wheels. Each axle is attached to the chassis by joints that are free to rotate in two degrees-of-freedom. The rover's maximum velocity is 0.9 m/s or 3.2 km/h. Zoë is powered by an array of solar panels that provide, on average, 27% efficiency in converting solar energy. Additionally, two 1500Whr Li-Polymer batteries charge when there is excess power and are



Figure 4.1: Zoë is an autonomous astrobiology rover developed by Carnegie Mellon University. The robot carries various instruments that allow for autonomous science.

drawn down when energy is needed.

Zoë carries a customized and integrated portable spectrometer, an Analytical Spectral Devices (ASD) FieldSpec Pro. The instrument has 2151 channels within a spectral range of 350-2500 nm. It is internal to the rover but connected by a multi-strand fiber-optic line to a 1 foreoptic mounted on a pan-tilt mechanism. This allows the spectrometer to collect spectra of nearby rocks and soils, or patches of more distant terrain.

Zoë carries a computer with an Intel i7 processor with a base frequency of 2.8 GHz and 4 cores, as well as 4GB of DDR4 memory. The rover's software architecture has evolved over the years. The latest architecture is implemented using the Robot Operating System (ROS) [140]. We pay special attention to the science planner in the architecture (Figure 4.2). The science planner enables the rover to reason about scientific objectives and make decisions about data collection. It is the module that contains the Bayesian models and runs the algorithms presented in this research. The science planner updates the Bayesian models as the rover collects new data and generates new science plans accordingly.

4.2 Geologic Surveys

Planetary rovers usually focus on performing geologic studies. They are equipped with special tools to study a diverse collection of dust, soil, and rocks. Furthermore, the record of past water activity can be found in rocks, minerals, and geologic landforms, particularly in those that can only form in the presence of water. In turn, water is of great importance for astrobiology since it is strongly correlated to habitability, biosignature potential, and ultimately life.

Next we describe the three geologic scenarios that are explored in this work: terrestrial sur-

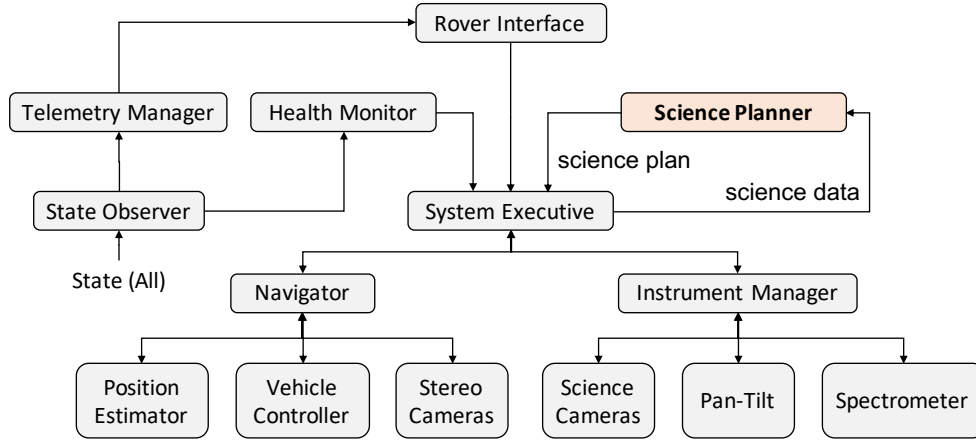


Figure 4.2: Software architecture of the Zoë rover. We emphasize the science planner module, which allows the rover to reason about science goals and make decisions about data collection. The science planner contains a model of the environment that updates as new science data is available. Afterward, new science plans are generated accordingly and sent for execution.

face studies at Cuprite Hills, Nevada, analog studies of airless planetary bodies, and Martian surface studies at Jezero Crater and Nili Fossae.

4.2.1 Terrestrial Surface Investigations

Cuprite Hills, Nevada is a well-studied region of high mineralogical diversity that is amenable to remote sensing [161], therefore making it an important test site for spectroscopic algorithms [168]. In this work we will refer to this as the *Cuprite scenario*. This scenario involves simulation studies and field experiments with the Zoë rover.

We used data from three different spectrometers: the Advanced Spaceborne Thermal Emission and Reflection Radiometer (ASTER) [49], the Airborne Airborne Visible Infrared Imaging Spectrometer Next Generation (AVIRIS-NG) [68], and the rovers’s ASD Fieldspec Pro (ASD). ASTER is a low-resolution orbital instrument, AVIRIS-NG is a high-resolution airborne spectrometer, and ASD is a high-resolution *in situ* device. The data consists of reflectance spectra, representing the fraction of incident light reflected in each wavelength. The AVIRIS-NG spectra was processed with the atmospheric correction approach described in [165]. ASTER and AVIRIS-NG data products were spatially aligned with manually-selected ground control points. They were co-registered using a first degree polynomial warping transformation, and then re-sampled to 15 m/pixel. Figure 4.3 shows an example of a couple of representative spectroscopic measurements in the scene as seen by ASTER and AVIRIS-NG instruments. Table 4.1 shows the wavelengths, channels, and resolutions that were used for each instrument.




Figure 4.3: Example of remote spectroscopic measurements of Cuprite, Nevada, as seen by ASTER and AVIRIS-NG.

Table 4.1: Measurement and instrument specifications

Instrument	Source	Used Wavelengths	Channels	Average Resolution
ASTER	Spacebourne	2.0 - 2.5 μm	5	80 nm
AVIRIS-NG	Airborne	2.0 - 2.5 μm	85	5 nm
ASD	Rover	2.0 - 2.5 μm	500	1 nm

4.2.2 Airless Surface Investigations

Toolbox for Research and Exploration (TREX) is a NASA project that aims to develop tools and research methods for exploration of airless surfaces in preparation for human and robotic *in situ* resource utilization (ISRU) missions. The TREX project is being led by the Planetary Science Institute and has several collaborators from other organizations including Carnegie Mellon University. TREX is taking current knowledge of spectral characteristics of candidate materials out into the field, to test algorithms for sample identification and selection, and to understand how

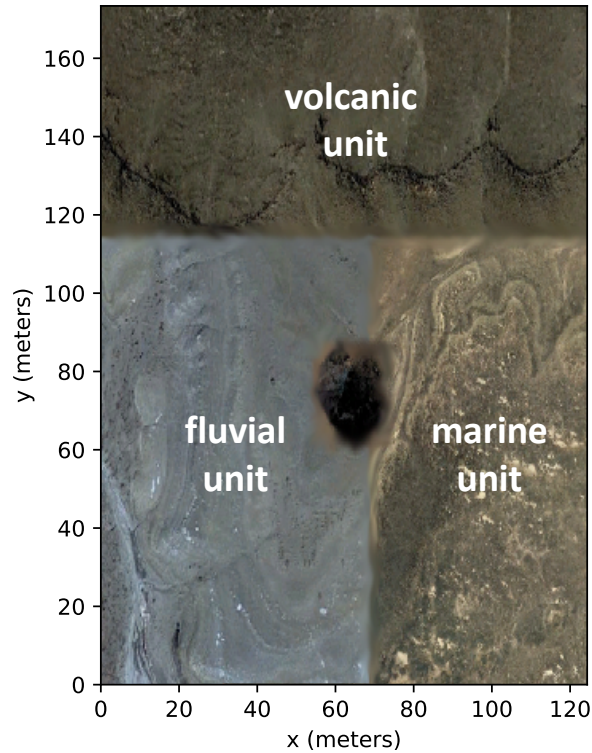


Figure 4.4: Synthetic geologic environment prepared by the TREX team for the operations readiness test in Gascola, Penn Hills, PA. The environment consists of a custom AVIRIS-NG spectroscopic image that merges three different geologic units: volcanic, fluvial, and marine.

these new tools can help to make future exploration more efficient. Experiments will be conducted at locations on Earth mimicking the fine-grained surface conditions typically found on the Moon and small bodies such as asteroids [131]. Herein we focus on an operations readiness test (ORT) performed in preparation for the upcoming field campaign in the Painted Desert, Arizona. In this work we will refer to this as the *TREX scenario*. This scenario involves simulation studies and field experiments with the Zoë rover.

The ORT was conducted between 26-30 July 2021 at Gascola, Penn Hills, Pennsylvania. To this end, the TREX science team prepared a synthetic geologic environment for the Gascola site. The geologic environment consisted of a synthetic AVIRIS-NG spectroscopic image in which three different geologic units were merged: volcanic, fluvial, and marine (Figure 4.4). The image has 640 rows, 620 columns, and 224 spectral channels; its spatial resolution is of approximately 0.2 m/pixel. This image was geographically projected onto the Gascola site. Additionally, we deployed a drone and collected aerial imagery to better assess the Gascola site, as well as to generate a traversability map that allowed the rover’s science planner to avoid artificial obstacles such as logs, power line poles, trees, and mounds (Figure 4.5). Specifically, we defined an elevation constraint of 200 meters and a slope threshold of 20 degrees.

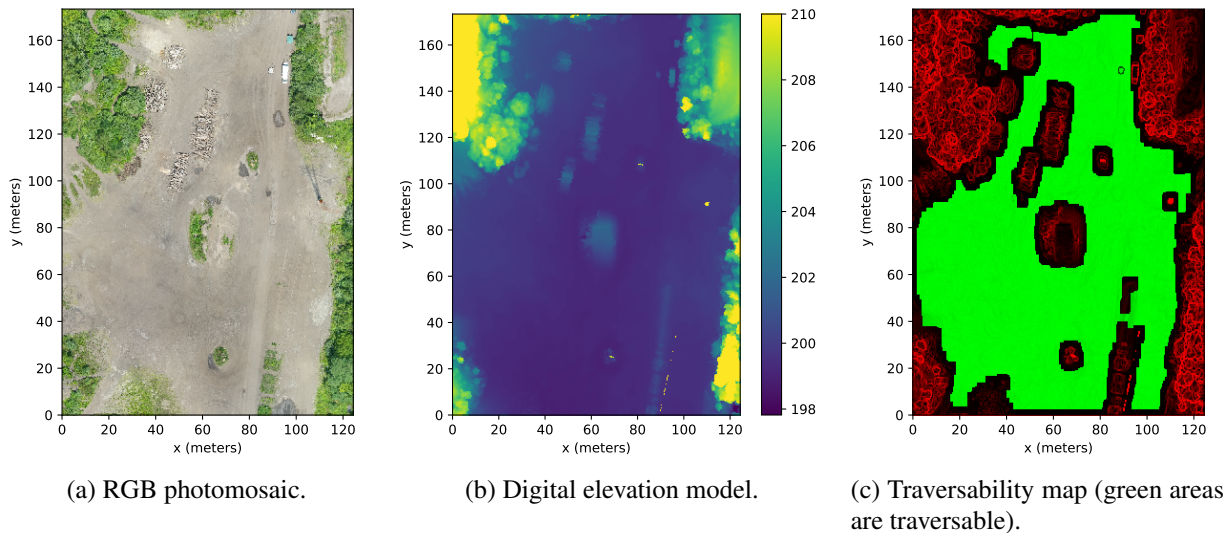


Figure 4.5: Drone data products for the test site at Gascola.

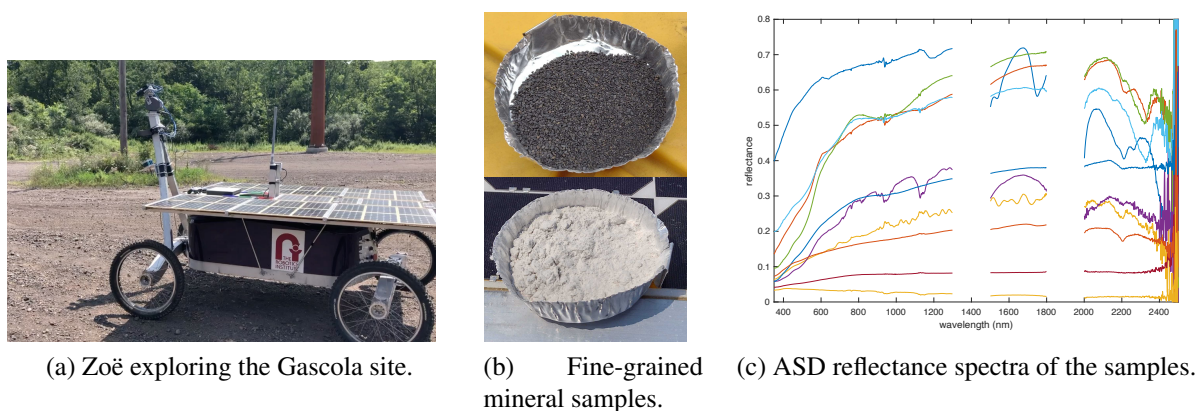


Figure 4.6: *In situ* data collection at Gascola.

The Zoë rover was deployed at this site in order to realistically simulate data collection from the synthetic geologic environment (Figure 4.6a). Around 20 mineral samples were prepared by the TREX science team: they contained different fine-grained materials such as feldspars, kaolinites, and hematites (Figure 4.6b). These samples were placed at specific geographic coordinates according to their associated locations in the synthetic geologic environment. Whenever the rover reached a particular waypoint, the corresponding sample was measured with manual assistance. Specifically, the rover collected reflectance spectra of predetermined samples using its onboard ASD spectrometer (Figure 4.6c).

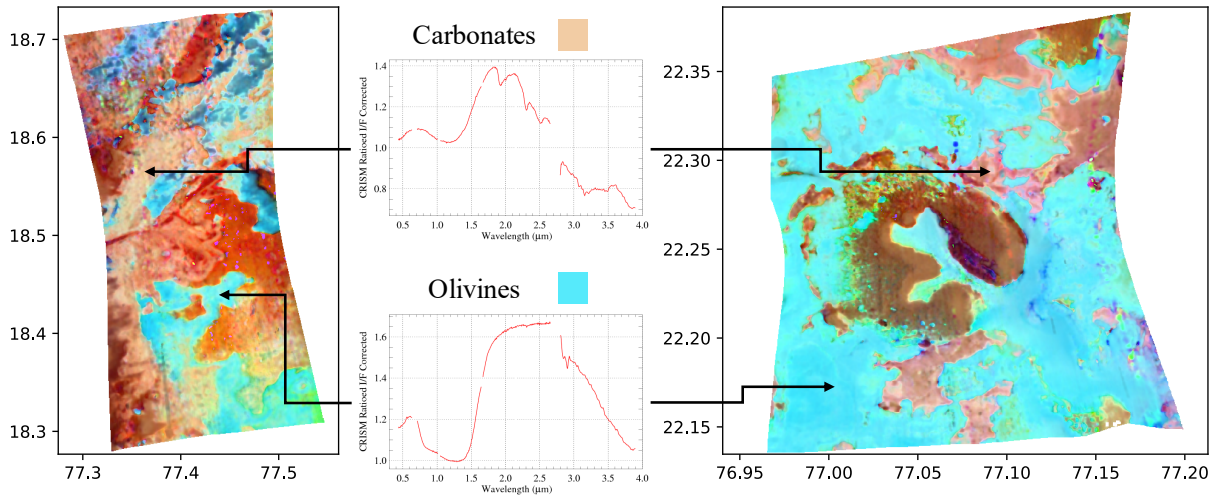


Figure 4.7: Spectroscopic maps of Martian regions at Jezero Crater (left) and Nili Fossae (right). Mineralogy is estimated by looking at spectral signatures using data from the orbital instrument CRISM (center). Cream tones represent carbonates, whereas cyan indicates olivine presence.

4.2.3 Martian Surface Investigations

The *Mars scenario* focuses on spectroscopic investigations and mineral mapping tasks at two locations of interest on Mars: Jezero Crater and Nili Fossae (Figure 4.7). This scenario involves simulation studies and it uses real or synthetic data from four different spectrometers.

- The Mars Express High-Resolution Stereo Camera (HRSC) instrument [91]. HRSC has global coverage of Mars, but its resolution is insufficient for composition analysis. It has an approximate resolution of 50 m/pixel and 4 color channels: blue 440 nm, green 530 nm, red 750 nm, and infrared 970 nm. HRSC image h5270_0000 is used for both sites.
- The Mars Reconnaissance Orbiter (MRO) Compact Reconnaissance Imaging Spectrometer for Mars (CRISM) instrument [128]. CRISM is an imaging spectrometer that provides high-resolution spectra (6.5 nm/channel) in a wide spectral range (436 - 3897 nm) at a spatial resolution of 18-36 m/pixel; unfortunately, it does not have a global coverage. The following CRISM images are used in this study (Figure 4.7): HRL000040FF (Jezero Crater), and FRT00003E12 (Nili Fossae).
- The Mars 2020 Perseverance rover MastCam-Z instrument [7]. MastCam-Z is a multispectral, stereoscopic imager. It also allows for mineral identification in the visible and near-infrared wavelengths (400 - 1100 nm). MastCam-Z consists of 2 cameras (left and right) with a total of 20 bandpass filters (channels). We simulate measurements by convolving CRISM data with the spectral response profile of MastCam-Z.
- The Mars 2020 Perseverance rover SuperCam instrument suite [177]. SuperCam provides a number of versatile sensing techniques that can be used at long distance as well as within robotic-arm workspace. These include laser-induced breakdown spectroscopy (LIBS), Raman




Figure 4.8: A mineral (Fe olivine) as measured by three Mars instruments: CRISM, MastCam-Z, and HRSC. Their channels and respective bandwidths are shown.

and luminescence spectroscopies, and visible and infrared reflectance spectroscopy. Additionally, a remote micro-imager (RMI) provides high-resolution color context imaging, and a microphone can be used as a complementary tool. Herein we focus on the infrared spectrometer which operates with 256 channels in the 1300-2600 nm spectral range. We simulate measurements by convolving CRISM data with the spectral response profile of SuperCam.

The HRSC and CRISM products were spatially aligned with manually-selected ground control points. They were co-registered using a first degree polynomial warping transformation, and then resampled to 36 m/pixel. CRISM data was preprocessed and ratioed using the standard procedure in [173]. Afterwards, the empirical line method [155] was used to find the correspondence between the HRSC and CRISM reflectance values. MastCam-Z and SuperCam spectra were estimated by convolving their spectral response profiles with CRISM data. For MastCam-Z, a subset of 12 channels were used since 8 of the filters either overlap (i.e. too redundant) or have very wide bandwidths. The relationship between the instruments' different wavelengths and bandwidths are shown in an example of this curated data set in the 400-110 nm range (Figure 4.8).

This study relies on the CRISM spectral library for mineral classification [173]. The CRISM spectral library contains many of the minerals that have been identified on Mars. Herein we focus on carbonates since they have been found on both sites and are of scientific interest due to their fossil preservation properties [85, 173]. We also consider olivines and silicas. Specifically, we use the following five minerals: Fe/Ca-carbonate, Mg-carbonate, Fe-olivine, Mg-olivine, and hydrated silica.

For our work regarding science and risk-aware planning for planetary rovers (Chapter 7), we focus on Jezero Crater and also employ data from MRO's High Resolution Imaging Science Experiment (HiRISE) [118]. HiRISE is an instrument that captures high resolution orbital imagery at 0.25 - 0.30 m/pixel. We also use DEMs that were generated from HiRISE stereo images and have a resolution of 1 m/pixel, allowing us to infer terrain properties such as slope. Furthermore,

we utilize a recent geologic map of Jezero Crater that differentiates between bedrock (low-risk) and surficial (high-risk) units [158].

4.3 Coral Reef Studies

This scenario focuses on spectroscopic investigations and coral reef mapping tasks. In this work we will refer to this as the *Coral scenario*, which involves simulation studies.

In addition to their significance in the marine biome, coral reefs are important to the cultural and economic lives of hundreds of millions of people around the world [34, 125]. It is incontrovertible that many coral reefs are in various stages of decline and may be unable to withstand the consequences of global climate change [56, 80, 81, 135, 157]. Yet a small fraction of the world's reef area has been studied quantitatively (i.e., 0.01-0.1%) as most reef assessments largely rely on local in-water surveys [78]. Therefore, our current understanding may not be representative of the reef under study, nor the regional and global reef ecosystem given existing data constraints [77].

Remote sensing from airborne and spaceborne platforms have proven to be a useful tool for aspects of reef science [72]. Early applications of remote sensing to coral reef environments focused on mapping reef geomorphology and ecological zonation [64, 106, 127]. In the past few decades, much of remote sensing has focused on mapping habitats using qualitative descriptors comprising various combinations of substrate (e.g., sand, limestone, rubble), benthic functional type (e.g., coral, algae, seagrass), reef type (e.g., fringing, patch, barrier), and/or location within the reef system (e.g., slope, flat) [72]. The recent NASA Earth Venture Suborbital-2 (EVS-2) CORAL Reef Airborne Laboratory (CORAL) mission focused on reef benthic functional type discrimination [77]. In this work we use high-resolution airborne data from the CORAL mission.

CORAL mapped portions of the Great Barrier Reef, main Hawaiian Islands, Marianas Islands, and Palau in 2016-2017. We focus on two flight lines from Heron Island, Australia on 17 September 2016 and Kaneohe Bay in Oahu, Hawaii on 6 March 2017 (Figure 4.10). Data is from the NASA/JPL Portable Remote Imaging SpectroMeter (PRISM) flown on the Tempus Applied Solutions Gulfstream-IV at 8.5 km altitude. PRISM provided water-leaving reflectance from 350-1050 nm (246 channels) at approximately 3 nm spectral resolution and approximately 8 m spatial resolution from which benthic reflectance and benthic functional type were derived. The benthic reflectance calculation used the shallow-water reflectance described by Maritorena et al. [115]. It was modeled using a linear nonnegative combination of a set of one or more basis endmembers from a library of bottom reflectances [167]. Here we evaluate benthic reflectance products since they provide invariance to water column properties; these products have a 420-680 nm spectral range and consist of 92 channels (Figure 4.9).

Benthic functional type corresponds to probabilities associated with coral, algae, and sand for each seafloor pixel, and was determined by logistic regression using mean reflectance spec-




Figure 4.9: Benthic reflectance of coral (left), algae (center), and sand (right) as estimated using the approach by Thompson et al. [167] applied to PRISM data.

tra from an existing spectral library [76, 78]. The focus of CORAL - and thus this work - was on coral, algae, and sand spectra (4.9). Examples of PRISM-derived benthic composition from the CORAL mission are shown in Figure 4.10, with each pixel representing the primary benthic functional type (i.e., greatest percentage of one type per pixel, wherein all types sum to one per pixel). Herein we evaluate PRISM-derived benthic reflectance since it provides invariance to water column properties. Utilization of *in situ* measurements of benthic reflectance are preferred; however, there were no coincident collections of such measurements with overflights under cloud-free conditions. Maps of PRISM-derived benthic composition (i.e., coral, algae, and sand) are used to validate our mapping/sampling results, serving as “ground truth”. Note that PRISM-derived products were validated by an extensive field collection as part of the CORAL mission and is not the focus herein (e.g., Thompson et al. [167]). For example, PRISM-derived benthic composition maps were validated by 10 m x 10 m photomosaics collected in the field at random. In each 10 m x 10 m plot approximately 500-1000 digital photographs were taken, mosaicked using structure-from-motion techniques (Agisoft Metashape) to provide a single orthomosaic, and then analyzed using coral reef science standard point-count methods to provide proportional benthic composition for benthic functional types.

We also use Level 2 surface reflectance products from the orbital instrument Landsat-8 (Figure 4.10). Landsat-8 provides global-scale coverage, but at coarser spatial and spectral resolution, serving as our source of remote data. We use the first four channels, which provide limited information in the visible wavelengths as compared to PRISM data. For Kaneohe Bay, we use image LC08_L2SP_064045_20170306_20200905_02_T1, which was collected on the same day as the corresponding CORAL flight line (6 March 2017). For Heron Island, we use image LC08_L2SP_091076_20161026_20200905_02_T1; due to cloud coverage issues, we had to use an image that was collected on a different day (26 October 2016).




Figure 4.10: Remote sensing data used for the coral studies. Landsat 8 data provides global-scale coverage but at coarse spatial and spectral resolutions that are often insufficient for benthic cover analysis. PRISM is an airborne imaging spectrometer with finer spatial and spectral resolution that was used by the CORAL mission to discriminate benthic functional types. The CORAL mission focused on three benthic functional types: coral (red), algae (green), and sand (blue). Abundance maps were estimated by Thompson et al. [167] and validated by the CORAL mission with photomosaics collected in the field.

4.4 Summary

This chapter has presented four scientifically relevant scenarios that can provably benefit from autonomous robotic exploration: a) a terrestrial geology scenario; b) an airless surface geology scenario; c) a Martian geology scenario; and d) a marine biology scenario. These scenarios for science-driven robotic exploration were used to develop and test the proposed Bayesian models in this research. It is worth noting that these scenarios share many elements in common. All of them rely on spectroscopic data for analysis and interpretation. However, each scenario uses a different combination of orbital, aerial, and surface data. A summary of this information is shown in Table 4.2.

Table 4.2: Science scenarios and relevant information.

Scenario Name	Scientific Discipline	Spatial Domain	Sites	Orbital Data	Aerial Data	Ground Data	Field Experiments
Cuprite	Geology	Terrestrial	Cuprite Hills, NV	ASTER	AVIRIS-NG	ASD	Yes
TREX	Geology	Airless Bodies	Gascola, PA	-	AVIRIS-NG Drone Imagery	ASD	Yes
Mars	Geology	Martian	Nili Fossae Jezero Crater	HRSC CRISM HiRISE	-	-	No
Coral	Biology	Marine	Heron Island Kaneohe Bay	Landsat-8	PRISM	-	No

Chapter 5

Deep Generative Models for Spectroscopic Data Analysis

5.1 Introduction

Earth and planetary sciences often rely upon the analysis of spectroscopic data. Spectrometers measure different wavelengths of the electromagnetic spectrum and thus capture more information than can be seen with the eye. A more detailed explanation regarding spectroscopic data can be found in Chapter 3. A key insight is that spectra contain spectral features that are used for composition analysis [32]. The detection of spectral features often requires high resolution spectroscopic measurements [32] as low-resolution spectra may sometimes lose valuable information (Figure 5.1). Unfortunately, high resolution spectra are often scarce, whereas low resolution spectra are usually easier to obtain [88, 184]. Another important problem is that spectra often contain mixtures of several components; therefore, estimating the correct combination of endmembers and their respective fractional abundances can be a challenging task.

Our research addresses three main aspects of spectroscopic data analysis: spectral features, spectral resolution, and mixtures of endmembers. To this end, we rely on the power of Bayesian models combined with deep learning, also known as *deep generative models*. We are interested not only in achieving high performance, but also in helping scientists make informed and analytical decisions. We start by looking at the problem of spectral feature extraction with the help of Variational Autoencoders, popular neural networks for density estimation and dimensionality reduction. Afterward, we focus on the problem of learning and exploiting the relationships between different spectrometers. To this end, we introduce the Deep Conditional Gaussian Model, an architecture that performs Gaussian parameter estimation for each different input, providing a practical tool for probabilistic regression [21]. Particularly, we show how it can be used to enhance low-resolution data; a task also known as super-resolution. Finally, we focus on the problem of spectral composition analysis by presenting the Deep Conditional Dirichlet Model, a neural network that learns a Dirichlet parametrization for each spectrum. Dirichlet density esti-




Figure 5.1: Shortwave infrared spectra from two different phyllosilicate minerals as measured by both low (red) and high (blue) resolution instruments. Low resolution measurements fail to capture important spectral features required for mineral identification.

mation allows for spectral classification, and more generally, for spectral endmember unmixing.

In this research we develop and evaluate the aforementioned deep generative models. To this end, we use data from three of the science scenarios described in Chapter 3: Cuprite, Mars, and Coral. Furthermore, we show how these generative models provide both mathematically-grounded and human-interpretable results, making them compelling spectroscopic analysis tools. The structure of this chapter is as follows. We start by explaining our work with Variational Autoencoders for spectral feature extraction. We then describe the Deep Conditional Gaussian Model regression and its applications to probabilistic regression and super-resolution. Finally, we discuss the Deep Conditional Dirichlet Model and its applications to spectral classification and endmember unmixing.

5.2 Related Work

In the past decade, deep learning [62] has achieved impressive results in a myriad of tasks. Some examples include: accurate image recognition for hundreds of diverse classes such as cars and cats [103], learning how to transfer artistic styles in paintings [57], and Atari game playing that surpasses human experts [124]. Despite their remarkable results, deep neural networks typically work as black boxes. In contrast, scientists require formal, transparent and interpretable methods for systematic analysis. In other words, scientists need to know how much to trust their tools. Recent efforts have tried to alleviate this problem by trying to quantify uncertainty in deep learning models. Gal and Ghahramani [53] use Monte Carlo dropout as a strategy to perform Bayesian approximations by randomly disabling network nodes during training and prediction.

The resulting empirical distribution is then used to obtain a mean value and a confidence measure in terms of the distributional variance. Kendall and Gal extend this idea and use Monte Carlo dropout to approximate epistemic and aleatoric uncertainties in neural networks [95]. Deep ensembles are another approach for generating distributions through empirical approximations [108]. They consist of multiple networks trained with different random initializations, working in a similar fashion to ensemble learning methods such as random forests. One major drawback of these random-based uncertainty quantification methods is that their results may be inconsistent or unstable if too few iterations are used. Another inconvenience is that deep neural networks often require substantial computational resources for training; consequently, randomized trials and ensembles are much more expensive to train. Furthermore, science exploration in remote environments usually operates with limited amounts of data, hence such approaches would most likely lead to overfitting.

Restricted Boltzmann machines, deep belief networks, and deep Bayesian networks are probabilistic graphical models structured in layers [75]. They have played an important role in the foundational work leading to deep learning as we know it today [151]. They are easy to formulate, but difficult for performing exact inference. Hence, inference approximations are usually required. They have fallen out of favor as networks have become deeper and more complex.

Deep generative models are an alternative for more transparent and interpretable deep learning architectures. They combine ideas from density estimation and deep learning. They usually have Bayesian representations in just a subset of the layers comprising the architecture, making them more scalable. Deep generative models have recently achieved impressive results, especially in tasks focused on unsupervised learning and data generation [86]. In general, there are two learning approaches for deep generative models depending on whether they optimize a likelihood function or not.

There are deep generative models for which the likelihood is not explicitly defined [181]. The main ones are generative adversarial networks (GANs) and Variational Autoencoders (VAEs). GANs consist of two neural networks contesting with each other in order to learn how to produce data, usually from a simple Gaussian random generator [61]. Despite their encouraging results in image generation, the density estimation process is highly opaque: it does not allow calculation of the likelihood of observed data, nor visualization of the learned statistical dependencies. Additionally, GANs are notoriously famous for being difficult to train. VAEs [97], also known as Deep Latent Gaussian Models (DLGMs) [145], are easier to train and have a more transparent probabilistic representation. This representation is learned by first encoding a data set into a latent space and then decoding it back into its original state. VAEs usually assume that latent variables are normally-distributed. The latent representation typically has a much lower dimensionality than the input and output spaces, hence VAEs are popular for dimensionality reduction and feature extraction. It is also possible to use VAEs to build conditional models connecting different types of inputs and outputs: they have been used for spectral unmixing [14, 136], and for the reconstruction of high resolution spectroscopic measurements from syn-

thetic low resolution data [18]. Although their probabilistic model is somewhat more tractable, it is usually impossible to compute the likelihood function analytically. That is why VAEs rely on variational approximations for training, hence their name. Additionally, the interpretation of the learned statistical dependencies is often difficult because they lie in a hidden representation.

There are deep generative models that are explicitly driven by likelihood functions. An example are normalizing flows [144], which also work with latent representations. They use change of variable transformations in order to generate complex models from simple probability distributions. In practice, they do not work well on high dimensional spaces, and the likelihood evaluation of data is dependent on how invertible and tractable these series of transformations are. Autoregressive networks are a family of algorithms that have gained notable popularity for image generation [172] and super resolution [38] tasks. This is because they are scalable, relatively simple to train, and allow computation of log-likelihoods. They work by defining an arbitrary sequence of pixels and learning a series of conditional distributions in the output layer. Consequently, they are unable to show correlations between any pair of non-sequential pixels. We are especially interested in architectures that learn less constrained probability representations in the output layer. Gaussian output layers have been explored in the literature, where neural networks predict both a mean and a variance for each target [174, 178]. Another relevant example involves Dirichlet prior networks for classification tasks [114].

In summary, deep generative models have achieved remarkable results in various tasks, but have made little progress in producing probabilistic models that directly address uncertainty quantification and interpretability.

5.3 Variational Autoencoder

5.3.1 Model and Architecture

Many channels and wavelengths in spectra are strongly correlated, especially adjacent ones. This allows for dimensionality reduction techniques that learn a set of representative features that efficiently compress and capture most of the information in spectra. We formulate the problem as follows. High-resolution spectra $\mathbf{y} \in Y \subset \mathbb{R}^n$ need to be converted into a set of low-dimensional features $\mathbf{z} \in Z \subset \mathbb{R}^d$, such that $d \ll n$. In this research we rely on the Variational Autoencoder (VAE) to achieve this task.

The VAE is a deep learning architecture that performs non-linear dimensionality reduction [97, 145]. The VAE is composed of two networks: an encoder that extracts features from high-resolution data, and a decoder that uses these learned features to reconstruct data back as $\hat{\mathbf{y}} \approx \mathbf{y}$ (Figure 5.2). The encoder of the VAE produces a latent feature representation parametrized by a conditional Gaussian distribution, that is:

$$\mathbf{z}|\mathbf{y} \sim \mathcal{N}_n(\boldsymbol{\mu}_z(\mathbf{y}), \boldsymbol{\Sigma}_z(\mathbf{y})), \quad (5.1)$$

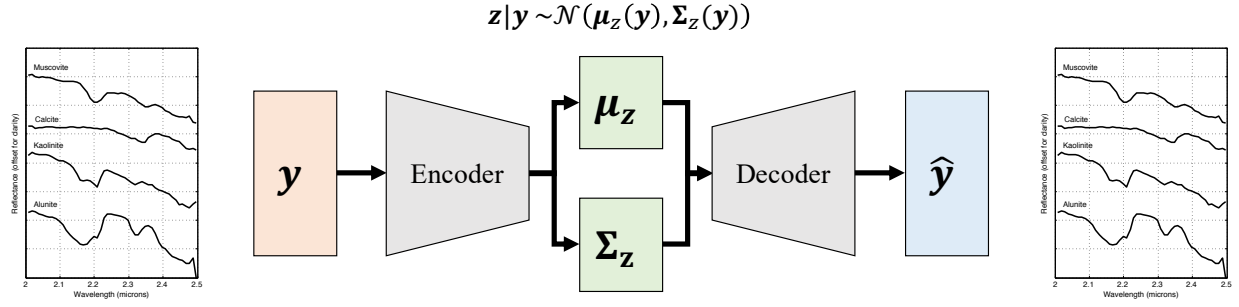


Figure 5.2: The architecture of the Variational Autoencoder (VAE). The input is a high-resolution spectrum $\mathbf{y} \in Y$ that is encoded into a low-dimensional Gaussian distribution $\mathbf{z}|\mathbf{y} \sim \mathcal{N}_n(\boldsymbol{\mu}_z(\mathbf{y}), \boldsymbol{\Sigma}_z(\mathbf{y}))$. Afterward, points are sampled from this distribution, passed through the decoder, and reconstructed back as $\hat{\mathbf{y}} \in Y$.

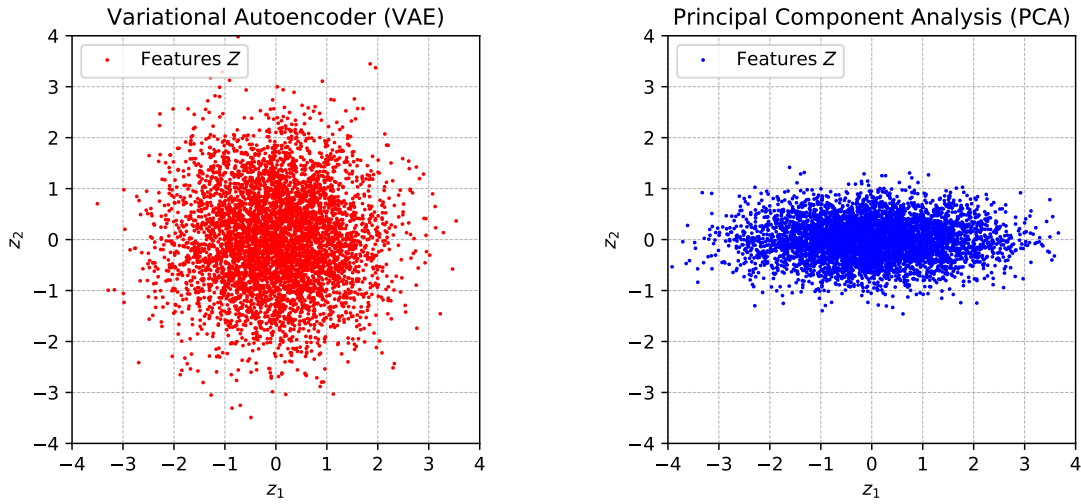


Figure 5.3: VAE vs. PCA feature representation. VAE decorrelates and normalizes features. PCA also decorrelates features, but it does not perform normalization.

where $\boldsymbol{\mu}_z \in \mathbb{R}^d$ is the mean and $\boldsymbol{\Sigma}_z \in \mathbb{R}^{d \times d}$ is the covariance. Furthermore, this representation uses a standard multivariate normal distribution as a prior, i.e., $\mathbf{z} \sim \mathcal{N}_d(\mathbf{0}, \mathbf{I}_d)$. This operation effectively decorrelates and normalizes features. Moreover, VAEs tend to spread information across the latent dimensions in a uniform manner while filtering noise. This yields an advantage against other feature extraction methods such as Principal Component Analysis (PCA) (Figure 5.3).

VAEs are often called unsupervised learning methods because they do not require labeled data for training. Strictly speaking, they are self-supervised learning methods as they use the same data as regression targets. Regardless, they offer a significant advantage over approaches that rely on manually-engineered features, which tend to be domain-specific, arbitrary, and laborious. Let us denote this data set as $\mathcal{D} = \{\mathbf{x}^i\}_{i=1}^N$. VAEs are trained using variational Bayesian methods, which are approximations for complex statistical models. In this particular case, it is

Table 5.1: Architecture of the VAE employed for spectral feature extraction and reconstruction.

Layer	Type	Activation Function	Number of Units
input	-	-	n
encoder 1	fully-connected	rectified linear unit	$(2/3) \times (n + d)$
encoder 2	fully-connected	rectified linear unit	$(1/3) \times (n + d)$
mean	fully-connected	linear	d
log-variance	fully-connected	linear	d
decoder 1	fully-connected	rectified linear unit	$(1/3) \times (n + d)$
decoder 2	fully-connected	rectified linear unit	$(2/3) \times (n + d)$
output	fully-connected	sigmoid	n

usually not possible to perform exact inference over the latent variables Z . VAEs approximate the posterior distribution by sampling features from the conditional distribution $p(\mathbf{z}|\mathbf{y})$. These sampled features are then passed through the decoder and reconstructed as $\hat{\mathbf{y}}$. Our specific model uses the conventional loss function for training VAEs, given by:

$$L = \sum_{\mathbf{y} \in \mathcal{D}} [\|\mathbf{y} - \hat{\mathbf{y}}\|_2^2 + \lambda D_{KL}(\mathcal{N}_n(\boldsymbol{\mu}_z(\mathbf{y}), \boldsymbol{\Sigma}_z(\mathbf{y})) \parallel \mathcal{N}_d(\mathbf{0}, \mathbf{I}_d))] . \quad (5.2)$$

This function adds two different loss terms: a) the reconstruction error between the predicted target and its true value; and b) a Kullback-Leibler divergence penalization of the latent space with respect to a standard unit Gaussian. This results in a trade-off between reconstructing data accurately and learning meaningful features; such trade-off is tuned with the parameter λ .

The specific architecture we employ to process spectra is quite simple. It consists of a series of sequential layers as specified in Table 5.1. The number of units of each layer depends on the input size n and the latent dimension d . Note that the variance layer operates in log-space to improve numerical stability.

5.3.2 Experiments and Results

For each science scenario, we employed a data set that consists of 3000 spectra sampled from high-resolution imagery. Specifically, CRISM spectra for the Mars scenario, AVIRIS-NG spectra for the Cuprite scenario, and PRISM spectra for the Coral scenario. More details can be found in Chapter 3. First, we performed a preprocessing step on the data to preserve spectral features and reduce the impact of misleading factors such as albedo and noise. Concretely, each spectrum was scaled in the range $[0,1]$ with min-max normalization. In each data set, 1000 spectra were used for the training set, 1000 for the validation set, and 1000 for the test set. The 3000 spectra were sampled using the k-means++ approach [4] to ensure they were diverse and representative. To avoid overfitting, we employed early stopping during training, which is a popular approach in

deep learning [62]. It consists in selecting the neural network weights that minimize the loss of the validation set while training (using the training set). Other approaches, such as k-fold cross-validation, are excellent to avoid overfitting, but are typically too expensive to use with neural networks. A 2.9 GHz Intel Quad-Core i7 laptop without a graphics processing unit (GPU) was used for training and testing.

In these experiments we compared two feature extraction algorithms: PCA and VAE. Each of these two models learned how to convert high-resolution spectra into a set of features, or dimensions in the latent space. Afterwards, the models were used to reconstruct spectra from features. In this experiment we observe how error changes as a function of number of features ranging from 1 to 12 for the Mars and Cuprite data sets; and ranging from 1 to 7 for the Coral data set (Figure 5.4). Reconstruction errors are similar for the training, validation, and test sets, suggesting a good generalization. VAE, a non-linear model, can compress data better than linear PCA in all cases since the reconstruction error is smaller and has a smaller variance. As one would expect, average error (for both PCA and VAE) tends to decrease as more features are used since less information is lost in the compression process. Consequently, the difference between PCA and VAE tends to reduce with higher dimensions. In those cases, PCA may be more appealing since it is much faster to train and apply, at least from an implementation standpoint. However, a disadvantage of PCA is that the first dimensions (features) learn most of the information while the last dimensions essentially learn just noise. In contrast, VAE tends to spread information evenly across the latent dimensions while filtering noise. Regarding the different scenarios, Mars/CRISM is the most noisy data set, followed by Cuprite/AVIRIS-NG. Cuprite is the scenario with the most spectral diversity, hence reconstruction errors tend to have larger variances.


Finally, examples of spectral reconstructions using only 3 features are shown in Figure 5.5. Overall, they seem quite accurate despite only using 3 dimensions in the latent space. Previous results are confirmed (Figure 5.4): Mars/CRISM data has more noise, whereas Coral/PRISM data has virtually no noise. It is worth noting that the VAE filters some of the noise in the data, especially for Mars and Cuprite spectra.

5.4 Deep Conditional Gaussian Model


5.4.1 Model and Architecture

The DCGM learns a conditional probability model that uses low resolution measurements in order to infer high resolution spectra [21]. Let $\mathbf{x} \in X \subset \mathbb{R}^m$ and $\mathbf{y} \in Y \subset \mathbb{R}^n$ denote the low and high resolution spectra, respectively. We define the conditional probabilistic relationship between X and Y as $p_{\theta}(\mathbf{y}|\mathbf{x})$. It is modeled using deep learning, where the network’s weights are defined as θ . This model assumes that the conditional probability can be approximated with a multivariate Gaussian distribution, hence the name *Deep Conditional Gaussian Model*. Formally,

Mars Scenario: CRISM spectra



Cuprite Scenario: AVIRIS-NG spectra



Coral Scenario: PRISM spectra






Figure 5.4: Reconstruction error as a function of feature dimension for the Mars, Cuprite, and Coral scenarios. Results show the average error plus-minus one standard deviation for each of the three sets: training (left), validation (center), and test (right). The Variational Autoencoder (VAE) performs better than principal component analysis (PCA) since its reconstruction error has a smaller mean and variance. All differences are statistically significant using a p-value of 0.05. Nonetheless, this difference in performance decreases as more dimensions are used in the feature space. Reconstruction errors are similar for the three sets, suggesting a good generalization.

Mars Scenario: CRISM spectra



Cuprite Scenario: AVIRIS-NG spectra



Coral Scenario: PRISM spectra




Figure 5.5: Reconstructed spectra using a VAE with 3 features or dimensions. Examples for the Mars, Cuprite, and Coral scenarios are shown. In all cases, the reconstructions are able to capture spectral shapes and features. Furthermore, VAE reconstructions reduce noise in the data, which is especially visible for the Mars scenario.

this means that:

$$\mathbf{y}|\mathbf{x} \sim \mathcal{N}_n(\boldsymbol{\mu}_{\mathbf{y}}(\mathbf{x}), \boldsymbol{\Sigma}_{\mathbf{y}}(\mathbf{x})) \quad (5.3)$$

where $\boldsymbol{\mu}_{\mathbf{y}} \in \mathbb{R}^n$ is the mean vector, and $\boldsymbol{\Sigma}_{\mathbf{y}} \in \mathbb{R}^{n \times n}$ is the covariance matrix. Note that both $\boldsymbol{\mu}_{\mathbf{y}}$ and $\boldsymbol{\Sigma}_{\mathbf{y}} \in \mathbb{R}^{n \times n}$ change as a function of \mathbf{x} . The probability density function is given by:

$$p_{\theta}(\mathbf{y}|\mathbf{x}) = \frac{\exp\left(-\frac{1}{2}(\mathbf{y} - \boldsymbol{\mu}_{\mathbf{y}}(\mathbf{x}))^T \boldsymbol{\Sigma}_{\mathbf{y}}^{-1}(\mathbf{x}) (\mathbf{y} - \boldsymbol{\mu}_{\mathbf{y}}(\mathbf{x}))\right)}{\sqrt{(2\pi)^n \det(\boldsymbol{\Sigma}_{\mathbf{y}}(\mathbf{x}))}}. \quad (5.4)$$

We rely on the well-known maximum likelihood estimation (MLE) method in order to estimate the parameters of the statistical model, that is, to tune the network's weights θ . The training data set consists of pairs of inputs and ground-truth outputs: $\mathcal{D} = \{(\mathbf{x}^i, \mathbf{y}^{i*})\}_{i=1}^N$. The network will try to maximize the following cumulative log-likelihood function:

$$\mathcal{L}(\theta|\mathcal{D}) = \sum_{(\mathbf{x}, \mathbf{y}^* \in \mathcal{D})} \log p_{\theta}(\mathbf{y}^*|\mathbf{x}). \quad (5.5)$$

Most deep learning frameworks work by minimizing a loss function L . We can easily adapt MLE by defining a negative log-likelihood (NLL) loss function as follows:

$$L = -\mathcal{L}(\theta|\mathcal{D}) \quad (5.6)$$

The log-likelihood function for a multivariate Gaussian distribution is given by:

$$\log p_{\theta}(\mathbf{y}|\mathbf{x}) = \log \left(\frac{\exp \left(-\frac{1}{2} (\mathbf{y} - \boldsymbol{\mu}_{\mathbf{y}}(\mathbf{x}))^T \boldsymbol{\Sigma}_{\mathbf{y}}^{-1}(\mathbf{x}) (\mathbf{y} - \boldsymbol{\mu}_{\mathbf{y}}(\mathbf{x})) \right)}{\sqrt{(2\pi)^n \det(\boldsymbol{\Sigma}_{\mathbf{y}}(\mathbf{x}))}} \right) \quad (5.7)$$

$$= -\frac{1}{2} (\mathbf{y} - \boldsymbol{\mu}_{\mathbf{y}}(\mathbf{x}))^T \boldsymbol{\Sigma}_{\mathbf{y}}^{-1}(\mathbf{x}) (\mathbf{y} - \boldsymbol{\mu}_{\mathbf{y}}(\mathbf{x})) - \frac{1}{2} \log \{ (2\pi)^n \det(\boldsymbol{\Sigma}_{\mathbf{y}}(\mathbf{x})) \}. \quad (5.8)$$

After analyzing Equation 5.8, one can observe a few interesting things. The first term is essentially the Mahalanobis distance

$$MD(\mathbf{y}) = \sqrt{(\mathbf{y} - \boldsymbol{\mu}_{\mathbf{y}})^T \boldsymbol{\Sigma}_{\mathbf{y}}^{-1} (\mathbf{y} - \boldsymbol{\mu}_{\mathbf{y}})}. \quad (5.9)$$

It is similar to the Euclidean distance, but scaled by the values in $\boldsymbol{\Sigma}_{\mathbf{y}}$ in order to measure distances in terms of standard deviations from the mean. We desire to minimize this distance by reducing the difference between the prediction $\boldsymbol{\mu}_{\mathbf{y}}$ and the real value \mathbf{y}^* , which is equivalent to a regression task. However, large values in $\boldsymbol{\Sigma}_{\mathbf{y}}$ will achieve the same result. Hence, the second term in Equation 5.8 serves as a regularization for $\boldsymbol{\Sigma}_{\mathbf{y}}$. As a matter of fact, this term is almost identical to the differential entropy of the Gaussian distribution, an information-theoretic metric used to quantify the uncertainty of the prediction [35]:

$$H(\mathbf{y}) = \frac{1}{2} \log \{ (2\pi)^n \det(\boldsymbol{\Sigma}_{\mathbf{y}}) \} + \frac{n}{2}. \quad (5.10)$$

One can observe that the only difference is the presence of the constant offset $\frac{n}{2}$. Note that the entropy is not defined by $\boldsymbol{\mu}_{\mathbf{y}}$, but rather by the volume of the covariance, which is given by $\det(\boldsymbol{\Sigma}_{\mathbf{y}})$. This means that entropy increases as a function of the volume of the covariance.

We conclude that the model minimizes a Mahalanobis distance loss by approximating $\boldsymbol{\mu}_{\mathbf{y}}$ to \mathbf{y}^* as closely as possible, i.e. regression. Simultaneously, the model minimizes an entropy

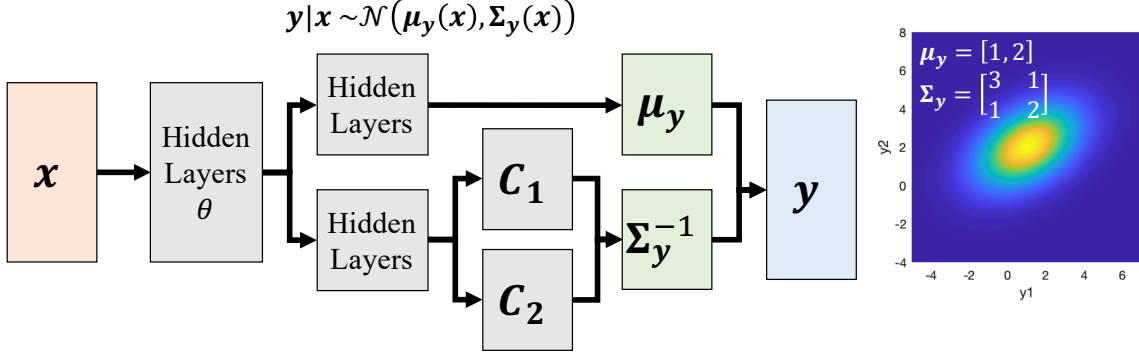


Figure 5.6: The architecture of the Deep Conditional Gaussian Model, based on an efficient Cholesky decomposition, is shown on the left. The input of the network is an available measurement $\mathbf{x} \in X \subset \mathbb{R}^m$. The output is the probabilistic prediction of the associated unavailable measurement $\mathbf{y} \in Y \subset \mathbb{R}^n$. This network learns the conditional probability function $p_\theta(\mathbf{y}|\mathbf{x}) = \mathcal{N}(\boldsymbol{\mu}_y(\mathbf{x}), \boldsymbol{\Sigma}_y(\mathbf{x}))$, where $\boldsymbol{\mu}_y \in \mathbb{R}^n$ and $\boldsymbol{\Sigma}_y \in \mathbb{R}^{n \times n}$. For efficiency, it operates with the Cholesky decomposition of $\boldsymbol{\Sigma}_y^{-1}$, represented as \mathbf{C}_1 and \mathbf{C}_2 . An example with $\mathbf{y} \in \mathbb{R}^2$ is shown on the right.

function reflecting the model’s confidence in its predictions, where accurate predictions should result in small covariances $\boldsymbol{\Sigma}_y$ and vice versa.

The DCGM draws inspiration from the VAE [97] and the DLGM [145], but differs in important ways. VAEs and DLGMs learn a probabilistic representation in a latent space by encoding data into a mean and a covariance. DCGMs, however, have a probabilistic representation that lies in the output layer (Figure 5.6). This enables our method to use an analytical likelihood function that allows for exact MLE, as opposed to VAEs and DLGMs, which require to learn via variational approximations.

VAEs typically work with diagonal covariances. DLGMs learn full covariance matrices by using a rank-1 approximation, which is just a slight improvement over a diagonal covariance [145]. We instead use a full-rank estimation that yields significantly better results, as will be shown later on. Additionally, Gaussian distributions in the output layer have been explored by other authors, but using either isotropic [38] or diagonal matrices [174].

Defining $\boldsymbol{\Sigma}_y$ as a full covariance matrix involves some practical challenges. First of all, $\boldsymbol{\Sigma}_y$ is a matrix of size $n \times n$, which could significantly increase the overall size of the output and the network. The network must also ensure $\boldsymbol{\Sigma}_y$ is always a symmetric positive definite matrix. The covariance matrix should also be numerically stable, since Equation 5.8 requires calculation of both its inverse and its determinant.

In order to face these challenges, we rely on the fact that the covariance can be represented with a Cholesky decomposition given by $\boldsymbol{\Sigma}_y = \mathbf{L}\mathbf{L}^T$, where \mathbf{L} is a lower triangular matrix. Numerical advantages related to this decomposition result in substantial computational simplifications. In order to avoid computing the inverse of the covariance matrix (first term of Equation

5.8), the model directly learns its Cholesky decomposition:

$$\Sigma_y^{-1} = CC^T. \quad (5.11)$$

Then, the determinant of the covariance matrix (which appears in the second term of Equation 5.8) can be easily calculated with just the main diagonal of the decomposition

$$\det(\Sigma_y) = \left(\prod_{i=1}^n C[i, i] \right)^{-2}. \quad (5.12)$$

For training and numerical stability purposes, the model splits the Cholesky decomposition C into two parts: its main diagonal $C_1 \in \mathbb{R}_{>0}^n$, and the rest of the elements $C_2 \in \mathbb{R}^{n(n-1)/2}$. Figure 5.6 shows the corresponding architecture of the DCGM. The rationale is that C_1 must include positive numbers exclusively, and thus ensure both Σ_y and Σ_y^{-1} are truly positive definite. Furthermore, using a small threshold $\lambda > 0$ can add numerical stability.

The network can be pre-trained by fixing all of the values in C_2 to 0. Once training convergence is achieved, the model will learn a conditional probabilistic representation where all the channels in the output are considered to be independent. This translates into a μ_y that gives accurate *maximum a posteriori* predictions, and also a diagonal covariance matrix Σ_y that captures uncertainty in each channel individually. Then the network can be fine-tuned by “unfreezing” C_2 , allowing it to also learn correlations between channels.

5.4.2 Experiments and Results

Experimental Setting

We conducted two different sets of experiments in order to evaluate the DCGM.

Experiments with Synthetic Data: The mapping function between x and y was known in advance, as well as the noise functions of x and y . For evaluation and visualization purposes, x and y were 1-D, i.e. $x \in \mathbb{R}, y \in \mathbb{R}$. We used Gaussian noise models for both x and y . Here the DCGM was implemented with the following network parameters (see Figure 5.6): one input layer, three sequential hidden layers (fully connected), and two output layers (one for μ_y and one for Σ_y). The Cholesky decomposition trick was not used since this is a 1-D problem. We used a dropout of 0.5 for all the hidden layers. The hidden layers used a rectified linear unit (ReLU) activation function, whereas μ_y and Σ_y used a sigmoid activation.

Experiments with Real Data: Real spectroscopic data was used for x and y , i.e. $x \in \mathbb{R}^m, y \in \mathbb{R}^n$. This was a more complex experiment since the mapping and noise functions were multidimensional and not known *a priori*. Specifically, we focused on the Cuprite scenario and used ASTER for low-resolution measurements and AVIRIS-NG for high-resolution spectra. Here the DCGM was implemented with the following network parameters (see Figure 5.6): one

input layer, two sequential hidden layers, and then three parallel convolutional decoders for μ_y , C_1 , and C_2 . We used a dropout of 0.7 for all the hidden layers. The hidden layers and C_1 used a rectified linear unit (ReLU) activation function, whereas μ_y and C_1 used a sigmoid activation.

In both experiments, we used a batch size of 4 and the Adam optimizer [96], with NLL as the loss function. The DCGMs were implemented in Keras [29]. We used early stopping to select the parameters that worked best with respect to the validation set in order to avoid overfitting. Overall, there was a good generalization. A 2.9 GHz Intel Quad-Core i7 laptop without a graphics processing unit (GPU) was sufficient to train the DCGM in both cases.

The DCGM was compared against four probabilistic baselines that are also based on Gaussian distributions (Figure 5.7).

- **Gaussian (G):** It concatenates the available and unavailable measurements, \mathbf{x} and \mathbf{y} respectively, and assumes both follow one Gaussian distribution, i.e. $p(\mathbf{x}, \mathbf{y}) = \mathcal{N}_{m+n}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$, where

$$\boldsymbol{\mu} = \begin{bmatrix} \boldsymbol{\mu}_x \\ \boldsymbol{\mu}_y \end{bmatrix}, \quad \boldsymbol{\Sigma} = \begin{bmatrix} \boldsymbol{\Sigma}_{xx} & \boldsymbol{\Sigma}_{xy} \\ \boldsymbol{\Sigma}_{yx} & \boldsymbol{\Sigma}_{yy} \end{bmatrix}. \quad (5.13)$$

MLE is performed by simply computing the joint population mean and covariance. Finally, the conditional mean $\boldsymbol{\mu}_y(\mathbf{x}) = \boldsymbol{\mu}_{y|x}$ and covariance $\boldsymbol{\Sigma}_y(\mathbf{x}) = \boldsymbol{\Sigma}_{y|x}$ are predicted using the standard formulas for conditional multivariate Gaussians [41]

$$\boldsymbol{\mu}_{y|x} = \boldsymbol{\mu}_y + \boldsymbol{\Sigma}_{yx} \boldsymbol{\Sigma}_{xx}^{-1} (\mathbf{x} - \boldsymbol{\mu}_x), \quad (5.14)$$

$$\boldsymbol{\Sigma}_{y|x} = \boldsymbol{\Sigma}_{yy} - \boldsymbol{\Sigma}_{yx} \boldsymbol{\Sigma}_{xx}^{-1} \boldsymbol{\Sigma}_{xy}. \quad (5.15)$$

- **Gaussian Mixture Model (GMM):** It concatenates \mathbf{x} and \mathbf{y} and learns a standard GMM for the joint distribution

$$p(\mathbf{x}, \mathbf{y}) = \sum_{i=1}^K w^i \mathcal{N}_{m+n}(\boldsymbol{\mu}^i, \boldsymbol{\Sigma}^i). \quad (5.16)$$

The conditional distribution also has the form of a GMM:

$$p(\mathbf{y}|\mathbf{x}) = \sum_{i=1}^K \bar{w}^i \mathcal{N}_n(\boldsymbol{\mu}_{y|x}^i, \boldsymbol{\Sigma}_{y|x}^i) \quad (5.17)$$

The parameters of each Gaussian component $\boldsymbol{\mu}_{y|x}^i, \boldsymbol{\Sigma}_{y|x}^i$ are given as in Equations 5.14 and 5.15, respectively. The conditional weights \bar{w}^i are updated using Bayes's theorem

$$\bar{w}^i = \frac{w^i p(\mathbf{x}|\boldsymbol{\mu}_{y|x}^i, \boldsymbol{\Sigma}_{y|x}^i)}{\sum_{j=1}^K w^j p(\mathbf{x}|\boldsymbol{\mu}_{y|x}^j, \boldsymbol{\Sigma}_{y|x}^j)}. \quad (5.18)$$

In our two experiments, we used 10-fold cross validation to estimate an adequate number of




Figure 5.7: 1-D example of the three baselines. The Gaussian predictor (black), defined by a single mean and a large variance, is most likely to underfit the data. The GMM (red) consists of multiple Gaussian distributions, and is likely to overfit the data. The Unimodal GMM is a combination of the previous baselines: it computes a single mean and variance that are biased toward the components of the GMM with larger weights.

mixtures.

- **Unimodal Gaussian Mixture Model (UGMM):** This method is a combination of the previous baselines. It estimates the overall mean $\mu_{y|x}$ and covariance $\Sigma_{y|x}$ of the conditional GMM, which are given by

$$\mu_{y|x} = \sum_{i=1}^K w^i \mu_{y|x}^i, \quad (5.19)$$

$$\Sigma_{y|x} = \sum_{i=1}^K w^i \left[\Sigma_{y|x}^i + \mu_{y|x}^i \mu_{y|x}^{i T} \right] - \mu_{y|x} \mu_{y|x}^T. \quad (5.20)$$

- **Variational Encoder (VE):** The spectral reconstruction method by Candela et al. that uses a modified VAE to predict high-resolution spectra from low-resolution inputs [18].

For Experiment #2, spherical, diagonal, and full covariance matrices were learned and tested for each of the previous algorithms. The differences between these covariances are illustrated in Figure 5.8.

We evaluated the performance of all these methods with the two following metrics that measure how well a model fits the data:

- **Root mean squared error (RMSE):** a measure of accuracy that computes differences between data and *maximum a posteriori* predictions.




Figure 5.8: 2-D example of spherical (left), diagonal (center), and full (right) covariance matrices. Covariances (red dotted lines) that are too simple may underfit the data (blue dots).

- **Negative log-likelihood (NLL):** a measure of how probable the data are with respect to the learned model, defined as the negative of the log-likelihood (Equation 5.5).

Experiments and Results using Synthetic Data

We started by conducting an experiment where a simple ground-truth model was known *a priori*. 1-D inputs were used to predict 1-D outputs. We designed this simple experiment for illustration purposes as it is easy to show how the ground-truth model look likes. In this experiment there was heteroscedastic uncertainty (noise) in the output and no uncertainty in the input. We compared the ground-truth model against DCGM and the four baselines methods that were described earlier. Since the input and output dimensions were 1-D, we trained a deep regression method instead of using a VE. It had the same architecture as the DCGM but without the variance output layer so it would perform non-probabilistic regression. From the ground-truth model we sampled 1000 data points for the training set, 1000 data points for the validation set, and 1000 data points for the test set. We used early stopping during training to avoid overfitting.

We show both a qualitative (Figure 5.9) and quantitative analysis of the results (Table 5.2). The Gaussian method is too simple and predicts a linear mean function and a constant variance. GMM does a better job, but it overfits the data. UGM is smoother than GMM and performs better when approximating the model from a visual perspective. Deep learning regression does a remarkable job when predicting the underlying function, but is unable to quantify uncertainty. DCGM has a slightly better performance in terms of regression, and more interestingly, is now able to learn the underlying noise function in the data.

We conducted a second experiment that was quite similar to the first one. It involved the same mapping function and output noise model, but the main difference was that it added heteroscedastic noise to the input as well. This forward propagation of uncertainty resulted in a

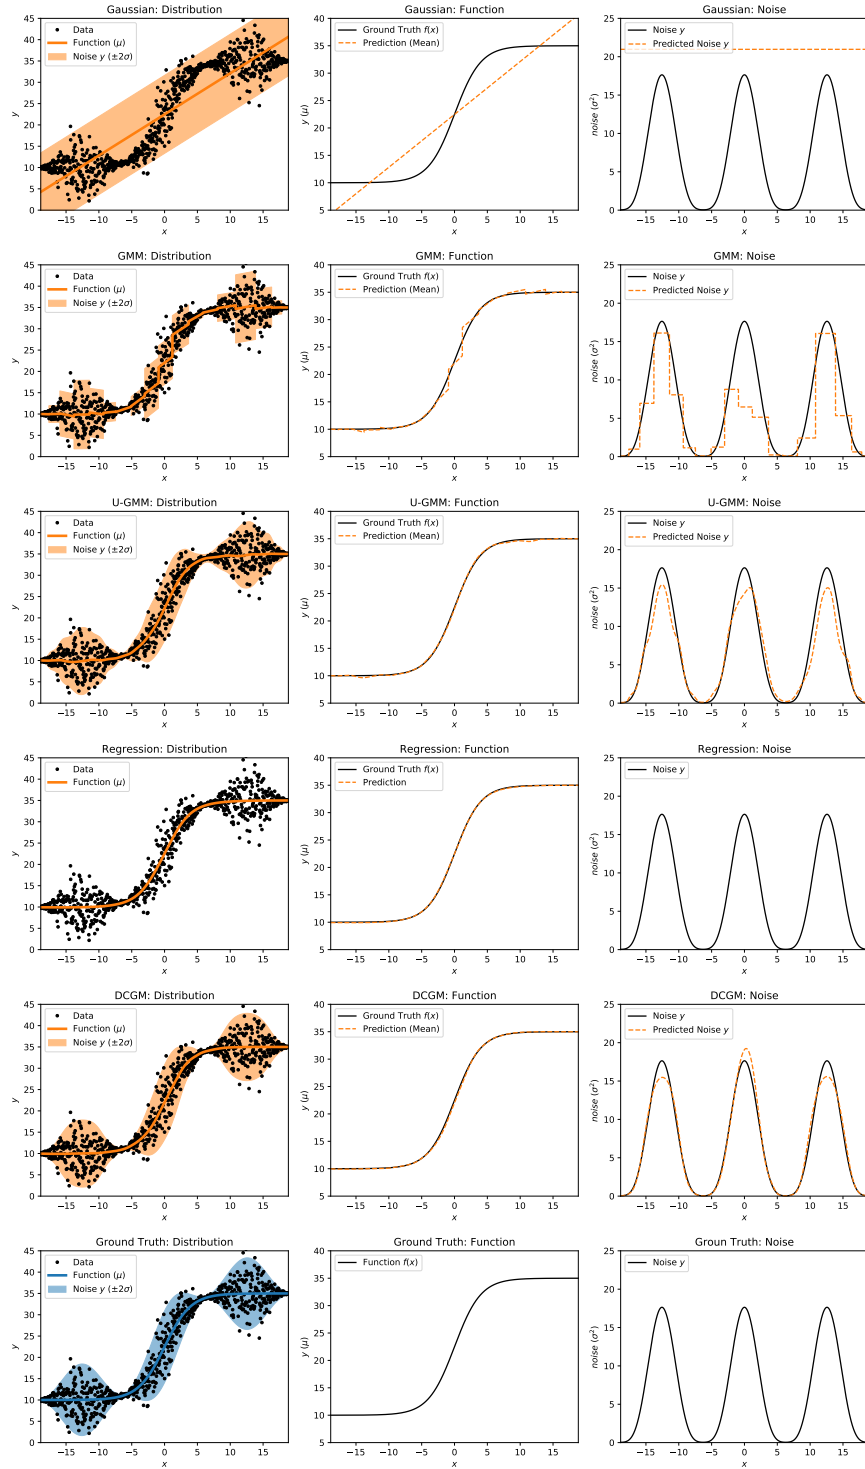


Figure 5.9: Experiment with synthetic 1-D data and heteroscedastic noise in the output. The objective is to learn, from the data, the underlying sigmoid function together with the noise model. Mean and noise predictions are shown for: Gaussian model, GMM, U-GMM, deep regression, DCGM, and ground truth.

Table 5.2: Experiments results for synthetic 1-D data set with heteroscedastic noise in the output. Results show the average performance plus-minus one standard deviation.

Metric	Gaussian	GMM	UGMM	Regression	DCGM	Ground Truth
RMSE	4.069 ± 2.077	20.994 ± 12.198	1.687 ± 1.845	1.684 ± 1.840	1.683 ± 1.849	1.681 ± 1.842
NLL	2.938 ± 0.433	1.876 ± 1.201	1.870 ± 1.103	-	1.823 ± 1.179	1.819 ± 1.179

Table 5.3: Experiments results for synthetic 1-D data set with heteroscedastic noise in both the input and output. Results show the average performance plus-minus one standard deviation.

Metric	Gaussian	GMM	UGMM	Regression	DCGM	Ground Truth
RMSE	4.266 ± 2.684	11.264 ± 6.689	2.576 ± 2.670	2.552 ± 2.683	2.536 ± 2.689	2.558 ± 2.720
NLL	3.036 ± 0.574	2.359 ± 1.132	2.494 ± 1.140	-	2.479 ± 1.491	2.507 ± 1.674

more complex noise model. In some cases, the overall ground-truth noise looked similar to the output noise. In other cases, it looked completely different. Qualitative results for the test set are shown in Figure 5.10, while quantitative performances appear in Table 5.3. In general, all methods perform worse than in the previous experiment because they have to learn a more complex and noisy underlying model. Again, we see that the Gaussian approach underfits the data. It is again clear that GMM overfits the data. UGMM is smoother and does a better job than GMM at learning the noise model from a visual perspective. Deep regression approximates the sigmoid function not as well as before, but still gets quite close; however, it is unable to model uncertainty. DCGM outperforms regression by a slight amount, and besides, it is able to capture the propagated aleatoric uncertainty in the data.

Experiments and Results using Real Data

We performed an experiment that involved the reconstruction (super resolution) of high-resolution spectroscopic data. The experiment focused on data of Cuprite, Nevada. Specifically, low-resolution ASTER data were used to infer high-resolution AVIRIS-NG spectra. More details can be found in Chapter 3. We performed a series of preprocessing operations on the data that seek to preserve spectral features [32] and reduce the impact of misguiding factors such as albedo and noise. Each spectrum was first scaled in the range $[0,1]$ with min-max normalization. Afterwards, we applied the MNF transformation to add noise robustness [63]. The used data set consists of over 6×10^6 spectra recorded by each instrument. However, most of them are redundant and can lead to overfitting due to strong correlations between neighboring locations. Moreover, there are a few dozen mineral classes at Cuprite with imbalanced instances [161]. We solved these issues by first building a more balanced data set with the help of the mineral classification algorithm Tetracorder [33]. We then eliminated spatially redundant measurements by randomly sampling 10,000 spectra from the scene. We divided them into three sets: training (5,000), validation (2,500), and test (2,500). The test set was sampled from the south region of Cuprite, whereas the other two sets from the north region.

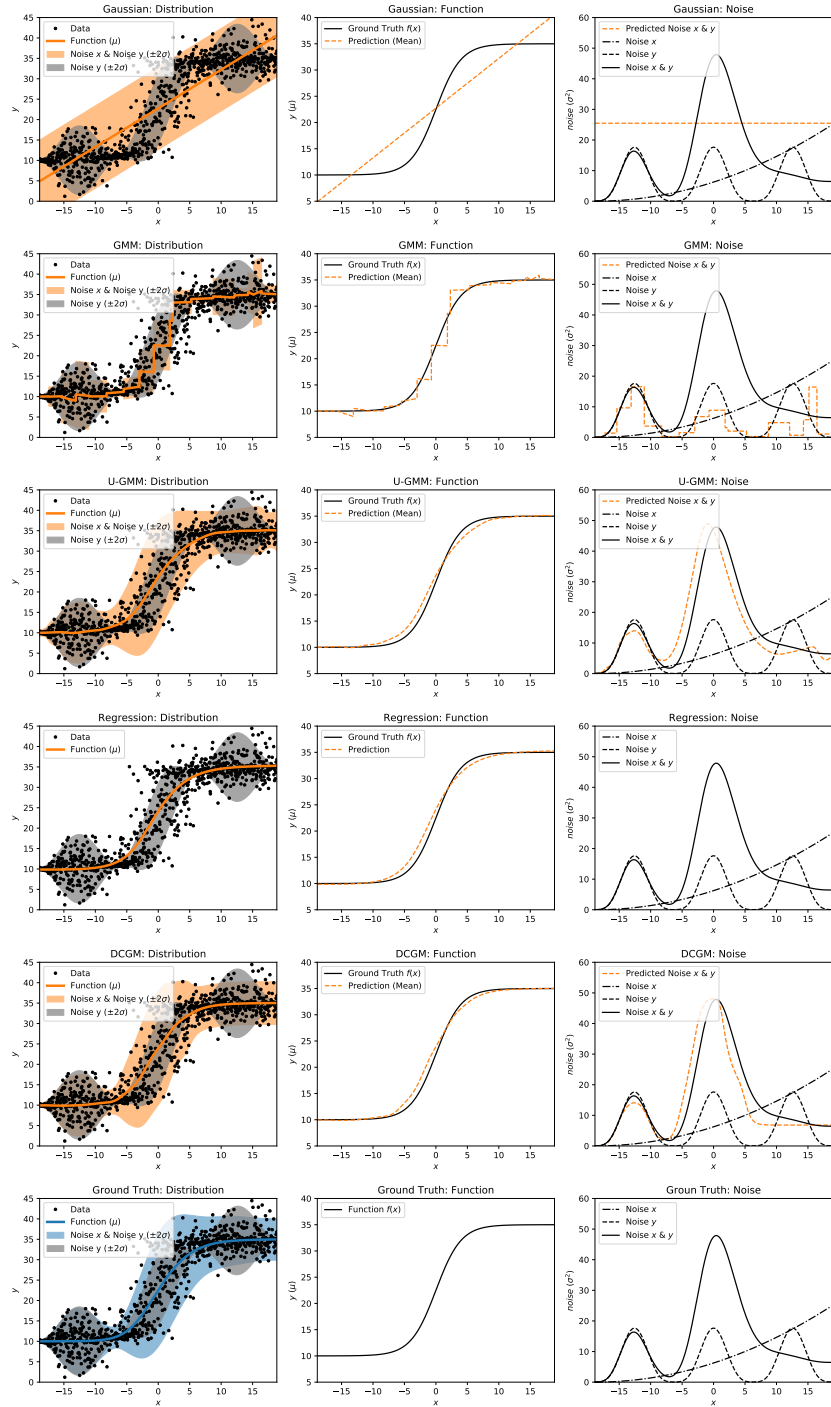


Figure 5.10: Experiment with synthetic 1-D data and heteroscedastic noises both in the input and output. The objective is to learn, from the data, the same underlying sigmoid function as in the previous experiment, but together with a more complex noise model resulting from forward uncertainty propagation. Mean and noise predictions are shown for: Gaussian model, GMM, U-GMM, deep regression, DCGM, and ground truth.

Table 5.4: Average model performance (smaller is better). The asterisk indicates cases where paired t-tests ($p > 0.05$) are not statistically significant.

Spherical Covariance Matrix					
Metric	G	GMM	UGMM	VE	DCGM
RMSE	0.0938*	0.0685	0.0616	0.0510	0.0479
NLL	-79.6	-158.5	-129.2	-	-144.3
Diagonal Covariance Matrix					
Metric	G	GMM	UGMM	VE	DCGM
RMSE	0.0938*	0.0602	0.0565	0.0479	0.0458
NLL	-95.1	-174.7	-149.6	-	-162.2
Full Covariance Matrix					
Metric	G	GMM	UGMM	VE	DCGM
RMSE	0.0566	0.0538	0.0509	0.0490	0.0447
NLL	-272.8	-308.5	-282.8	-	-297.3

The average performance of the different methods is shown in Table 5.4. Simpler covariance matrices lead to a deterioration in performance because they ignore valuable information. The DCGM has the best performance in terms of RMSE because of its superior *maximum a posteriori* predictions. This is consistent across different types of covariance matrices. GMM is the best in terms of NLL, probably because such low resolution inputs produce highly ambiguous relationships that are best modeled with a multimodal distribution (which needs to be tuned with the “right” number of components). However, DCGM still outperforms the other unimodal distributions. VE performs well in terms of RMSE, whereas its NLL is conventionally not defined. VE is consistent because of two reasons: it explicitly uses a RMSE loss function and its prior is a (spherical) standard Gaussian distribution. G, GMM, and UGMM perform poorly in terms of RMSE when using spherical and diagonal covariance matrices; but work well otherwise, especially GMM and UGMM because of the number of components they use. As expected, GMM is multimodal and thus outperforms UGMM in terms of NLL, but has more unstable *maximum a posteriori* predictions that result in a higher RMSE.

Let us discuss the potential of the DCGM as an analysis tool for geologists and spectroscopists via its human-interpretable results. We show how the DCGM uses ASTER to infer AVIRIS-NG spectra of two common mineral mixtures at Cuprite: mica with calcite and kaolinite with alunite (Figure 5.11). There are clear differences between both instruments in terms of resolution, but also regarding offsets and noise. The predictions are highly accurate since the ground truth spectra are well within the error bars. But more importantly, the covariance matrices contain useful information about the mineral and spectral features. For instance, there are high covariances near the main diagonal, which is to be expected because adjacent channels are strongly correlated. In the first mineral mixture, mica has a characteristic absorption feature around $2.2 \mu\text{m}$, whereas calcite around $2.35 \mu\text{m}$. Many fractional abundances are possible, re-

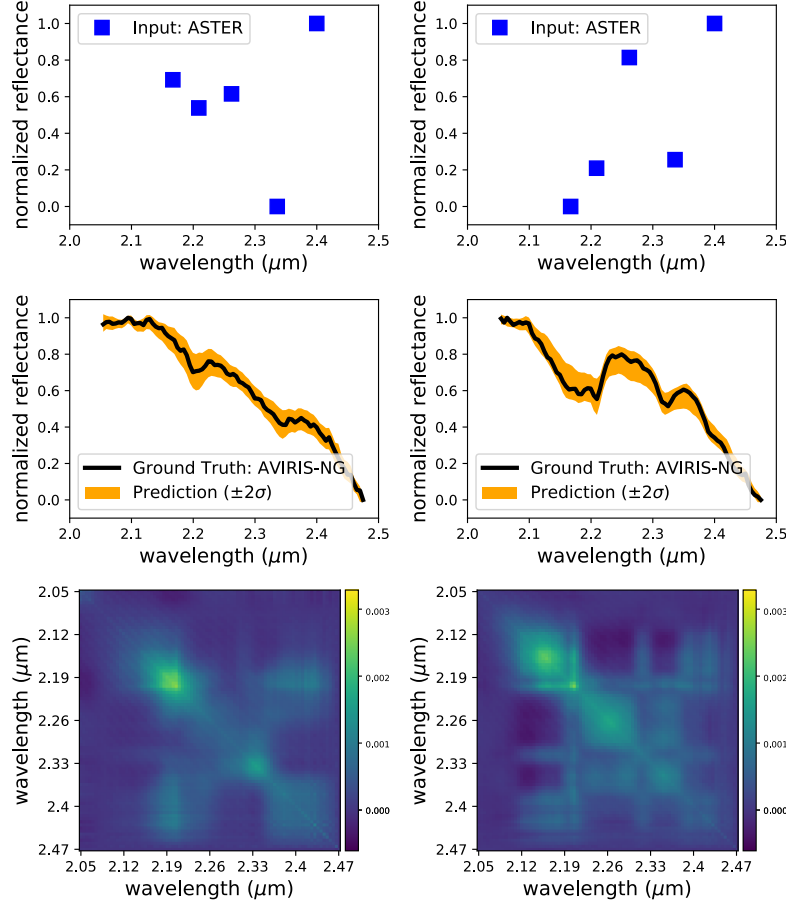


Figure 5.11: ASTER observations (top) being used to infer AVIRIS-NG measurements (middle) and their associated covariance matrices (bottom). This example shows two mineral mixtures at Cuprite: mica with calcite (left), and alunite with kaolinite (right).

sulting in features and wavelengths with high variances. Since this is a common mineral mixture, the features are strongly correlated. In the alunite and kaolinite mixture, both constituent minerals have broader and more complex features with overlapping wavelengths. The identification of these minerals, as well as the estimation of relative abundances, involves looking at subfeatures, hence the more complex correlations in the covariance matrix.

All these notions are common for a scientist during the analysis and interpretation of spectroscopic data, and are certainly visible in the model. In summary, the DCGM not only has accurate predictions, but also provides human-comprehensible explanations regarding correlated features and feasible alternatives.

5.5 Deep Conditional Dirichlet Model

5.5.1 Model and Architecture

The Dirichlet distribution is commonly used to model likelihoods over mixtures of components, and more generally over discrete probability distributions. It is often called “a distribution over distributions”. The domain can be viewed as the probabilities associated to K categories c_1, \dots, c_K , where the categories may be classes, components, etc. Formally, its support is $\mathbf{r} \in R \subset \mathbb{R}^K$, with $\sum_{k=1}^K r_k = 1, r_k \in (0, 1), \forall \mathbf{r} \in R$; where R takes the shape of an open standard $(K - 1)$ simplex. The Dirichlet distribution has a probability density function given by:

$$p(r_1, \dots, r_K | \alpha_1, \dots, \alpha_K) = \frac{\Gamma(\sum_{k=1}^K \alpha_k)}{\prod_{k=1}^K \Gamma(\alpha_k)} \prod_{k=1}^K r_k^{\alpha_k - 1}, \quad (5.21)$$

where Γ is the gamma function. The Dirichlet distribution is defined by the concentration parameters $\alpha_1, \dots, \alpha_K > 0$. A larger α_k means that category c_k is more likely, and vice versa. The mean and variance are given by:

$$\mathbb{E}[R_k] = \frac{\alpha_k}{\sum_{k=1}^K \alpha_k} \quad (5.22)$$

$$\mathbb{V}[R_k] = \frac{\frac{\alpha_k}{\sum_{k=1}^K \alpha_k} \left(1 - \frac{\alpha_k}{\sum_{k=1}^K \alpha_k}\right)}{\sum_{k=1}^K \alpha_k + 1} \quad (5.23)$$

Formally, we model the relationship between mixing ratios and spectra as a probabilistic conditional distribution $p_\theta(\mathbf{r}|\mathbf{y})$ where

$$\mathbf{r}|\mathbf{y} \sim \text{Dir}(\boldsymbol{\alpha}_\theta(\mathbf{y})). \quad (5.24)$$

The weights of the neural network are given by θ , and $\mathbf{y} \in Y \subset \mathbb{R}^n$ are high-resolution spectra. An illustration of the DCDM architecture is shown in Figure 5.12.

The DCDM directly maximizes a cumulative log-likelihood function where the training data set $\mathcal{D} = \{(\mathbf{y}^i, \mathbf{r}^i)\}_{i=1}^N$ consists of pairs of inputs \mathbf{y} (inputs) and ground-truth mixing ratios \mathbf{r}




Figure 5.12: The architecture of the Deep Conditional Dirichlet Model (DCDM) is shown on the left. The network receives a spectrum $\mathbf{y} \in Y$ as the input. The network then predicts the values of α , the concentration parameters of the Dirichlet distribution. These values are used to infer the mixing ratios \mathbf{r} for the components c_1, \dots, c_K . An example with three components ($K = 3$) is shown on the right.

(outputs). The log-likelihood function for a Dirichlet distribution is given by:

$$\mathcal{L}(\theta|\mathcal{D}) = \sum_{(\mathbf{y}, \mathbf{r} \in \mathcal{D})} \log p_{\theta}(\mathbf{r}|\mathbf{y}) \quad (5.25)$$

$$= \sum_{(\mathbf{y}, \mathbf{r} \in \mathcal{D})} \log \left(\frac{\Gamma(\sum_{k=1}^K \alpha_k(\mathbf{y}))}{\prod_{k=1}^K \Gamma(\alpha_k(\mathbf{y}))} \prod_{k=1}^K r_k^{\alpha_k(\mathbf{y})-1} \right) \quad (5.26)$$

$$= \sum_{(\mathbf{y}, \mathbf{r} \in \mathcal{D})} \left[\log \Gamma \left(\sum_{k=1}^K \alpha_k(\mathbf{y}) \right) - \sum_{k=1}^K \log \Gamma(\alpha_k(\mathbf{y})) + \sum_{k=1}^K (\alpha_k(\mathbf{y}) - 1) \log r_k \right]. \quad (5.27)$$

The first two terms control the magnitude of α , serving as regularizers. The last term measures the difference between the predicted and true mixing ratios in a similar way as cross-entropy.

Mixing ratios with values close to 0 or 1 lead to unstable likelihood values. Hence, we recommend to use a small threshold to avoid these situations. Furthermore, for training and numerical stability, we recommend to add a second term to the NLL loss function L . Specifically, one term that measures the error between the predicted and true mixing ratios:

$$L = -\mathcal{L}(\theta|\mathcal{D}) + \lambda \sum_{\mathbf{y}, \mathbf{r} \in \mathcal{D}} D_{KL}(\mathbf{r} || \hat{\mathbf{r}}(\mathbf{y})), \quad (5.28)$$

where $\lambda > 0$ is a tuning parameter and $\hat{\mathbf{r}}(\mathbf{y})$ is the predicted mean of the learned Dirichlet distribution.

5.5.2 Experiments and Results

Experimental Setting

We conducted two different sets of experiments in order to evaluate the DCDM.

Experiments with Synthetic Data: The mapping function between \mathbf{y} and \mathbf{r} was known in advance, as well as the noise functions of \mathbf{y} and \mathbf{r} . For evaluation and visualization purposes, $\mathbf{y} \in \mathbb{R}^2$, and \mathbf{r} belonged to the standard 1 simplex (the Dirichlet model turn into a simpler Beta distribution). We used Dirichlet noise models for both \mathbf{y} and \mathbf{r} , although other noise models for \mathbf{y} would have been possible as well. Here the DCGM was implemented with the following network parameters (see Figure 5.12): one input layer, three sequential hidden layers, one layer for $\alpha \in \mathbb{R}^2$, and one layer for the output \mathbf{r} . All the layers were fully-connected. We used a dropout of 0.5 for all hidden layers. The hidden layers used a rectified linear unit (ReLU) activation function. The layer for α used a modified sigmoid function with an output between 0 and 100. The output layer for α used a softmax activation function, ensuring that mixing ratios are inside a simplex. We used NLL as the loss function.

Experiments with Real Data: Real spectroscopic data was used for $\mathbf{y} \in \mathbb{R}^n$, while extensively validated mixing ratios were used for \mathbf{r} . This was a more complex experiment since the mapping and noise functions were multidimensional and not known *a priori*. Specifically, we focused on the Coral scenario, thus we used PRISM high-resolution spectra and their corresponding mixing ratios estimated during the CORAL mission for coral, algae, and sand. The DCDM architecture was essentially the same as the one we used for the experiment with synthetic data, but with an input of 92 spectral channels and an output of 3 endmembers. In order to achieve numerical stability, we used the modified NLL loss function shown in Equation 5.28 with $\lambda = 0.1$.

In both experiments, we used a batch size of 4 and the Adam optimizer [96]. The DCGMs were implemented in Keras [29]. We used a batch size of 4 and the Adam optimizer [96]. The DCGM was implemented in Keras [29]. We used early stopping to select the parameters that worked best with respect to the validation set in order to avoid overfitting. Overall, there was a good generalization. A 2.9 GHz Intel Quad-Core i7 laptop without a graphics processing unit (GPU) was sufficient to train the DCGM in both cases.

The DCDM was compared against the following methods for classification or unmixing.

- **Classification model:** It assumes that endmembers cannot be mixed together. Here we specifically use a logistic regression classifier.
- **Linear mixing model:** It assumes that mixtures between endmembers can be represented with a linear modal. We focus on three common methods for linear unmixing: least squares (LS), non-negative least squares (NNLS), and fully-constrained least squared (FCLS). More details can be found in Chapter 2.
- **Deep learning regression:** A neural network that learns how to unmix endmembers using a

Table 5.5: Experiment results for linear unmixing model. Results show the average performance plus-minus one standard deviation.

Metric	Classification	Least Squares	Regression	Discrete	DCDM	Monte Carlo
RMSE ($\times 10^{-2}$)	19.26 ± 10.98	5.178 ± 4.080	5.058 ± 4.141	5.227 ± 4.204	5.067 ± 4.136	5.178 ± 4.080
ACC	0.9360	0.9357	0.9355	0.9356	0.9357	0.9358
KLD ($\times 10^{-2}$)	48.71 ± 3.441	1.261 ± 1.700	1.221 ± 1.691	1.301 ± 1.793	1.227 ± 1.693	1.261 ± 1.700
NLL ($\times 10^{-1}$)	-	-	-	-1.365 ± 0.844	-1.403 ± 0.794	-1.380 ± 0.843

loss function based on the Kullback-Leibler divergence between the true and predicted mixing ratios. It has the same architecture as the DCDM, with the exception that it does not have a layer learning α parameters.

- **Discrete Dirichlet model:** The input space is discretized into regular intervals, and then Dirichlet parameter estimation is performed on each interval. 10-fold cross validation is used to estimate an adequate discretization of the input space.

We evaluated the performance of all these methods with the two four following metrics that measure how well a model fits the data:

- **Root mean squared error (RMSE):** a measure of accuracy that computes differences between data and *maximum a posteriori* predictions.
- **Classification accuracy (ACC):** a measure that selects the ground-truth mixing ratio with the highest value as the correct class to be predicted.
- **Kullback-Leibler Divergence (KLD):** a measure that treats true and predicted mixing ratios as probability distributions.
- **Negative log-likelihood (NLL):** a measure of how probable the data are with respect to the learned model, defined as the negative of the log-likelihood (Equation 5.27).

Experiments and Results using Synthetic Data

We started by conducting an experiment where a linear mixing model is known *a priori*. 2-D inputs were used to predict mixing ratios between 2 different components. We designed this simple experiment for illustration purposes as it is easy to show how the ground-truth model look likes. Each dimension in the input was directly related to each endmember. In this experiment there was heteroscedastic uncertainty (noise) in both the input and the output; these uncertainties were modeled as Dirichlet noise functions. Despite using a simple ground-truth model, it was still difficult to perform exact inference; hence we used results from Monte Carlo approximations for comparison. We compared this ground truth model against DCDM and the other methods that were described earlier. Since the mixing model is linear by design, LS, NNLS, and FCLS had the same performance; hence we only show results for LS. From the ground-truth model we sampled 1000 data points for the training set, 1000 data points for the validation set, and 1000 data points for the test set. We used early stopping during training to avoid overfitting.

Table 5.6: Experiment results for nonlinear unmixing model. Results show the average performance plus-minus one standard deviation.

Metric	Classification	Linear	Regression	Discrete	DCDM	Monte Carlo
RMSE ($\times 10^{-2}$)	12.85 ± 10.73	28.31 ± 12.89	3.938 ± 4.201	4.088 ± 4.394	3.955 ± 4.190	4.013 ± 4.374
ACC	0.9577	0.6281	0.9573	0.9562	0.9574	0.9563
KLD ($\times 10^{-2}$)	42.68 ± 39.69	26.07 ± 15.67	1.018 ± 1.761	1.098 ± 1.916	1.025 ± 1.739	1.070 ± 1.963
NLL ($\times 10^{-1}$)	-	-	-	-1.890 ± 1.229	-1.921 ± 1.171	-1.911 ± 1.263

Qualitative results for the test set are shown in Figure 5.13, while the corresponding quantitative performances appear in Table 5.5. The classification method is the best method in terms of classification accuracy, but the worst method with respect to the other metrics. This is because it fails to capture mixtures between endmembers. Least squares and deep learning regression achieve almost the same classification accuracy while doing a great job at predicting mixing ratios, but they are unable to learn the aleatoric noise in the data. The discrete Dirichlet model learns uncertainties, hence NLL can be computed; but it overfits the data so it has a somewhat worse performance in terms of the other metrics. DCDM not only models uncertainty, but is also more similar to the Monte Carlo approximation and achieves the lowest NLL scores overall. In terms of the other metrics, it is slightly outperformed by deep learning regression. This is most likely due to the regularization term in the log-likelihood (Equation 5.27). Note that the Monte Carlo method is not always the winner since it is an approximation of the ground truth model. It is also possible that the other methods learned biases present in the pseudo-random number generation process.

We conducted a second experiment that was quite similar to the previous one. The main difference was that it involved a nonlinear mixing model where the first component had a stronger presence. Qualitative results for the test set are shown in Figure 5.13, while the corresponding quantitative performances appear in Table 5.5. The results are similar to ones from the previous experiment. Overall, we confirm that DCDM adapts well to nonlinear models while learning the corresponding noise in the data. The most noticeable difference is that linear unmixing performs poorly, as would be expected for nonlinear models. One interesting fact is that classification is biased toward the first component as it is more dominant, representing a common scenario when using unbalanced data sets.

Experiments and Results using Real Data

We carried out an experiment that involved the spectral classification and unmixing of real data. The experiment focused on data from the CORAL mission. More details can be found in Chapter 3. We employed a data set that consists of 3000 spectra sampled from both PRISM flight lines (Heron Island and Kaneohe Bay), together with their respective mixing ratios of coral, algae, and sand. We performed a preprocessing step on the data to preserve spectral features and reduce the impact of misguiding factors such as albedo and noise. Concretely, each spectrum was scaled in

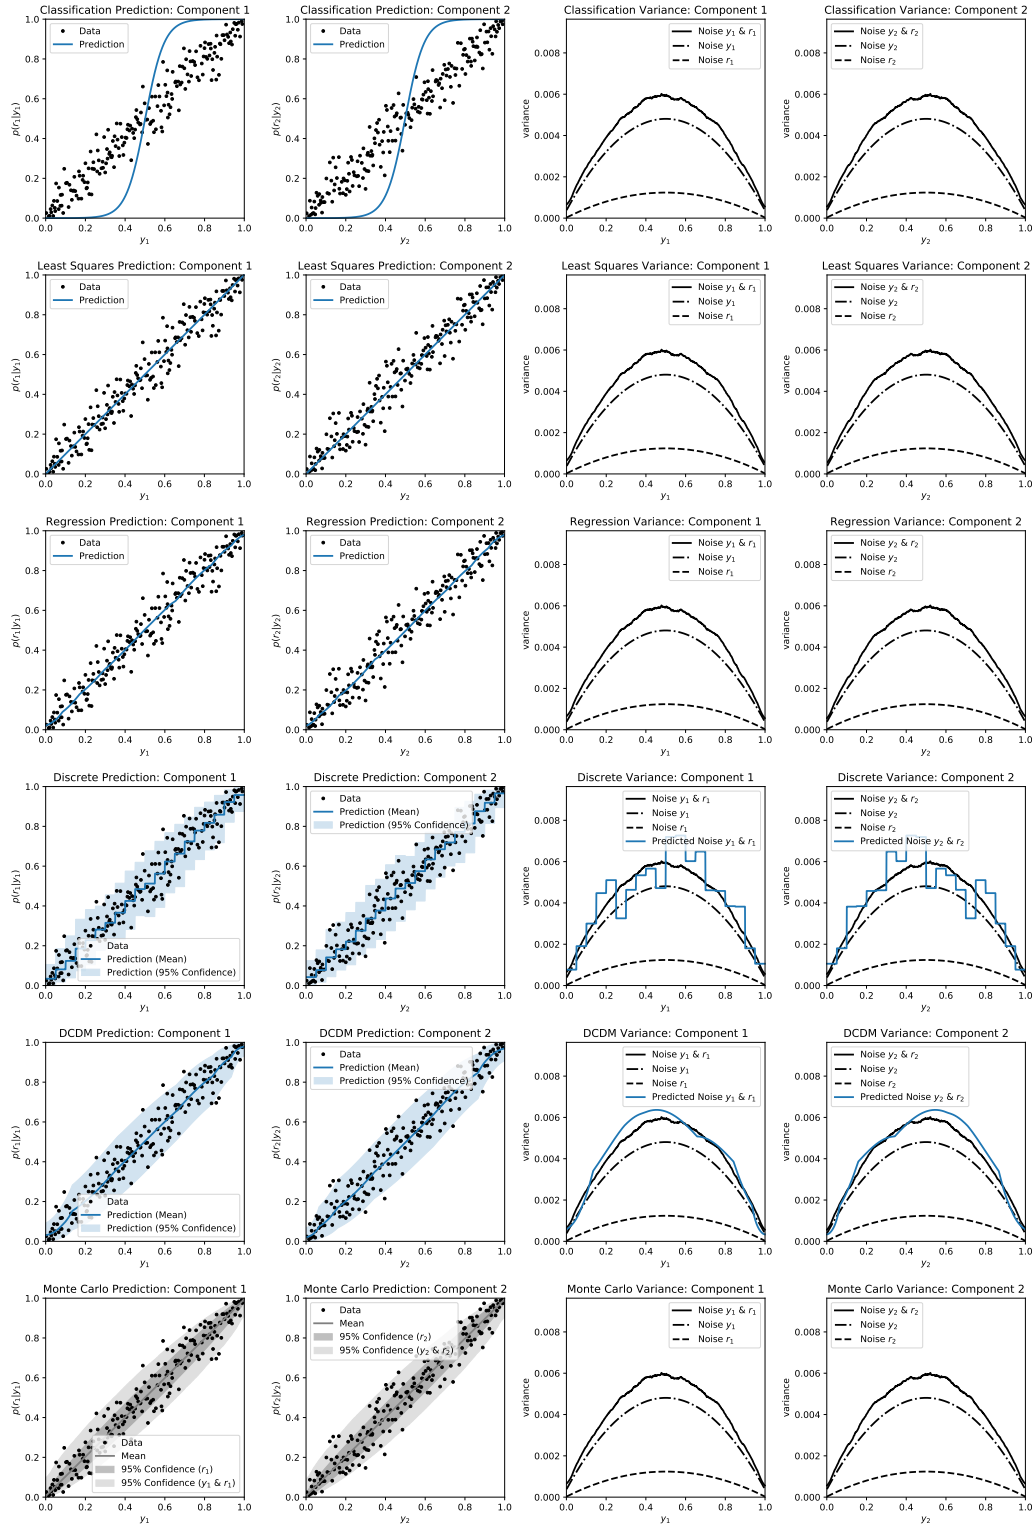


Figure 5.13: Experiment results for linear mixing model. The model consists of 2-D inputs that are used to unmix 2 components. Different methods use data to learn the mean and the variance of the model.

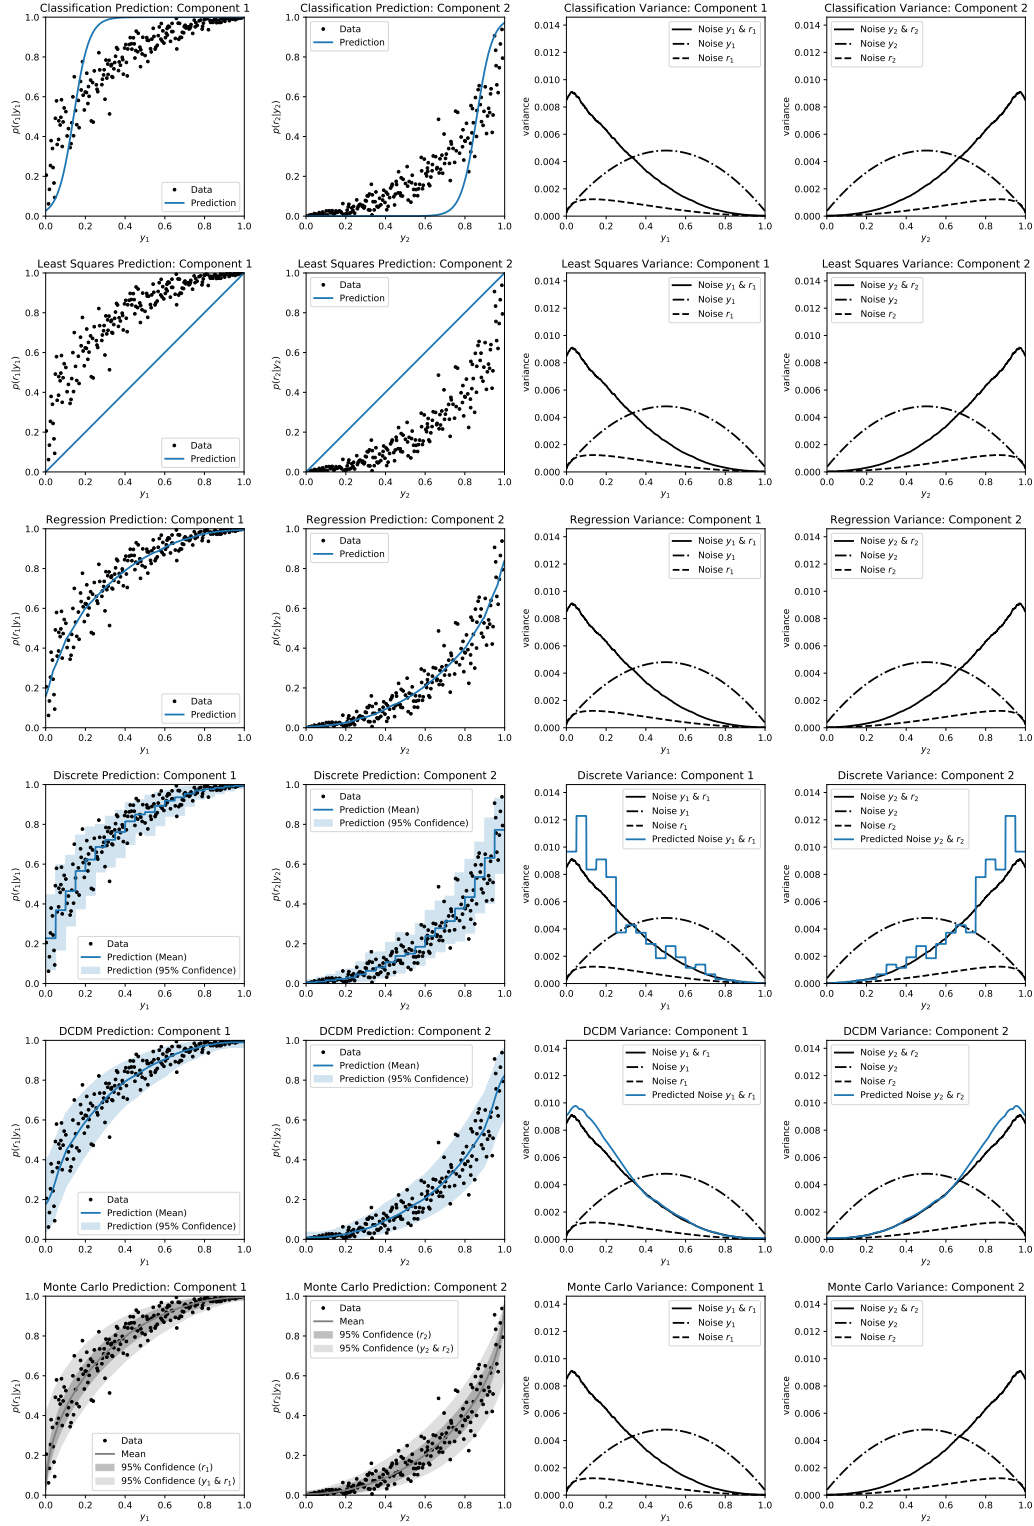


Figure 5.14: Experiment results for nonlinear mixing model. The model consists of 2-D inputs that are used to unmix 2 components. Different methods use data to learn the mean and the variance of the model.

Table 5.7: Experiment results using data from the Coral scenario. Results show the average performance plus-minus one standard deviation.

Metric	Classification	LS	NNLS	FCLS	Regression	DCDM
RMSE	0.263 ± 0.161	0.352 ± 0.519	0.464 ± 0.630	0.521 ± 0.875	0.105 ± 0.088	0.091 ± 0.086
ACC	0.747	0.539	0.489	0.346	0.923	0.936
KLD	0.119 ± 0.174	0.163 ± 0.195	0.243 ± 0.301	0.962 ± 0.797	0.031 ± 0.048	0.027 ± 0.055
NLL	-	-	-	-	-	0.193 ± 0.674

the range [0,1] with min-max normalization. 1000 spectra were used for the training set, 1000 for the validation set, and 1000 for the test set. The 3000 spectra were sampled using the k-means++ approach [4] to ensure that both their spectra and fractional abundances are diverse and representative. To avoid overfitting, we employed early stopping during training.

We compared the spectral unmixing performance of the aforementioned methods for spectral classification and unmixing. Given that the ground truth model was not known *a priori*, we are unable to provide the associated results. Additionally, we did not use the discrete Dirichlet model because of the high dimensionality of the data (92 channels). However, we employed and compared three different least squares methods. The associated quantitative results are shown in Table 5.7. Linear unmixing improves as more constraints are enforced, but at the cost of losing physical meaning (e.g., negative fractional abundances are impossible in practice). Interestingly, classification outperforms all the linear unmixing methods. Deep learning regression performs even better while still providing valid and consistent solutions because of its softmax output layer. These results suggest that the underlying model is strongly nonlinear. Another possible explanation is that linear methods are susceptible to the quality of the endmember extraction process, as well as to variability within classes. Finally, DCDM is the best method overall, with the particular benefit that NLL can be computed.

Let us discuss the potential of the DCDM as an analysis tool for marine biologists and spectroscopists via its human-interpretable results. To this end, we show how DCDM is using PRISM spectra to infer the corresponding fractional abundances of coral, algae, and sand (Figures 5.15 and 5.16). First, the DCDM seems to be accurate since true values lie inside the confidence intervals. When one endmember is essentially found in a pure state (Figure 5.15), predictions tend to reflect this with high confidence. More importantly, when spectra are mixed (Figure 5.16), predictions and their respective confidence vary accordingly. It is worth noting that this decreased confidence is also possible due to the effects of other endmembers besides coral, algae, and sand. These results indicate that DCDM is a more general tool than conventional classification. This is further confirmed by looking at predictions in the form of abundance maps (Figure 5.17). We see that conventional classification does a great job when dealing with pure spectra, but it often biases its predictions and fails to correctly estimate the abundances in mixed spectra. In contrast, the DCGM succeeds at predicting such varied abundances.




Figure 5.15: DCDM predictions of pure spectra. Spectra (top) are used to estimate the abundance (bottom) of coral (red), algae (green), and sand (blue). The DCDM performs accurate predictions and shows the confidence it has on them. The ground truth abundances are shown as white points. 95% confidence intervals were approximated via Monte Carlo sampling and appear as colored regions. These spectra are almost pure so predictions tend to be close to the corners with high confidence, resembling conventional classification tasks.

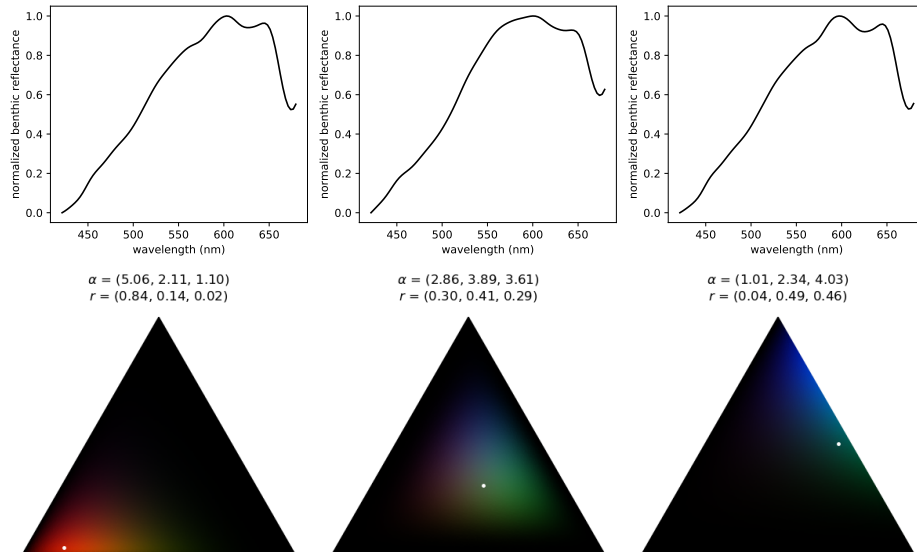

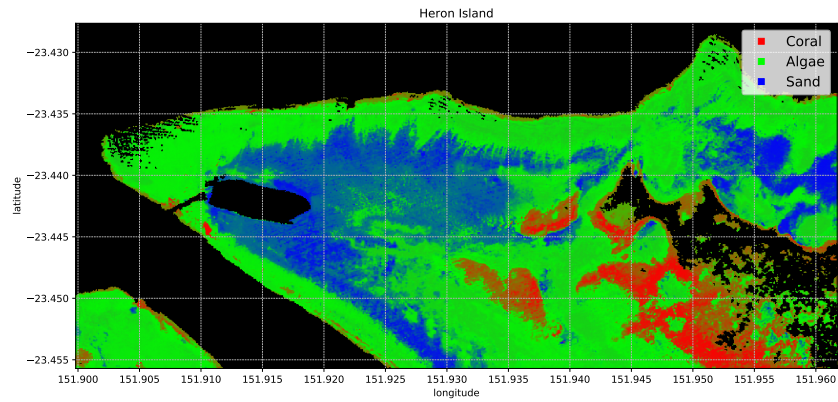



Figure 5.16: DCDM predictions of mixed spectra. Spectra (top) are used to estimate the abundance (bottom) of coral (red), algae (green), and sand (blue). The DCDM performs accurate predictions and shows the confidence it has on them. The ground truth abundances are shown as white points. 95% confidence intervals were approximated via Monte Carlo sampling and appear as colored regions. These spectra are mixed so predictions and confidences tend to be varied.



(a) Coral reef map produced during the CORAL mission. It serves as the ground truth in this work. Pure pixels appear as intense red, green and blue tones. Mixed pixels are shown as combinations of these three colors.



(b) Coral reef map as predicted by the DCGM. The map shows the mean of the learned Dirichlet distribution for each point. It resembles the ground truth and succeeds at mapping mixtures.



(c) Coral reef map generated via classification. It fails to model mixtures by biasing its predictions, resulting in more intense tones.

Figure 5.17: Coral, algae, and sand abundance maps of Heron Island, Australia.

5.6 Discussion and Conclusion

This chapter presents three deep generative models for scientific data analysis, with a special focus on spectroscopic data. It starts by presenting a version of the VAE that is used for learning spectral features, and consequently for dimensionality reduction. It then presents the DCGM, which can be used for probabilistic regression; and the DCDM, which can be used for classification, and more generally, unmixing. The DCGM and DCDM truly perform log-likelihood estimation, as opposed to the VAE and many other deep generative models, which rely on approximations. Furthermore, the DCGM and DCDM allow for the immediate visualization and interpretation of learned statistical dependencies.

Regarding the VAE, we observe that its nonlinear structure allows for richer feature representations when compared to linear methods such as PCA, especially when just a few number of features are extracted. It is worth noting that the VAE has a tendency to spread information across features in a uniform manner while filtering noise, whereas PCA first extracts features containing most of the information and then tends to learn increasing amounts of noise. This normalization process of the latent space of the VAE will result especially useful when we combine remote and *in situ* data in Chapter 5.

Our experiments reveal that DCGM and DCDM excel at making point predictions, like deep learning methods usually do. And more importantly, they learn models that can capture and show complex propagations of aleatoric uncertainties in the data. Both quantitative and qualitative results suggest that the DCGM and DCDM show promise as scientific analysis tools. Additionally, despite using unimodal distributions, they generate accurate and often better predictions than multimodal distributions, such as GMMs, which are prone to overfitting. Note that for the DCDM there is an important distinction between the number of endmembers (dimensionality of the Dirichlet distribution) and the number of probabilistic mixtures, for instance, of a Dirichlet mixture model.

We believe there are many interesting applications for the deep generative models presented in this work. For instance, they could be used for NASA’s Earth Surface Mineral Dust Source Investigation (EMIT), whose goal will be to study the role of atmospheric dust in Earth’s climate using an imaging spectrometer onboard the International Space Station [65]. Specifically, we could learn and visualize mineralogical and spectroscopic correlations between visible short-wave and thermal infrared measurements. The methods presented here also have direct applicability and benefit to NASA’s upcoming mission Surface Biology and Geology (SBG) Designated Observable [129]. The mission will use visible to shortwave infrared imaging spectroscopy and multispectral or hyperspectral thermal infrared imagery to address terrestrial and aquatic ecosystems and other elements of biodiversity, geology, volcanoes, and the water cycle.

Chapter 6

Combining Remote and In Situ Measurements for Adaptive Exploration

6.1 Introduction

The study of planetary surfaces has been made possible through the analysis of data collected by spacecraft and orbiters. For example, the instruments carried by the Mars Reconnaissance Orbiter (MRO) have been crucial in the mapping of landforms, minerals, and ice of Mars [186]. Despite the fact that these instruments provide useful information, factors such as sparsity, resolution, and noise leave uncertainty in the analysis of relatively low-resolution (10s of meters) remote sensing from orbit. For more definitive results, robotic explorers, such as Curiosity and ExoMars, are needed to collect high resolution, *in situ* measurements. Nonetheless, rovers face many operational challenges and constraints [50], so it is important to identify locations that maximize information value.

Remote sensing from airborne and spaceborne platforms have proven to be a useful tool for aspects of reef science [72]. Early applications of remote sensing to coral reef environments focused on mapping reef geomorphology and ecological zonation [64, 106, 127]. In the past few decades, much of remote sensing has focused on mapping habitats using qualitative descriptors comprising various combinations of substrate (e.g., sand, limestone, rubble), benthic functional type (e.g., coral, algae, seagrass), reef type (e.g., fringing, patch, barrier), and/or location within the reef system (e.g., slope, flat) [72]. The recent NASA Earth Venture Suborbital-2 (EVS-2) Coral Reef Airborne Laboratory (CORAL) mission focused on reef benthic functional type discrimination [77]. CORAL remote spectroscopic mapping was shown to be informative but lacked global and temporal coverage. Moreover, current remote sensing methods give only indirect information about reef condition, making *in situ* data critical.

Complete coverage of remote areas, while capturing relevant science features, can be achieved only through combination of remote sensing and *in situ*. Local observations are limited to a small area, but when they are combined with remote sensing, these observations can be put into

a broader context and extrapolated out to other regions. *In situ* samples tend to be scarce; therefore, it is important to select samples that refine the accuracy and resolution of remote sensing by adapting to the current state of knowledge of the environment.


In a few words, the motivation of this work is to develop an approach for wide-area mapping and optimal sampling. To this end, this research applies new techniques in orbital data analysis, probabilistic modeling for mapping, and decision theory for sample selection. Specifically, this chapter describes a machine learning model that actively combines *remote* and *in situ* measurements in order to efficiently map wide-areas via the extraction and and extrapolation of relevant science features. This chapter also develops automatic mission and path planning that can be executed by robotic explorers to refine and validate such maps. Finally, this chapter demonstrates the feasibility and performance of our approach via spectroscopic investigations. Simulation results are shown for the Cuprite, Mars, and Coral scenarios (more details appear in Chapter 4). Additionally, field experiments were conducted at Cuprite with the rover Zoë.

6.2 Related Work


Currently most works regarding adaptive exploration in robotics fail to exploit the power of remote sensing; consequently, they focus on mapping simple quantities such as salinity, temperature, altimetry, or dissolved oxygen measurements [11, 46, 159]. High-quality seafloor mapping has been achieved through the collection of *in situ* hyperspectral measurements and stereo imagery [13], but only at local scales through dense sampling. Recent efforts have developed robotic exploration models with high-level science objectives [3, 17]. However, they work with highly discretized models and fail to exploit valuable information that is available from remote sensing data.

Information from remote sensing data can be a powerful tool for extrapolation and wide-scale coverage. Fossum et al. [48] use remote sea-surface temperature images to infer spatial patterns in the evolution of oceanographic conditions. Rao et al. [141] and Shields et al. [153] combine remote bathymetry with *in situ* imagery for benthic habitat mapping. Thompson et al. present a linear model that combines remote and *in situ* measurements for spatio-spectral mapping, but at the cost of downsampling high-resolution data and thus losing valuable information [166]. Foil and Wettergreen [47] and Thompson et al. [168] effectively utilize contextual information for surface classifications with spectroscopic data.

Gaussian process (GP) regression has been widely used in spatial statistics [36, 142]. GPs also provide a convenient framework for adaptive sampling [102]. Hence, GPs have been very popular for robotic exploration; applications include binary terrain classification [162], ocean temperature mapping [10, 11] and plant phenotyping [107]. However, most work regarding GPs (especially for AUVs) involves mapping scalar fields [11, 46, 159]. In contrast, the objective of this work is to map more complex science features.



(a) The mapping process starts with poor predictions.



(b) The mineral map improves as more *in situ* samples are collected and extrapolated.

Figure 6.1: Example of active robotic exploration for mineral mapping.

Feature extraction approaches have been used for non-scalar benthic habitat mapping. Rao et al. use a deep learning architecture that does not allow for adaptive sampling [141]. Shields et al. use a similar deep learning approach, but they have recently taken some first steps toward enabling adaptive sampling by quantifying model uncertainty through numerical approximations [153].

6.3 Probabilistic Mapping Model

This chapter describes an approach to spectral feature and composition mapping over large areas. To this end, it combines two different types of spectral measurements: low-resolution (or multispectral) remote data $\mathbf{x} \in X \subset \mathbb{R}^m$, and high-resolution (or hyperspectral) *in situ* data $\mathbf{y} \in Y \subset \mathbb{R}^n$. Remote spectra are available beforehand for many spatial locations $\mathbf{l} \in L \subset \mathbb{R}^2$ (latitude and longitude), but in general their spectral resolution does not permit composition analysis. In contrast, *in situ* spectra tend to have a spectral resolution that is sufficient for composition analysis, but unfortunately only a scarce number of samples can be collected by a robotic explorer such as a rover or an AUV. The objective not only is to extrapolate *in situ* samples over large areas with the assistance of remote sensing, but also to construct maps that can easily adapt and improve with new information. A relevant example involving mineral mapping is shown in Figure 6.1.

This approach consists of two main steps: spectroscopic mapping and composition mapping (Figure 6.2). First, remote sensing is used as a prior to extrapolate high-resolution *in situ*




Figure 6.2: Pipeline for spectroscopic and composition mapping. First, high-resolution spectra are collected by the robotic explorer. Then, features are extracted and extrapolated to generate feature maps. Afterward, spectroscopic maps are reconstructed from these feature maps. Finally, predicted spectra are used to estimate composition.

spectral measurements over large areas. Then, predicted high-resolution spectra is employed for composition analysis. These steps are decoupled to allow for alternative composition analysis techniques. Specifically, both spectroscopic and composition mapping are achieved by combining different probabilistic machine learning algorithms. GP regression is used to spatially extrapolate *in situ measurements*. However, the objective is to perform spectroscopic mapping , i.e., map non-scalar data. This problem is solved with feature extraction techniques, which reduce the dimensionality of spectra by deriving a subset of non-redundant features. Once spectra are reconstructed from extrapolated features, this approach employs learning-based methods to estimate composition from spectra. The rest of this section describes spectral feature extraction, Gaussian process regression, and learning-based composition analysis in more detail.

6.3.1 Spectral Feature Extraction

In situ spectroscopic measurements tend to have a high resolution, which may lead to the *curse of dimensionality*. However, many of these channels and wavelengths are highly correlated, allowing for the application of *dimensionality reduction* techniques. To this end, this approach uses a variational autoencoder (VAE) [97, 145], a neural network that converts a set of high-resolution observations $\mathbf{y} \in Y \subset \mathbb{R}^n$ into a set of lower dimensional features $\mathbf{z} \in Z \subset \mathbb{R}^d$, where $d \ll n$. More details can be found in Chapter 5. Note that automated feature extraction requires the exact same input size. Nonetheless, it is possible to generalize and scale these algorithms to data from similar spectrometers via resampling or interpolation; this is shown later in the experimental results validating this approach.

It is important to mention that feature extraction ignores spatial information and neighboring context; it only considers individual spectra. The Location Guided Autoencoder (LGA) by Yamada et al. considers spatial autocorrelations by assuming that spatial information is known *a priori* and by making a simple modification to regular VAEs, potentially improving performance [182]. However, for generalization purposes, spectral features are assumed to be learned from spectral libraries that may not necessarily contain the corresponding spatial information in ad-

vance. The next section describes how our approach models and learns spatial autocorrelation between spectra in an adaptive manner using GPs.

6.3.2 Gaussian Processes for Spatio-Spectral Regression

This approach relies on GPs [142] for spatio-spectral regression, that is, for learning the spatial distribution of spectra throughout the scene. GPs are a powerful technique for extrapolation, as well as for the refinement of maps as new data is collected. GPs are typically used for mapping scalar values, but spectroscopic mapping involves multivariate regression. This challenge is overcome by using GP regression to learn the distribution of low-dimensional features Z instead. Moreover, dimensionality reduction with VAEs uncorrelates the learned feature representation. This property simplifies the problem substantially by allowing for the utilization of a small number of $i = 1, 2, \dots, d$ independent GPs.

We next provide a brief explanation regarding the specific elements of our GP regression model. If needed, extensive and canonical documentation can be found in [142]. Formally, we define an input vector that concatenates spatial coordinates and remote measurements as $\mathbf{v} = (\mathbf{l}, \mathbf{x}) \in V \subset \mathbb{R}^{2+m}$. This formulation is based on the work by Thompson [162]. We assume there exists a latent function $f^i : \mathbb{R}^{2+m} \rightarrow \mathbb{R}$ that maps an input \mathbf{v} to each feature z^i :

$$z^i = f^i(\mathbf{v}) + \epsilon_i. \quad (6.1)$$

Each GP learns a distribution over the values that f^i can take. A GP is defined by a prior mean μ^i and covariance function K_θ^i , that is:

$$f^i(\mathbf{v}) \sim \mathcal{GP}(\mu^i(\mathbf{v}), K_\theta^i(\mathbf{v}, \mathbf{v}')). \quad (6.2)$$

We assume that the prior mean is zero because of the way features are normalized by the VAE. For the covariance matrix, we rely on the widely used *squared exponential kernel*. Similarly as Thompson [162], we define an anisotropic kernel that distinguishes between spatial and spectral dimensions as follows:

$$K_\theta^i(\mathbf{v}, \mathbf{v}') = \theta_0^i \exp \left(-\frac{\|\mathbf{l} - \mathbf{l}'\|_2^2}{2(\theta_l^i)^2} - \frac{\|\mathbf{x} - \mathbf{x}'\|_2^2}{2(\theta_x^i)^2} \right), \quad (6.3)$$

where $\theta^i = (\theta_0^i, \theta_l^i, \theta_x^i)$ are the kernel hyperparameters for each GP. Once all of these variables are specified, the posterior (or predictive) distribution of f^i can be computed given a training set consisting of pairs of inputs and features $\mathcal{D} = \{(\mathbf{v}^i, z^i)\}_{i=1}^N$. The posterior distribution is defined by a posterior mean $\hat{\mu}^i$ and posterior covariance $\hat{\Sigma}^i$ as follows:

$$f^i(\mathbf{v}|\mathcal{D}) \sim \mathcal{GP}(\hat{\mu}^i(\mathbf{v}|\mathcal{D}), \hat{\Sigma}^i(\mathbf{v}, \mathbf{v}'|\mathcal{D})). \quad (6.4)$$

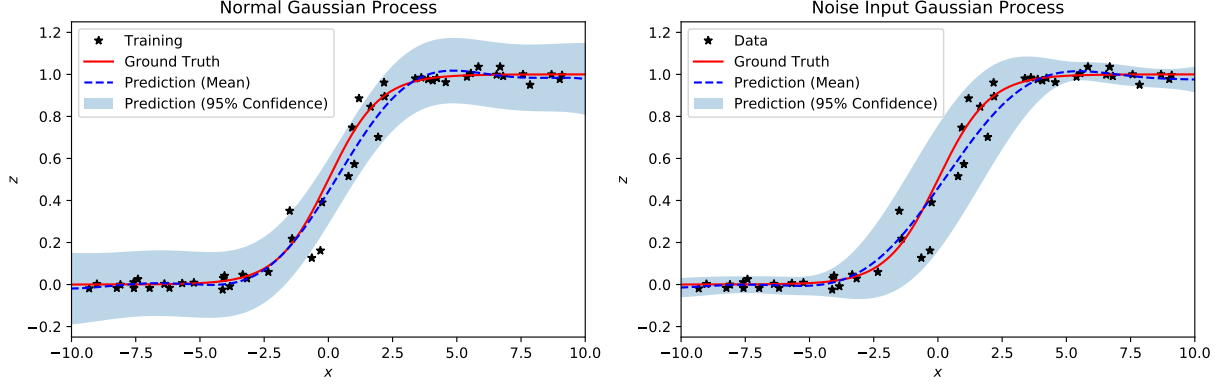


Figure 6.3: Noise modeling in Gaussian processes. GPs typically learn one constant noise value for the output, but this sometimes underfits the data (left). GPs that account for input noise tend to better fit the data [119] (right).

Here we use the GP variant for noisy observations [142]. Therefore, we include a hyperparameter representing the standard deviation of the output (feature) noise $\sigma_{noise,out}^i$. It is worth noting that this noise is assumed to be homoscedastic (constant). Modeling of the overall aleatoric uncertainty can be improved by taking into account the input noise too. For this we use the method by McHutchon and Rasmussen [119], which adds and learns an input noise hyperparameter $\sigma_{noise,in}^i$. This method works with a homoscedastic input noise model and uses first-order Taylor expansions to model the process of forward uncertainty propagation from $\sigma_{noise,in}^i$ and $\sigma_{noise,out}^i$. Noise modeling is improved with this approach that learns both the input and output noise in the data (Figure 6.3). Finally, all the GP hyperparameters are estimated by maximizing the log-likelihood of the observed data [142]. This function is typically optimized with gradient descent techniques. We specifically use the L-BFGS-B algorithm [126] since it handles simple bound constraints on variables.

6.3.3 Learning-Based Spectral Composition Analysis

In this work we explore different learning-based methods for composition analysis and apply them to separate scenarios. We especially focus on probabilistic methods. For the Cuprite scenario, we use a logistic regression classifier to estimate mineralogy; it was trained with scarce data from the USGS spectral library version 7 [100]. For the Mars scenario, a Gaussian classifier is trained using scarce data from the CRISM spectral library [173]. For the Coral scenario, there is plenty of labeled data. Hence, we employed the Deep Conditional Dirichlet Model (DCDM) to perform spectral unmixing (Chapter 5). More details regarding these composition analysis methods are provided in the sections describing the experiments and results for each of the three scenarios.

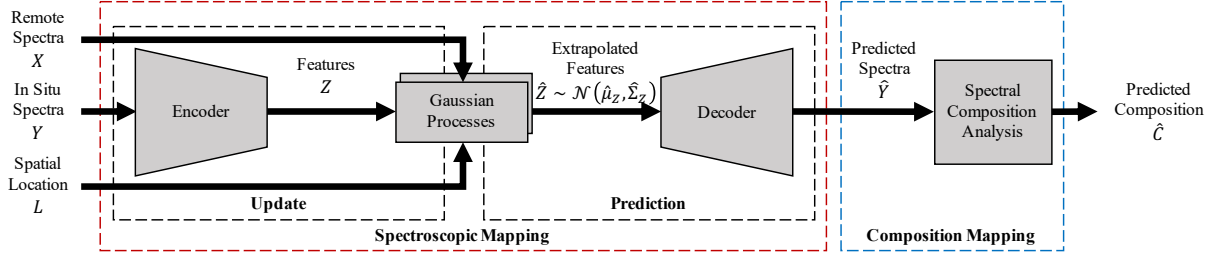


Figure 6.4: The active mapping model. It combines spatial coordinates, remote and *in situ* spectra. It integrates feature extraction, Gaussian process regression, and learning-based composition analysis.

6.3.4 Integrated Model

The overall model integrates feature extraction, GP regression, and learning-based composition analysis (Figure 6.4). The whole learning process, as in recursive Bayesian estimation, consists of two steps: update and prediction [170] (Figure 6.5). In the update step, the learning model is improved when the robot collects *in situ* measurements (Figure 6.5a). The d independent GPs are updated using the features that are extracted with the encoder of the VAE, along with the associated spatial coordinates and remote measurements. In the prediction step, the model uses this new knowledge to better reconstruct the scene (Figure 6.5b). First, the GPs predict the features of each point in the map using a normal distribution. Then, these features are passed through the decoder of the VAE and reconstructed as high resolution spectra. Finally, predicted spectra are used to estimate composition via classification or spectral unmixing.

6.4 Informative Exploration


The robotic explorer aims to collecting the most meaningful science measurements, that is, the ones that better explain and reconstruct the entire scene. In Bayesian experimental design, information-driven action selection can be formulated as the minimization of *posterior entropy*, which measures the uncertainty of a set of variables of interest A after collecting new information B [25, 35, 102].

The uncertainty of the model throughout the scene is quantified using differential entropy [35], which is given by [102]:

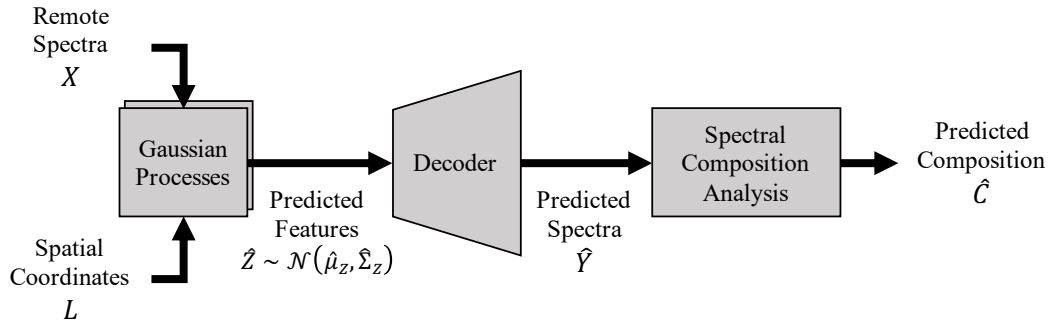
$$H(A|B) = \frac{1}{2} \log \left\{ (2\pi e)^{|A|} \left| \hat{\Sigma}(A|B) \right| \right\}. \quad (6.5)$$

Entropy computation (Equation 6.5) becomes expensive as the cardinality of A increases, specifically the calculation of $|\hat{\Sigma}(A|B)|$. To address challenge, we minimize the upper bound on the entropy instead, that is:

$$H(A|B) \leq \hat{H}(A|B). \quad (6.6)$$



(a) The update step. Each time the robot collects an *in situ* spectrum, the corresponding features, remote spectrum, and spatial coordinates are used to update the multiple Gaussian processes in the model.



(b) The prediction step. The Gaussian processes use remote spectra X and their corresponding spatial coordinates to extrapolate a map of predicted features. These predicted features are passed through the decoder to reconstruct high-resolution spectra. Finally, predicted spectra are used to estimate composition.

Figure 6.5: Learning process of the active mapping model. As in recursive Bayesian estimation, it consists of two steps: update and prediction [170].

An upper bound can be easily constructed with just the main diagonal of the predicted covariance matrix. A corollary of Hadamard's inequality [67] states that if $\hat{\Sigma}(A|B)$ is a positive-semidefinite matrix, then its determinant is bounded by the product of its main diagonal:

$$\left| \hat{\Sigma}(A|B) \right| \leq \prod_{j=1}^{|A|} \hat{\sigma}_j(A|B)^2. \quad (6.7)$$

Consequently, the following bound is derived:

$$\hat{H}(A|B) = \frac{1}{2} \log \left\{ (2\pi e)^{|A|} \prod_{j=1}^{|A|} \hat{\sigma}_j(A|B)^2 \right\} \quad (6.8)$$

$$= \sum_{j=1}^{|A|} \frac{1}{2} \log \left\{ (2\pi e) \hat{\sigma}_j(A|B)^2 \right\}, \quad (6.9)$$

which is equivalent to assuming that all points in A are conditionally independent given B . This procedure simplifies computation substantially because it avoids calculating the full covariance matrix, as well as its determinant. Furthermore, this upper bound does not accumulate error as it is renewed every time the GP is updated.

In this work, the random variable of interest is the spatial distribution of features throughout a map, which we assume is composed of a large, yet finite set of points $\mathcal{M} = \{\mathbf{v}_1, \dots, \mathbf{v}_p\}$. The new information is given by *in situ* coordinates and measurements collected by the robot, i.e. $\mathcal{P} = \{[l_1, \mathbf{y}_1], \dots, [l_k, \mathbf{y}_k]\}$. Since the overall model consists of d independent GPs, the posterior entropy of the map is additive across features [35]. Then, the upper bound is given by the following expression:

$$\begin{aligned} \hat{H}(\mathcal{M}|\mathcal{P}) &= \sum_{i=1}^d \left[\sum_{j=1}^{|\mathcal{M}|} \frac{1}{2} \log \left\{ (2\pi e) \hat{\sigma}_j^i(\mathcal{M}|\mathcal{P})^2 \right\} \right] \\ &= \sum_{i=1}^d \sum_{j=1}^{|\mathcal{M}|} \frac{1}{2} \log \hat{\sigma}_j^i(\mathcal{M}|\mathcal{P})^2 + \frac{d|\mathcal{M}|}{2} \log 2\pi e. \end{aligned} \quad (6.10)$$

where $\hat{\sigma}_j^i(\mathcal{M}|\mathcal{P})^2$ is the predicted variance of each feature z^i in every point of the map. This entropy bound is used to compute and represent the uncertainty in the map given a robot's path (Figure 6.6).

This research formulates exploration as an optimization problem where the robot's path should minimize the posterior entropy of the map, or more precisely, its upper bound:

$$\begin{aligned} \min_{\mathcal{P}} \quad & \hat{H}(\mathcal{M}|\mathcal{P}) \\ \text{s.t.} \quad & \text{Cost}(\mathcal{P}) \leq \text{Budget} \end{aligned} \quad (6.11)$$

There are many informative path planning algorithms that could be used to solve this type of optimization problem, each with its own advantages and drawbacks. There are methods that assume independence between sampling locations, which is usually an oversimplification in informative exploration scenarios [27, 183]. There are also near-optimal greedy algorithms that work well on Gaussian processes [102, 107]. Other approaches may be computationally intensive, but po-

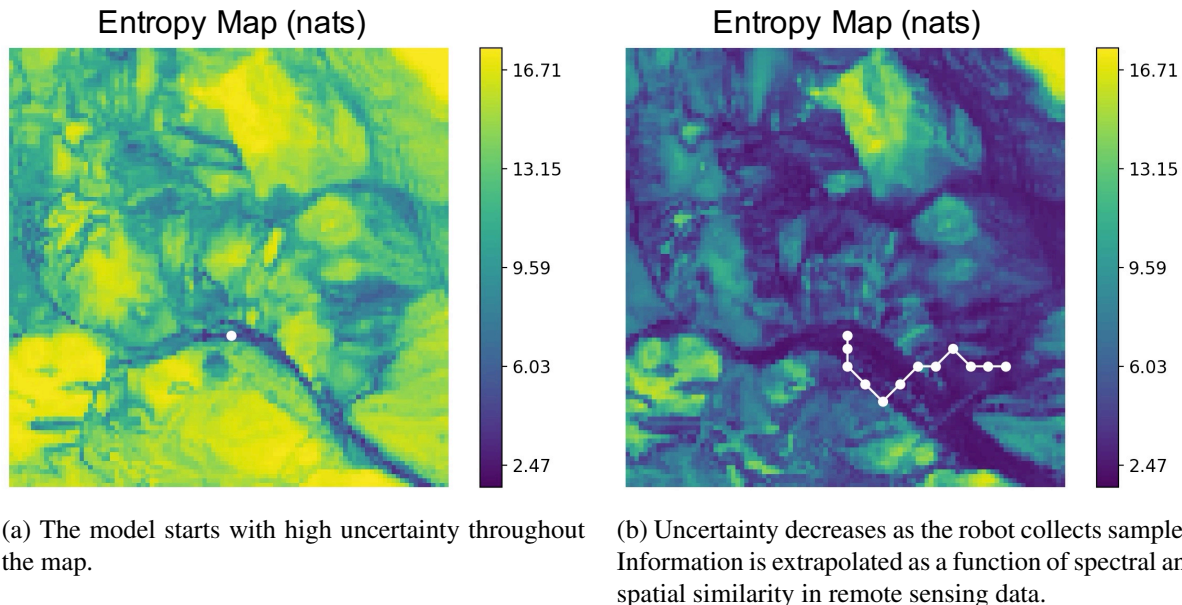


Figure 6.6: Visualization of model uncertainty in terms of differential entropy. Light/yellow tones indicate high uncertainty areas, whereas dark blue tones denote low uncertainty.

tentially closer to optimality, such as branch and bound (B&B) techniques for both discrete [10] and continuous [83] space representations. Finally, we are especially interested in Monte Carlo tree search (MCTS) planners that have been applied to geologic exploration scenarios [3, 99].

6.5 Cuprite Scenario

The feasibility and performance of our high-level autonomy approach were demonstrated via spectroscopic investigations at Cuprite, Nevada [20, 169]. More details regarding the Cuprite scenario can be found in Chapter 4.

6.5.1 Experimental Setting

The experiments were designed to evaluate the performance of both the learning and exploration strategies in terms of science productivity. We first perform a thorough analysis in simulations, and then show field results with the rover. Zoë. In the simulation experiments, the ASTER instrument provided low-resolution remote measurements, whereas AVIRIS-NG spectra were proxy for *in situ* measurements. In the field experiments, remote data also came from ASTER, but the rover collected *in situ* spectra with an onboard ASD instrument. Since ASD samples were sparse, AVIRIS-NG served as the ground truth for validation purposes.

The exploration space was represented as an 8-connected grid using a pixel step size of 2,

where each pixel has a resolution of 15 meters. We focused on three subregions of Cuprite that are mineralogically diverse and relatively traversable, and called them A, B, and C. Both the VAE and the GPs were trained with a subset of the data that was withheld from the experiments. The VAE learned how to extract $d - 6$ features from AVIRIS-NG spectra. We found this value to work well in practice. Then, the model consisted of 6 independent GPs that were pre-trained with the same data set, and later fine-tuned online in order to better adapt to the incoming flow of data. A logistic regression classifier was used for composition analysis. It was trained using data from the USGS spectral library version 7 [100] in order to identify three dominant minerals in the selected subregions: alunite, kaolinite, and opal (also known as hydrated silica).

Three informative path planning algorithms were compared in this study:

- **Random:** It sequentially samples a random neighboring location until the sampling budget is exhausted. It corresponds to an uninformed sampling strategy. This is a *science-blind* baseline.
- **Greedy:** It makes the locally optimal choice using a one-step lookahead. This is a *myopic* exploration strategy.
- **MCTS:** A Monte Carlo tree-search planner that creates a long term strategy by planning ahead for possible observations in the future. The planner was implemented by Kodgule et al. [99] using a four-step lookahead. This is a *non-myopic* path planner.

Three metrics were used to evaluate the performance of the planners. For normalization purposes, we compute the averages with respect to the total number of points in the map.

- **Entropy:** It is a measure of uncertainty in the model. It is *directly* minimized by the planners (Equation 6.11) and is calculated without a ground truth.
- **Error:** It is a measure of the spectral reconstruction error of the scene in terms of root mean squared error (RMSE). It should be *indirectly* minimized by the planners since it requires a ground truth for comparison.
- **Kullback-Leibler divergence:** It is a smooth function that measures the difference between the true and predicted class probabilities (Equation 2.8). It should be *indirectly* minimized by the planners.

6.5.2 Simulation Experiments and Results

Three hundred different traverses were simulated using random starting locations. Starting locations were evenly spaced throughout each subregion; end goals were not specified to allow for more freedom during exploration. We imposed a constraint of 20 samples per traverse. We defined a constraint of 20 samples per traverse. Additionally, a digital elevation model (DEM) from ASTER was used to apply a slope constraint of 18° .

Training and simulation were performed using a laptop computer with an Intel i7 processor (2.9 GHz quadcore) and 16GB of memory. Each waypoint was computed within just a few seconds or minutes, depending on the complexity of the path planner.




Figure 6.7: Scatter plots of average reconstruction error vs. entropy for the three sites: A (left), B (center), and C (right). In all cases, there is a strong correlation as measured by the Pearson correlation coefficient.

We first calculated the correlation between entropy and reconstruction error throughout the simulations. The three sites present high Pearson correlation coefficients, (0.862, 0.903, 0.919) respectively, indicating there exists a strong positive correlation (Figure 6.7). This suggests that entropy is a suitable objective function for reconstructing spectra in the scene. We also see high correlation coefficients for the three sites (0.769, 0.886, 0.874) comparing entropy and the Kullback-Leibler divergence (Figure 6.8). These coefficients are slightly lower, probably due to the additional step between spectral prediction and composition analysis. Regardless, these results reaffirm the convenience of using entropy to guide exploration.

We then evaluated the performance of the different planners. In all cases, entropy, reconstruction error, and the Kullback-Leibler divergence show decreasing trends as more samples are collected (Figures 6.9, 6.10, and 6.11). It is clear that Random is the worst planner. We compare Greedy (myopic) against MCTS (non-myopic) via paired t-tests (Figures 6.12 and 6.13). Results indicate that MCTS is not necessarily better at the beginning, but it outperforms Greedy in the long term. Overall, we confirm that algorithms with farther planning horizons tend to perform better as they exploit more information from the model.

6.5.3 Field Experiments and Results

The rover Zoë performed a series of autonomous traverses in the field in order to validate the simulation results. In general, computation time was not an issue. Depending on the complexity of each path planner, the planning process for each waypoint took just a few seconds or minutes to run.

The rover carried an ASD Fieldspec Pro spectrometer with a 1° foreoptic mounted on a pan-tilt mechanism. Each measurement consisted in a *panorama* that acquired and averaged 16 spectra, moving the pan-tilt actuator in a 5×5 m raster pattern. The spectrometer was calibrated for each panorama using an onboard white reference. Panoramas were monitored both automatically and manually; whenever faulty (e.g. there was a shadow on the calibration




Figure 6.8: Scatter plots of average Kullback-Leibler divergence vs. entropy for the three sites: A (left), B (center), and C (right). In all cases, there is a strong correlation as measured by the Pearson correlation coefficient.




Figure 6.9: Plots and 1-sigma error bars for entropy as a function of collected samples per traverse. Results are shown for the three sites: A (left), B (center), and C (right). Smaller is better.




Figure 6.10: Plots and 1-sigma error bars for reconstruction error as a function of collected samples per traverse. Results are shown for the three sites: A (left), B (center), and C (right). Smaller is better.

target), they were recalibrated and retaken. The ASD spectra were interpolated and downsampled so they would have the same resolution as AVIRIS-NG. This was done in order to use the same models from the simulated experiment.




Figure 6.11: Plots and 1-sigma error bars for Kullback-Leibler divergence as a function of collected samples per traverse. Results are shown for the three sites: A (left), B (center), and C (right). Smaller is better.




Figure 6.12: Paired entropy t-tests between the Greedy and MCTS planners. Results are shown for the three sites: A (left), B (center), and C (right). Points below the threshold indicate that MCTS performs better with a 95% confidence.




Figure 6.13: Paired reconstruction error t-tests between the Greedy and MCTS planners. Results are shown for the three sites: A (left), B (center), and C (right). Points below the threshold indicate that MCTS performs better with a 95% confidence.




Figure 6.14: Paired reconstruction error t-tests between the Greedy and MCTS planners. Results are shown for the three sites: A (left), B (center), and C (right). Points below the threshold indicate that MCTS performs better with a 95% confidence.

Zoë used GPS localization to save its geographic coordinates whenever it collected an ASD panorama. Zoë used a planar homography to convert from geographic coordinates (latitude and longitude) to map pixel coordinates (row and column), enabling the spatial registration of *in situ* and remote spectra, respectively. Inherent GPS error was negligible since the used spectroscopic maps have a resolution of 15 m/pixel. The rover considered that it had reached a waypoint whenever it was less than 5 m away from it.

Safe rover navigation was one main challenge. The terrain was not always traversable at Cuprite. As in the simulations, we used the ASTER DEM to try to enforce a slope constraint of 18° . It proved somewhat useful, but not sufficient because of the poor DEM resolution (15 m/pixel). Human supervision and manual overriding were occasionally required for risk and obstacle avoidance.

As opposed to the simulation studies, Zoë was unable to execute hundreds of traverses; instead, it generated and traversed three paths, each one running a different planner. This process was repeated at the three different sites.

We show a few ASD spectra that were collected by Zoë at Cuprite, as well as ASTER and AVIRIS-NG measurements that correspond to the same geographic coordinates (Figure 6.15). The ASD and AVIRIS-NG spectra align well, showing that the latter is an adequate validation source. Additionally, *maximum a posteriori* predictions of the model not only are accurate, they also remove some of the noise in the ASD measurements.

The rover traverses at the three sites in Cuprite are illustrated in Figures 6.16, 6.17, and 6.18; along with their respective entropy, reconstruction error, and Kullback-Leibler divergence plots. The plots are not as smooth as the ones from the simulations because they correspond to a single trial. While entropy is monotonic, this is not always true for reconstruction error nor Kullback-Leibler divergence. Reasons include *in situ* measurement noise, slight underfitting or overfitting of the model, and small differences between ASD and AVIRIS-NG spectra (Figure 6.15). Nevertheless, they all show decreasing trends. Overall, we observe that the non-myopic




Figure 6.15: Examples of spectroscopic measurements from the field experiment at Cuprite, Nevada. These plots correspond to three distinctive minerals in the scenes: alunit (left), opal (center), and kaolinite (right).

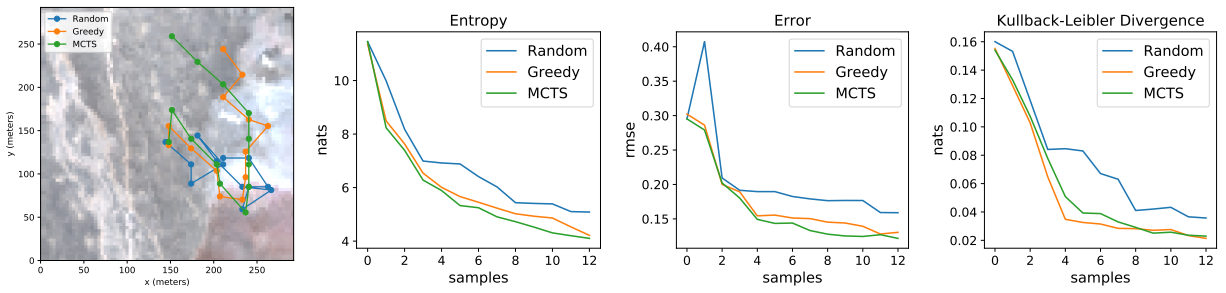


Figure 6.16: Results from the field experiment with the rover Zoë at site A: traversed paths, entropy, reconstruction error, and Kullback-Leibler divergence.

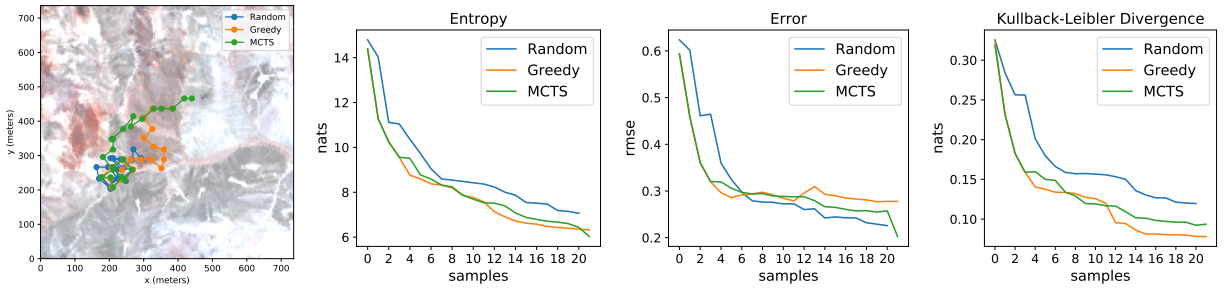


Figure 6.17: Results from the field experiment with the rover Zoë at site B: traversed paths, entropy, reconstruction error, and Kullback-Leibler divergence.

MCTS planner has the best performance in most cases, followed by the greedy planner.

Finally, we include a visual representation of the active learning process during the traverse where Zoë was running MCTS at site A (Figure 6.19). At first, when just one sample has been collected, the performance of spatio-spectral regression is quite poor. There is still plenty of uncertainty in the map, especially in places that are significantly different to the sampled location. After the rover has collected diverse samples, we observe that the model's predictions resemble the AVIRIS-NG ground truth even at unsampled locations, that the mineral map has improved,




Figure 6.18: Results from the field experiment with the rover Zoë at site C: traversed paths, entropy, reconstruction error, and Kullback-Leibler divergence.

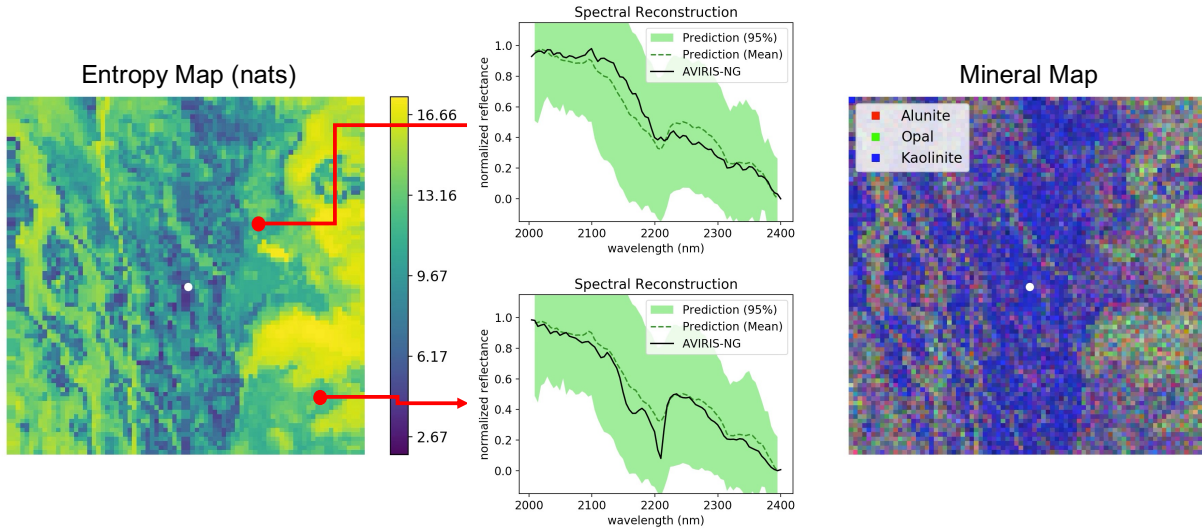
and that entropy has decreased substantially throughout the map.

6.6 Mars Scenario

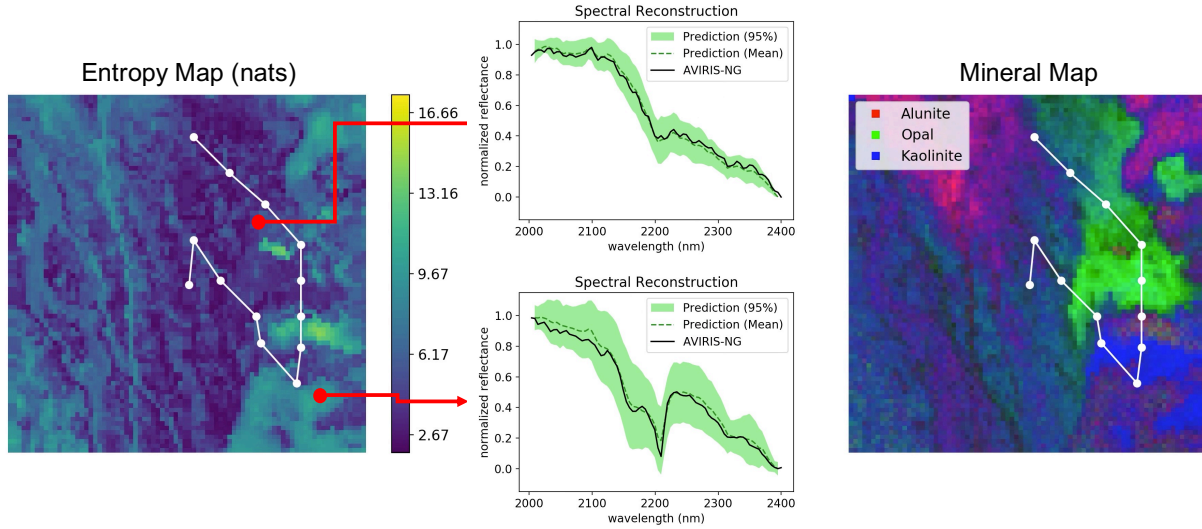
Our approach was further validated in a simulation study that involved spectroscopic investigations on Mars [19]. More details regarding the Mars scenario can be found in Chapter 4.

6.6.1 Experiments

The HRSC instrument provided low-resolution remote measurements, whereas CRISM spectra were proxy for *in situ* MastCam-Z measurements. Three hundred rover traverses were simulated at each of the two sites. For each site, we selected a 3×3 km mineralogically diverse subregion. The exploration space was represented as an 8-connected grid with a step size of 150 meters (3 pixels). Starting locations were evenly spaced throughout each subregion; end goals are not specified to allow for more freedom during exploration. We imposed a constraint of 20 samples per traverse. In this study, both the VAE and the GPs were trained with a subset of the data that was withheld from the experiments. L2 normalization was applied to both HRSC and MastCam-Z spectra to allow the model to focus on spectral features rather than albedo values. The VAE learned how to extract $d = 3$ features from MastCam-Z spectra. We found this value to work well in practice. Then, the model consisted of 3 independent GPs that were pre-trained with the same data set, and later fine-tuned online in order to better adapt to the incoming flow of data. A Gaussian classifier was used for composition analysis, consisting of a uniform prior distribution, 5 means corresponding to the 5 selected minerals from the CRISM spectral library (Fe/Ca-carbonate, Mg-carbonate, Fe-olivine, Mg-olivine, and hydrated silica), and isotropic (spherical) covariances $\Sigma = (0.02)^2 I_3$. For validation, we assumed that the ground truth classes are given by the output of the Gaussian classifier when applied to simulated MastCam-Z spectra, which is also assumed to be the ground truth *in situ* spectra. Training and simulation were performed using a laptop computer with an Intel i7 processor (2.9 GHz quadcore) and 16GB of memory.



(a) At the beginning of the rover traverse, entropy is high and the model's spectral and mineral predictions are poor.



(b) By collecting and extrapolating *in situ* samples, the robot reduces entropy in the map and improves the model's spectral and mineral predictions.

Figure 6.19: Evolution of entropy, spectral reconstruction, and mineral mapping at site A while running MCTS. Confidence bounds were approximated via Monte Carlo sampling.

Each waypoint was computed within just a few seconds or minutes, depending on the complexity of the path planner.

Three informative path planning algorithms were compared in this study:

- **Random:** It sequentially samples a random neighboring location until the sampling budget is exhausted. It corresponds to an uninformed sampling strategy. This is a *science-blind* baseline.
- **Greedy:** It makes the locally optimal choice using a one-step lookahead. It resembles oppor-




Figure 6.20: Scatter plots of average reconstruction error vs. entropy for the two sites: Jezero Crater (left) and Nili Fossae (right). In all cases, there is a strong correlation as measured by the Pearson correlation coefficient.

tunistic observations that are sometimes performed by Mars rovers. This is a *myopic* exploration strategy.

- **MCTS:** A Monte Carlo tree-search planner that creates a long term strategy by planning ahead for possible observations in the future. The planner was implemented by Kodgule et al. [99] using a four-step lookahead. This is a *non-myopic* path planner.

Three metrics were used to evaluate the performance of the planners. For normalization purposes, we compute the averages with respect to the total number of points in the map.

- **Entropy:** It is a measure of uncertainty in the model. It is *directly* minimized by the planners (Equation 6.11) and is calculated without a ground truth.
- **Error:** It is a measure of the spectral reconstruction error of the scene in terms of root mean squared error (RMSE). It should be *indirectly* minimized by the planners since it requires a ground truth for comparison.
- **Kullback-Leibler divergence:** It is a smooth function that measures the difference between the true and predicted class probabilities (Equation 2.8). It should be *indirectly* minimized by the planners.

6.6.2 Results

An example of a rover traverse at Nili Fossae together with its corresponding mineral mapping process is shown in Figure 1.3 (Chapter 1). Mineral signatures are successfully identified and extrapolated throughout the entire scene. Just a few samples (20) are sufficient to map the mineralogy of most of the 3×3 km area.




Figure 6.21: Scatter plots of average Kullback-Leibler divergence vs. entropy for the two sites: Jezero Crater (left) and Nili Fossae (right). In all cases, there is a strong correlation as measured by the Pearson correlation coefficient.

We first calculated the correlation between entropy and reconstruction error throughout the simulations (Figure 6.20). The Pearson correlation coefficients are 0.885 and 0.826 for Jezero Crater and Nili Fossae, respectively. These values indicate a positive correlation between HRSC and CRISM data, and confirm that entropy is a suitable objective function for spectroscopic mapping. We see a similar results when calculating the correlation between entropy and the Kullback-Leibler divergence (Figure 6.21). The two sites also have high correlation coefficients (0.825, 0.788). These coefficients are slightly lower, probably due to the additional step between spectral prediction and composition analysis. Regardless, these results reaffirm the convenience of using entropy to guide exploration.

We then evaluated the performance of the three planners (Figure 6.22). In all cases, entropy, reconstruction error, and KLD show decreasing trends as more samples are collected. Error and divergence converge faster than entropy. It is clear that Random is the worst planner. MCTS, the algorithm with the farthest planning horizon, performs best at the end. Note that Greedy outperforms MCTS during the first few samples; this is to be expected because Greedy selects points that immediately provide high rewards, whereas MCTS computes a long-term strategy.

We include a visual representation of the active learning process during a traverse at Nili Fossae while running MCTS (Figure 6.23). Specifically, we show evolution of the entropy, reconstruction error, and Kullback-Leibler divergence (KLD) maps. At first, when just one sample has been collected, the performance of is quite poor. There is still plenty of uncertainty in the map, especially in places that are significantly different to the sampled location. After the rover has collected diverse samples, we observe that entropy, error, and KLD have decreased substantially throughout the map.

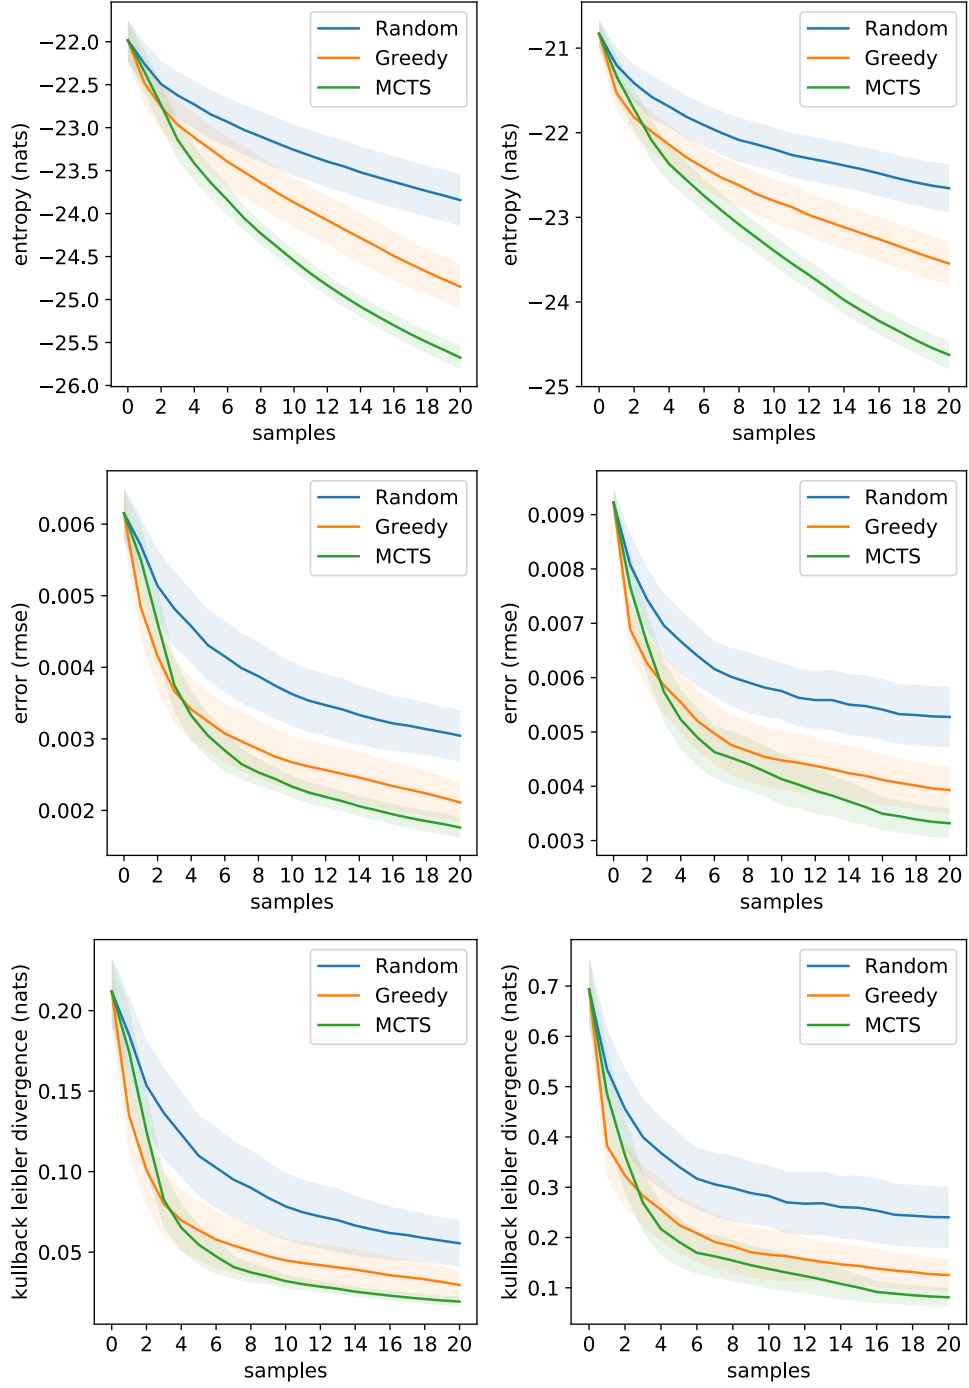




Figure 6.22: Performance plots and 1-sigma error bars for Jezero Crater (left column) and Nili Fossae (right column). Average entropy (top row), reconstruction error (middle row), and KLD (bottom row) as a function of collected samples per traverse. All results have a statistically significant difference ($\alpha = 0.05$).



(a) At the beginning of the traverse entropy, error, and KLD are high throughout the map.



(b) By collecting and extrapolating *in situ* samples, entropy is reduced in the map and the model's spectral and class predictions improve.

Figure 6.23: Evolution of entropy, spectral reconstruction, and Kullback-Leibler divergence (KLD) at Nili Fossae while running MCTS.

6.7 Coral Scenario

The presented approach was further explored and tested in a simulation study that involved spectroscopic investigations for coral reef mapping [22]. More details regarding the Coral scenario can be found in Chapter 4.

6.7.1 Experiments

The Landsat-8 instrument provided low-resolution remote measurements, whereas PRISM spectra were proxy for *in situ* measurements due to high quality and resolution (both spectral and spatial). In this experiment, we used a VAE for feature extraction and the DCDM for spectral unmixing (coral, algae, and sand) because of their superior results. The data set that was used for training was withheld from this experiment. Min-max normalization was applied to both Landsat-8 and PRISM spectra to allow the model to focus on spectral features rather than albedo values. The VAE learned how to extract $d = 3$ features from PRISM spectra. We found this value to work well in practice.

For Heron Island, we focused on a 6×3 km subregion that is diverse in terms of benthic cover; for Kaneohe Bay, we focused on a smaller subregion of size 1.5×1.5 km (Figure 6.27).

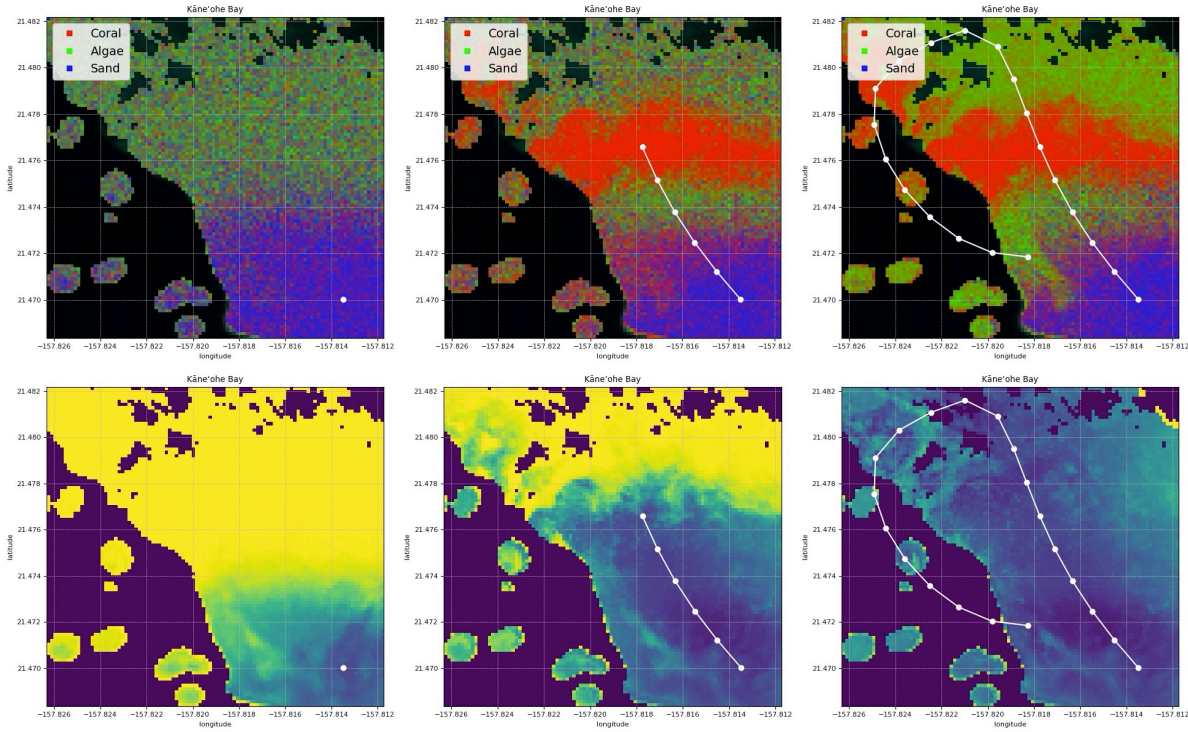


Figure 6.24: An example of a simulated traverse at Kaneohe Bay. Coral reef maps are shown on the top, while entropy maps appear on the bottom. Signatures are successfully identified and extrapolated throughout the scene. The map is refined as more samples are collected. A few samples (20) are sufficient to map the benthic cover of most of the 1.5 x 1.5 km area.

Forty-nine different starting locations were evenly spaced throughout each subregion; end goals were not specified. In total, 343 traverses were simulated at each site. Additionally, we imposed a constraint of 20 samples per traverse. An example of a simulated traverse at Kaneohe Bay together with its corresponding coral reef mapping process is shown in Figure 6.24.

Training and simulation were performed using a laptop computer with an Intel i7 processor (2.9 GHz quadcore) and 16GB of memory. Each waypoint was computed within just a few seconds or minutes, depending on the complexity of the path planner.

Four automated sampling approaches were compared in this study (Figure 6.25).

- **Random Sampling:** It is a baseline that does not quantify nor use information from the model whatsoever, hence poorer performance is to be expected. In order to generate trajectories that are somewhat smooth, random sampling occurred inside a fixed radius at each step. Two different random algorithms were used (Figure 6.26):
 - **Random Uniform (RU):** It samples from a uniform distribution inside the radius.
 - **Random Edge (RE):** It samples from a uniform distribution along the perimeter of the circular region.
- **Bayesian Experimental Design:** It consists of two greedy heuristics that achieve near-optimal




Figure 6.25: Four sampling strategies for coral reef mapping: random sampling, Bayesian experimental design, Monte Carlo tree search, and ergodic optimal control. Random sampling is the simplest approach since it ignores how useful future samples might be. Bayesian experimental design provides a probabilistic framework for identifying the most informative samples and planning paths accordingly. Monte Carlo tree search combines random sampling with a tree search that focuses on the most promising actions. Ergodic optimal control not only selects informative samples, but also generates smooth trajectories that can be suitable for boats or AUVs.

results in Gaussian processes [102]. Both rely on a local-window sampling approach, similarly to the previous random sampling methods.

- **Maximum entropy (ME):** It selects the point with the largest individual entropy. ME sampling is relatively fast to execute, but it tends to select outliers.
- **Maximum Information Gain (MIG):** An improved strategy that selects the point with the largest information gain. MIG sampling is more expensive since it requires computation of an expectation, but it tends to outperform ME sampling by selecting more representative samples.
- **MCTS:** A Monte Carlo tree-search planner that creates a long term strategy by planning ahead for possible observations in the future. The planner was implemented by Kodgule et al. [99] using a four-step lookahead. It performs a grid-search, as opposed to the previous sampling methods that rely on local windows.
- **Ergodic Optimal Control:** Ergodic theory is a framework for studying the statistical properties of dynamical systems and stochastic processes [138]. Ergodic optimal control leverages such concepts to derive sampling trajectories that are both smooth and informative. Regarding smoothness, our GP-based formulation is a continuous model with the inherent benefit of




Figure 6.26: Visualization of the random sampling strategies. Random uniform selects any point inside a local region defined by a fixed radius r . Random edge selects any point along the region's perimeter, potentially increasing traversed distance and coverage..


allowing samples anywhere. Regarding informativeness, the key idea behind ergodic optimal control is to compute trajectories such that the amount of time spent in a region is proportional to the expected information gain in that region. Based on the work by Edelson [42], we examined two algorithms specifically designed for information gathering:

- **Spectral Multi-scale Coverage (SMC):** Proposed by Mathew and Mezić [116], SMC provides distinct benefit of balancing exploration and exploitation. Additionally, it produces trajectories that are distinctly smooth. Formally, the objective function is designed to maximize the rate of decay of an ergodicity metric that measures the difference between the time-averaged behavior of the trajectory and the uncertainty of the model in terms of entropy. An informative trajectory will visit high-entropy regions frequently and low-entropy regions occasionally. The model is updated after collecting a sample and a new trajectory is computed at every step to produce a path.
- **Projection-based Trajectory Optimization (PTO):** Proposed by Miller and Murphy [121], PTO directly optimizes ergodicity over the entire trajectory. SMC, on the other hand, improves the rate of change of ergodicity over a single step. PTO tends to generate trajectories that collect even more informative samples than SMC, but at the expense of less smoothness, more total traversed distance, and little control on step size.

Furthermore, we compared these automated sampling methods against two additional approaches: an optimal bound in which every pixel in the scene is sampled (ultimately thousands of points), and a sampling strategy that consists of the coordinates from which actual samples were collected by scuba divers during the CORAL field campaign (Figure 6.27).

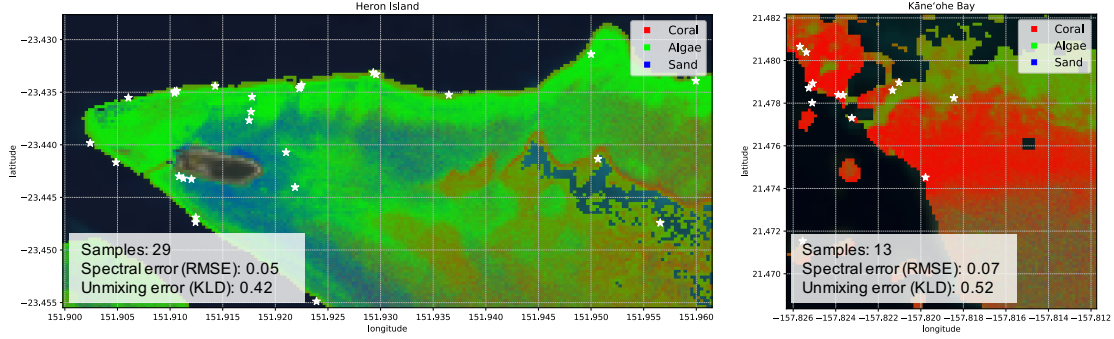
Three metrics were used to evaluate performance in terms of science return. For normalization purposes, we compute the averages with respect to the total number of points in the map.

PRISM-Derived Benthic Cover




(a) Ground truth abundance maps estimated by Thompson et al. [167] and validated during the CORAL mission.

CORAL Mission Sampling Strategy



(b) Predicted abundances using our model along with the sampling strategy followed during the CORAL mission.

Optimal Sampling Strategy



(c) Predicted abundances using paths generated by two optimal sampling strategies. In this example, the optimal sampling strategies produce more accurate maps both as a function of spectral reconstruction error (RMSE) and unmixing error (KLD).

Figure 6.27: Experimental setting for the simulation study. A subregion of 6 x 3 km in Heron Island is shown on the left, whereas a region of 1.5 x 1.5 km in Kaneohe Bay appears on the right. White asterisks indicate samples.

- **Entropy:** It is a measure of uncertainty in the model. It is *directly* minimized by the planners (Equation 6.11) and is calculated without a ground truth.
- **Error:** It is a measure of the spectral reconstruction error of the scene in terms of root mean squared error (RMSE). It should be *indirectly* minimized by the planners since it requires a ground truth for comparison.
- **Kullback-Leibler divergence:** It is a smooth function that measures the difference between the true and predicted class probabilities (Equation 2.8). It should be *indirectly* minimized by the planners.

Additionally, used three cost metrics: total computation time, total traversed distance, and roughness of the traverse. The latter two are directly related to energy and time consumption. The roughness cost quantifies sudden turns along the trajectory, where perfectly straight paths receive a score of 0 degrees. Roughness is based on the smoothness scores for robots proposed by Hidalgo-Paniagua et al. [74] and Guillén Ruiz et al. [150].

6.7.2 Results

We first calculated the correlation between metrics throughout the simulations. There are high Pearson correlation coefficients between reconstruction error (RMSE) and model uncertainty (entropy) (Figure 6.28); they are 0.881 and 0.948 for Heron Island and Kaneohe Bay, respectively. These values indicate a positive correlation between Landsat and PRISM data; they also confirm that entropy is a suitable objective function for spectral mapping. A similar result is observed when comparing KLD with entropy (Figure 6.29). The correlation coefficients for Heron Island and Kaneohe Bay are of 0.937 and 0.836, respectively. This demonstrates that entropy is also quite useful for building accurate abundance maps.

The science performance plots of the different sampling strategies are shown in Figure 6.30, whereas the associated costs appear in Table 6.1. Note that Table 6.1 does not include cost results for the scuba diving strategy. This is because optimal sampling strategies generate paths, whereas the scuba diving samples were collected throughout multiple days, sometimes by more than one person; hence comparisons may not be fair. In all cases, entropy, reconstruction error, and KLD show decreasing trends that approach the optimal bound as more samples are collected. Random sampling strategies perform worst overall since they lack the ability to identify the most informative samples. Random edge is better than random uniform, apparently because it covers more distance. However, long traverses are not enough to achieve good performance since random edge is outperformed by the rest of the sampling algorithms. The greedy strategies, maximum entropy and maximum information gain, do better than random sampling. As expected, information gain is a more useful reward than bare entropy. Non-greedy approaches perform best since they look farther. PTO has the best performance scores, but at the cost of significantly longer traverses that cannot be controlled nor bounded. MCTS has similar performance scores

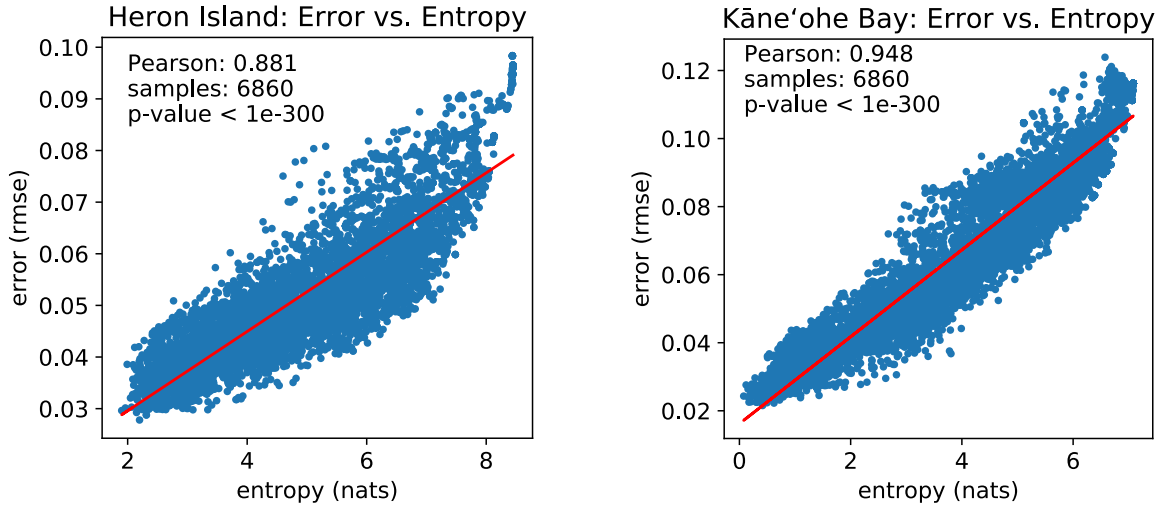


Figure 6.28: Spectral reconstruction error (RMSE) vs. model uncertainty (entropy). Results for Heron Island are shown on the left, while results for Kaneohe Bay appear on the right. Both variables are strongly correlated according to Pearson’s correlation coefficient.

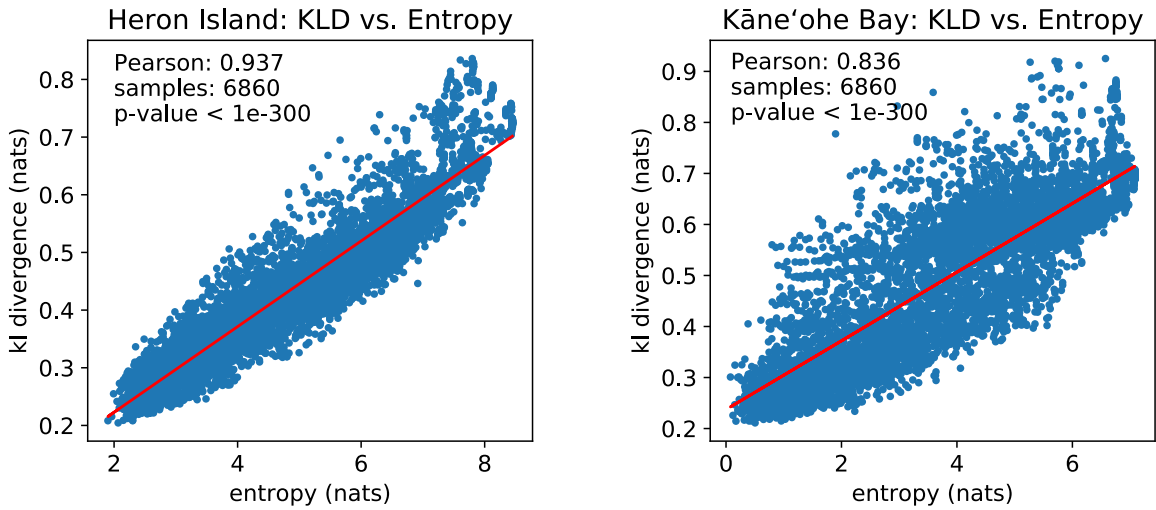


Figure 6.29: Unmixing error (KLD) vs. model uncertainty (entropy). Results for Heron Island are shown on the left, while results for Kaneohe Bay appear on the right. Both variables are strongly correlated according to Pearson’s correlation coefficient.

with much more reasonable traverses, but with expensive computation times. SMC seems like an intermediate alternative with the appeal of having the smallest roughness scores overall. Interestingly, the divers’ strategy tends to score somewhere in between the random methods and the intelligent algorithms in terms of mapping accuracy. This seems reasonable since the divers did not formulate sample selection in terms of entropy, let alone optimized it directly; nonetheless, they were certainly executing a more intelligent strategy than mere random sampling. Another

Table 6.1: Associated costs for each sampling algorithm in terms of computation time, traversed distance, and smoothness. Seven sampling strategies are evaluated: random uniform (RU), random edge (RE), maximum entropy (ME), maximum information gain (MIG), Monte Carlo tree search (MCTS), Spectral Multi-scale Coverage (SMC) and Projection-based Trajectory Optimization (PTO). Results show the average performance plus-minus one standard deviation for 49 simulated traverses per strategy per site.

Metric	Site	RU	RE	ME	MIG	MCTS	SMC	PTO
Time (seconds)	Heron Island	36.58 ± 1.24	33.35 ± 2.14	30.16 ± 3.61	40.38 ± 1.24	454.6 ± 35.4	38.45 ± 0.59	43.63 ± 0.34
	Kaneohe Bay	22.07 ± 0.71	21.45 ± 0.09	19.56 ± 1.31	25.29 ± 0.85	350.0 ± 29.1	23.38 ± 0.44	27.51 ± 0.05
Distance (km)	Heron Island	7.12 ± 0.65	10.62 ± 0.03	8.70 ± 0.39	9.68 ± 0.46	10.18 ± 0.33	9.25 ± 0.11	12.96 ± 1.46
	Kaneohe Bay	2.47 ± 0.19	3.69 ± 0.10	2.95 ± 0.47	3.31 ± 0.32	3.51 ± 0.08	3.2 ± 0.02	4.95 ± 0.46
Roughness (degrees)	Heron Island	56.0 ± 13.5	55.4 ± 12.3	25.4 ± 8.7	24.7 ± 9.1	15.9 ± 6.3	7.7 ± 1.8	21.9 ± 7.8
	Kaneohe Bay	55.3 ± 10.5	53.6 ± 10.5	35.7 ± 11.7	29.3 ± 12.1	31.3 ± 9.0	8.6 ± 1.7	21.7 ± 4.7

observation is that divers prioritized sampling corals during the field campaign, whereas most of our automated sampling strategies favor benthic diversity.

6.8 Discussion and Conclusion

This chapter presents an approach to wide-area mapping and optimal robotic sampling. Through combination of remote sensing and *in situ* data, our method achieves wide-scale coverage by extrapolating relevant spectral features and by adapting to new information. This is done through a probabilistic machine learning model that integrates spectral feature extraction, spatio-spectral Gaussian process regression, and learning-based composition analysis. Furthermore, our approach identifies and quantifies the most valuable samples by utilizing well-defined principles of decision theory, information theory, and Bayesian experimental design for sample selection. Finally, various sampling techniques are applied and compared. Specifically, random sampling, greedy heuristics, Monte Carlo tree search, and in some instances ergodic optimal control.

Our studies and experiments demonstrate that the combination of these machine learning methods, together with remote sensing, enable a robotic explorer to leverage a few *in situ* samples by extrapolating their features to many locations in large areas. Consistent results from the Cuprite, Mars, and Coral experiments suggest a good generalization of our approach. Regarding decision theory for informative sample selection, our results indicate that probabilistic modeling leads to substantial benefits. Concretely, we observe that entropy is strongly correlated to spectral reconstruction error, as well as to composition mapping error; indicating that entropy is

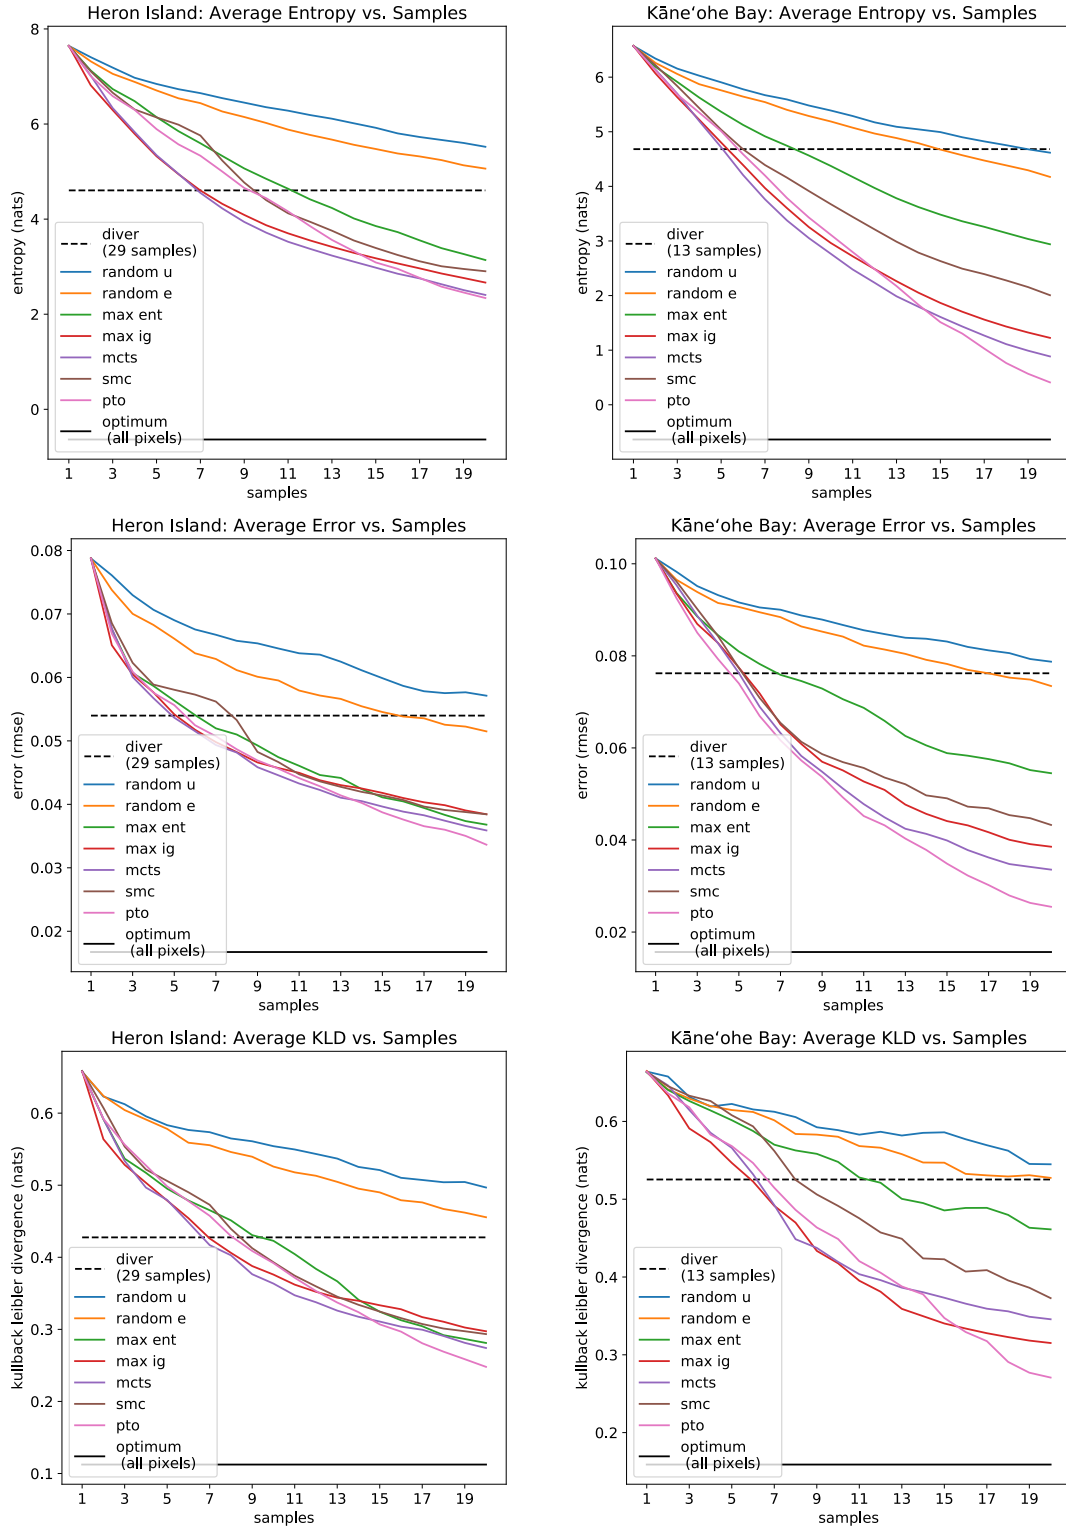


Figure 6.30: Average model uncertainty (entropy, top row), spectral reconstruction error (RMSE, middle row), and unmixing accuracy (KLD, bottom row) as a function of samples.

a suitable objective function for efficient wide-area mapping. We evaluate various state-of-the-art sampling strategies and conclude that they can accommodate diverse needs and constraints. Greedy approaches are fast and achieve decent results. If computation time is not an issue, we recommend MCTS since it is non-myopic. In the Coral scenario experiments, PTO has the best mapping accuracy, but at the expense of very long traverses that are hard to bound. SMC appears to be a well-balanced method: it has a competitive performance with fast computation times, its traversal distances are straightforward to control, and it generates the smoothest trajectories overall, potentially resulting in significant energy and time savings.

The results from the field experiment with the rover Zoë demonstrate the viability of our approach in a real exploration scenario using a robotic platform. Important lessons were learned, especially those related to processes that tend to be oversimplified in simulations, such as *in situ* data collection and safe terrain navigation.

Chapter 7

The Science Hypothesis Map for Co-Exploration

7.1 Introduction

Modern planetary robotic exploration is guided by scientists specifying the waypoints on a path that they believe will best address the investigation questions. The path is formed from expert knowledge of the site and expectations about where to gather mission-critical information. 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 [79]. 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 [44, 164]. However, many exploration scenarios occur under low bandwidth and high latency communications, leaving limited opportunities to revise the exploration plan.

This chapter presents 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 scientist initially 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. It introduces the *science hypothesis map*, a 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. Next, it describes information gain-based path planning strategies that can be integrated with the science hypothesis map to optimize the rate and productivity of discovery. Finally, it demonstrates the application of these methods to a simulated geologic mapping problem of a well-studied site: Cuprite Hills, Nevada.

Table 7.1: Examples of simplified Science Traceability Matrices.

Investigation objectives	Physical properties	Measurements
Map geologic formation processes in a region	Abundance of key mineral classes (quartz, kaolinite, calcite, etc.)	Surface reflectance, 350-2500 nm, at 10 nm spectral resolution and >300 SNR
Map water quality in a river	Salinity of the water	Water conductivity in $\mu\text{S}/\text{cm}$
Map air quality in a city	Abundance of particles such as carbon monoxide, sulfur dioxide, etc.	Particle concentration in $\mu\text{g}/\text{m}^3$

7.2 Related Work

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, abstract knowledge in a direct way that enables true robotic understanding. However, it is often 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*. Some examples of Science Traceability Matrices for geologic, oceanographic [11], and pollution mapping are shown in Table 7.1.

A few deployed robots can perform automatic science data analysis [11, 28, 44, 180], 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 [11]), or detecting transient features (dust devils on Mars [24]), are defined long in advance. Some robots have used Bayesian networks for tasks such as mineral classification [54], or meteorite identification [137], 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 Bayesian experimental for robotic exploration, and consequently build probabilistic graphical models connecting high-level concepts with low-level observations [3, 59]. 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.

7.3 Probabilistic Mapping Model

7.3.1 Science Hypothesis Map

The *science hypothesis map* extends the idea of a Science Traceability Matrix to a probabilistic model with spatial extent, granting robustness under uncertainty. It is also a framework for calculating the information that collected measurements, directly interpretable by the robot, provide with respect to the fundamental investigation objective. Specifically, the science hypothesis map has the following hierarchical components (Figure 7.1):

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

For mapping simplification purposes, the explored environment is partitioned into disjoint regions, each associated with its own independent conditional distributions relating investigation objectives, physical properties, and measurements (Figure 7.1).




Figure 7.1: 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 $z \in Z$.

7.3.2 Bayesian Inference

The next step is to use the previous probabilistic model to infer the corresponding investigation objective H for each region $r \in R$, given a measurement Z . In other words, the goal is to 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:

$$P_R(H, Y, Z) = P_R(H)P_R(Y|H)P_R(Z|Y). \quad (7.1)$$

Integrating over variables:

$$P_R(H, Z) = P_R(H) \sum_Y P_R(Y|H)P_R(Z|Y), \quad (7.2)$$

$$P_R(Z) = \sum_H P_R(H) \sum_Y P_R(Y|H)P_R(Z|Y). \quad (7.3)$$

From the definition of conditional probability:

$$P_R(H|Z) = \frac{P_R(H, Z)}{P_R(Z)}. \quad (7.4)$$

Finally, we get the desired closed-form solution:

$$P_R(H|Z) \propto P_R(H) \sum_Y P_R(Y|H)P_R(Z|Y). \quad (7.5)$$

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.

7.3.3 Geologic Model

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. This work considers a geologic exploration model in which (Figure 7.2):

- The investigation objectives H are *geologic classes* defining different ages and formation processes. An example would be a hydrothermal formation.
- The physical properties Y are *minerals*, distinctive chemical compositions that diagnose the geologic formation conditions of rocks. For example, hydrothermal formations are characterized by the presence of minerals such as chlorites and serpentines.
- The measurements Z , *reflectance spectra*, representing the fraction of incident light reflected




Figure 7.2: Visual example of the relation between geologic classes, minerals and reflectance spectrum measurements. Elaborated with information from the United States Geologic Survey (USGS) [31].

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 [32].

7.4 Informative Exploration

This work focuses on the specific 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 in a path $\mathcal{P} = \{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_n\}$, where spatial coordinates are denoted as $\mathbf{x} \in X \subset \mathbb{R}^2$. The corresponding possible combinations of the n spectrum measurements is denoted as Z^n .

The robot aims to optimize *information gain*, which is equivalent to *mutual information*. It is a classical objective function for information-driven action selection [9, 113]. Specifically, it is defined as the expected reduction in uncertainty after collecting new information, measured with Shannon entropy. The Shannon entropy of an independent region r is:

$$I_R(H) = - \sum_H P_R(H) \log P_R(H). \quad (7.6)$$

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

$$I_R(H|Z^n) = - \sum_{Z^n} P_R(Z^n) \sum_H P_R(H|Z^n) \log P_R(H|Z^n). \quad (7.7)$$

From Equations 7.6 and 7.7, the information gain for each region can be calculated as:

$$IG_R(H|Z^n) = I_R(H) - I_R(H|Z^n). \quad (7.8)$$

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:

$$IG(R|Z^n) = \sum_R IG_R(H|Z^n). \quad (7.9)$$

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:

$$|Z^n| = \binom{m+n-1}{n}. \quad (7.10)$$

It can be noticed that m , the total number of possible different minerals, is constant. Therefore, the combinatorial explosion when calculating the expectancy over alternatives (Equations 7.7 and 7.8) is due to the number of measurements n .

While information gain can be estimated with Markov chain Monte Carlo approximations [152], here the distributions $P(Z)$ and $P(H|Z)$ are not known *a priori*. An alternative is to perform Monte Carlo integration [8], with the possibility of using simple sampling strategies. But another problem arises: the approximation accuracy deteriorates as n increases. However, given the problem formulation in this work, 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:

$$\lim_{n \rightarrow \infty} IG_R(H|Z^n) = I_R(H). \quad (7.11)$$

as shown by Haussler and Oppel [71]. Since information gain (Equations 7.7 and 7.8) is calculated by marginalizing Y (equation 7.5), these properties hold true.

This means that it is plausible to fit a curve to the first few data points (i.e., a small number of measurements n), which are simpler to calculate or 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 [146]. Specifically, the used GLF is parametrized as follows:

$$GLF(n; a, b, c) = a (1 - e^{-bn})^{-c}, \quad (7.12)$$

where a, b, c are the hyperparameters that need to be estimated. Therefore, at least three points are required for fitting.

An example is shown for the information gain calculation for a Bayesian network with 30




Figure 7.3: Information gain calculation using three different methods: direct, Monte Carlo integration, and GLF fitting.

investigation objectives and 20 physical properties (Figure 7.3). 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 6 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 [156]. 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.

Finally, this work 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:

$$\begin{aligned} \max_{Z^n} \quad & IG(R|Z^n) \\ \text{s.t.} \quad & Cost(Z^n) \leq Budget \end{aligned} \tag{7.13}$$

Information gain is submodular. In this context, it means that sampling the same region repeatedly leads to diminishing expected returns. From the objective function in equation 7.9 and the submodularity of information gain, there is a resulting exploitation vs. exploration trade-off that translates to taking many samples from a few highly-rewarding regions, or to visiting as many different regions as possible. The science hypothesis map integrated with this formulation in equation 7.13 can accommodate many different path planning strategies. There exist straightforward methods that use mixed integer linear programming (MILP) to generate optimal informative paths, but only when the rewards between sampled points are independent [183]. There is a family of algorithms that use the submodular property of information gain to pro-

duce near-optimal solutions [101, 154]. There are also some Monte Carlo tree search (MCTS) methods for informative path planning [3]. Other approaches may be computationally intensive, but potentially closer to optimality, such as using branch and bound (B&B) techniques for both discrete [10] and continuous [83] space representations.

7.5 Cuprite Scenario

The science hypothesis map for geologic exploration was first validated in a simulation study based on the Cuprite scenario. This section first discusses the model that was constructed for Cuprite; it then describes the experiments that were conducted as well as the respective results.

7.5.1 Geologic Model

We next explain the procedure that was followed to train the hierarchical probabilistic model for the Cuprite scenario, i.e., learn the association between geologic classes and minerals, $P(Y|H)$, and the relationship between minerals and reflectance spectrum measurements, $P(Z|Y)$. The different elements of the science hypothesis map are shown in Figure 7.4. Geologic classes H and minerals Y are based on expert-drawn geologic and mineral maps by Swayze et al. [161], 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 class 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. These proportions are used to train a conditional probability table of minerals given geologic classes, i.e., $P(Y|H)$. This table consists of 32 geologic classes and 20 minerals. The training set uses up to 2000 random samples from every predefined region. These datasets are withheld from the analysis that follows.

Representative reflectance spectra of some key minerals in Cuprite are also shown in Figure 7.4. *In situ* 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) [68]. It assigns a unique reflectance spectrum measurement to every location (pixel) in the scene. AVIRIS-NG mapped the area 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 [166].

A Gaussian Mixture Model (GMM) is used as a simple but effective probabilistic classifier that predicts the corresponding mineral from a spectrum measurement. 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




Figure 7.4: Cuprite, NV. Left: geologic classes H . Center: minerals Y and their abundances. Right: Representative reflectance spectra Z of some key minerals. Each pixel is associated with a full spectrum. Geologic classes and mineral abundances are based on the work by Swayze et al. [161].

estimator described in [82]. 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 Z as a discrete variable simplifies information gain calculations. Consequently, the classification output of the GMM classifier is treated as a categorical answer. Figure 7.5 shows the corresponding 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 the estimated mineral Z to differ from the true physical variable Y . The used training and validation sets each use up to 1000 samples from every mineral type. These datasets are withheld from the analysis that follows.

Finally, information gain was computed using GLF-based approximations. Specifically, we fitted a GLF to each disjoint spatial region as a function of its prior distribution (initial hypothesis). Similarly to the example described in Section 7.4, the GLF fitting process used Monte Carlo approximations for the first 5 measurements n .

7.5.2 Experiments and Results

Bayes Learning

This group of experiments evaluated the model's ability to predict and recover the true geologic class of a region (with a specific age, formation process, etc.) from high-resolution AVIRIS spectral measurements with the Bayesian update. For each of the 32 regions, 200 random sampling




Figure 7.5: Empirical conditional probability table relating actual to estimated minerals for each of 20 possible mineral types in the Cuprite scene.

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 class of each region, this experiment tracked two fundamental variables: the posterior probability of the correct geologic class, and the updated entropy (Figure 7.6). As more and more measurements are collected, the updated probability of the true underlying geologic class 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 three representative scenarios, explained next.

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 class, while the remaining probability mass is distributed to alternatives. This prior allows ample margin for converging to either a correct or incorrect answer. Most beliefs converge to the right answer with few samples, while entropies reduce significantly. The three exceptions with poor outcomes correspond to the smallest regions trained with the fewest data points.

Starting with complete uncertainty: In this case, a situation of complete uncertainty was represented with a uniform prior over geologic class labels. Despite the more challenging situation, most beliefs still show significant improvement. As expected, there is more variance and an overall slower convergence rate.

Starting with an inaccurate prior: Finally, this part represents a situation where the scientist’s initial beliefs are actually incorrect. For that, a prior probability of 50% was assigned to an incorrect unit: the most similar geologic class according to the Hellinger distance for probability distributions – a challenging error to correct. The remaining probability mass is distributed uniformly. There is an even slower improvement with a higher variance. But in most cases the

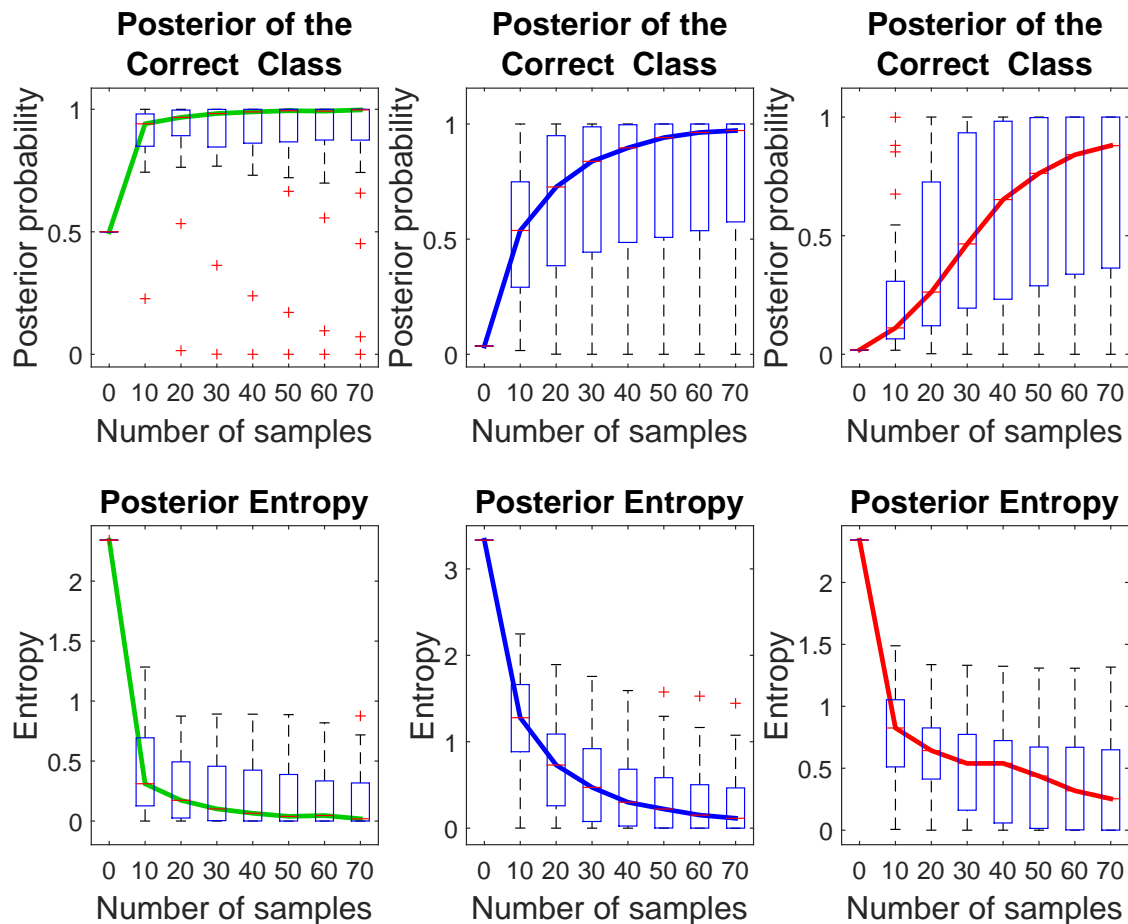


Figure 7.6: Posterior probability of the correct class (top row) and posterior entropy (bottom row), starting with an accurate prior (left column, green medians), under complete uncertainty (center column, blue medians), and an inaccurate prior (right column, red medians). Each data point in the box and whisker plots corresponds to one of the 32 different regions.

model still recovers.

This results indicate that the model is able to recover the true geologic class most of the times as spectral measurements are collected, even under misleading initial conditions.

Path Planning

This group of experiments integrated 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.

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 class 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.

As a control case, this work evaluates three science-blind planners that ignore information gain.

- **Direct (D)**: A path planner that selects the sequence of waypoints that minimizes the total path length with Dijkstra’s algorithm [39].
- **Random(R)**: A path planner that sequentially selects a set of random neighboring waypoints that do not violate the budget on path length.
- **Max-N (M)**: A path planner that maximizes the number of visited nodes using a classic recursive greedy approach for orienteering [27]. 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 work compares three science-aware planners that incorporate measurements’ information gain.

- **Greedy (G)**: A path planner that adds at each step the neighboring waypoint that maximizes the objective function, this without exceeding the budget.
- **Branch and bound (B)**: A path planner as a nonmyopic alternative. It selects a sequence of waypoints using the algorithm described by Binney and Sukhatme [10]. In this particular case, implementing a three-step calculation look ahead.
- **pSPIEL (P)**: A planner that combines submodular orienteering algorithms by Singh et al. [154] with this work’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 (Figure 7.7).

We computed the average scores per traverse of three simple but useful metrics: path length, number of collected measurements, and number of explored regions (Table 7.2). 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




Figure 7.7: Example paths from the 6 planners. The start is on the upper left, and the end on the lower right. Paths can differ notably from one another depending on the planning algorithm.

Table 7.2: Average scores for three relevant path planning metrics: path length, collected measurements, and explored regions.

Site	Metric	D	R	M	G	B	P
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

planners, apparently favoring exploration over exploitation in this scenario.

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

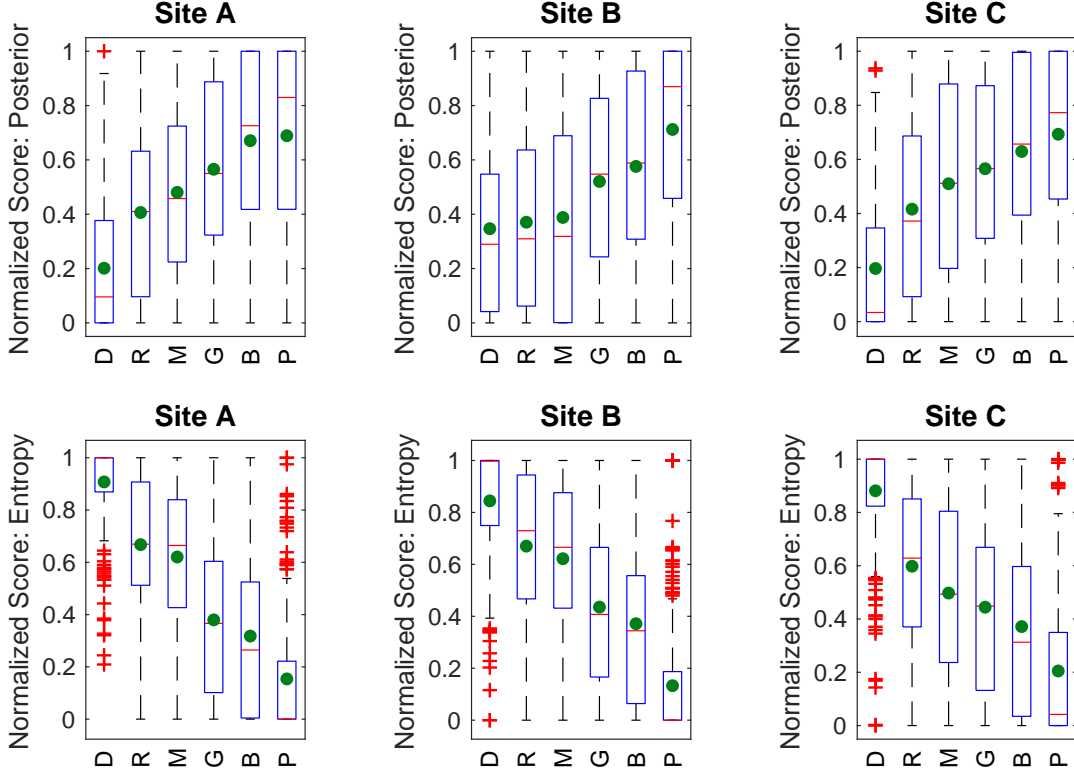


Figure 7.8: Comparison of planners based on their ability to recover correct geologic classes (top row) and the final entropy (bottom row) per traverse using a normalized score. The green dots are the means.

each given pair of start-end locations. The resulting scaled scores for prediction accuracy and entropy are shown in Figure 7.8. On the other hand, Table 7.3 shows a set of paired t-tests comparing the performance among planners for both metrics with their raw unnormalized scores.

These results demonstrate a couple of things. 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.

Table 7.3: Paired t-tests between planners’ performance using a significance level of 5%. In most cases, there is a statistically significant difference (Y/**N**) in both inference (first entries) and entropy (last entries).

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	-	N , Y	N , Y	Y, Y	Y, Y	Y, Y	-	Y, Y	Y, Y	Y, Y	Y, Y	Y, Y
R	-	-	Y, Y	Y, Y	Y, Y	Y, Y	-	-	N , N	Y, Y	Y, Y	Y, Y	-	-	N , Y	Y, Y	Y, Y	Y, Y
M	-	-	-	Y, Y	Y, Y	Y, Y	-	-	-	Y, Y	Y, Y	Y, Y	-	-	-	N , N	Y, Y	Y, Y
G	-	-	-	-	Y, Y	Y, Y	-	-	-	-	N , Y	Y, Y	-	-	-	-	Y, Y	Y, Y
B	-	-	-	-	-	N , Y	-	-	-	-	-	Y, Y	-	-	-	-	-	Y, Y
P	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-

7.6 TREX Scenario

The science hypothesis map for geologic exploration was further validated using the TREX scenario. More details about this scenario can be found in Chapter 4. Both simulation and field experiments were conducted. This section first discusses the model that was constructed for the TREX studies. Afterwards, it describes the results from the simulation and field experiments.

7.6.1 Geologic Model

The TREX science team specified 10 geologic origins (investigation objectives) and 502 minerals (physical properties) for the model. The geologic origins are: lacustrine/marine, evaporitic/playa, metamorphic/hydrothermal, hydrothermal, pedogenic/diagenic/weathering, igneous, biogenic, ambiguous, ice, and manmade. The mineral list was derived from the USGS spectral library [31] and it included minerals that can be detected in the 1 μm and 2 μm spectral bands. The science team elaborated a simple table relating minerals to geologic origins. The table consisted of binary values (zeros and ones). In order to construct the probability distribution $P(Y|H)$, we normalized the table and used additive (Laplace) smoothing to avoid numerical issues such as zero division errors.

The Tetracorder algorithm [33] was used to estimate mineralogy from reflectance spectra. After analyzing a spectrum, Tetracorder returns three results for each mineral in the USGS spectral library: goodness of fit, spectral feature depth, and the product fit \times depth. All of these variables have a value between zero and one, where a higher value indicates a stronger mineral match. For our model we used fit \times depth and treated it as the probability distribution $P(Z|M)$. Additionally, Tetracorder returns these three values for each spectral band or “group”. Our analysis focused on the spectral groups 1 and 2 that correspond to the 1 μm and 2 μm spectral bands, respectively.

A subset of the science team did not have access to the whole synthetic geologic environment, just to the corresponding AVIRIS-NG image. Their objective was to perform a preliminary analysis based on contextual remote sensing data and provide an initial hypothesis regarding the geologic origins of the site. This prior, or initial hypothesis, is shown in Figure 7.9.

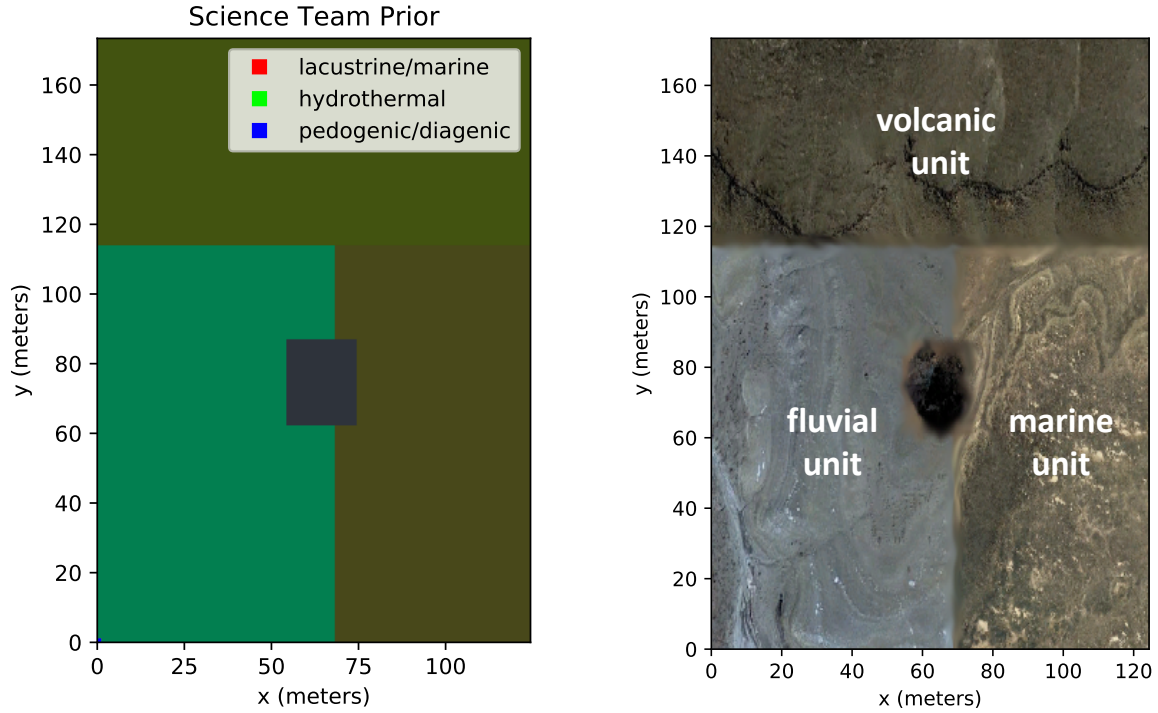


Figure 7.9: Initial hypothesis regarding geologic origins in the synthetic environment (left). This figure shows the three ground truth geologic origins: lacustrine/marine for the marine unit, hydrothermal for the fluvial unit, and pedogenic for the volcanic unit (right). The levels of confidence in the prior are somewhat low.

A simple neural network (multilayer perceptron) was trained in order to approximate information gain in a much faster manner. The reason for this choice is simple: the TREX science model suffers from a combinatorial explosion because of its 502 different minerals. For instance, the computation of information gain for a number of samples $n = 5$ requires the evaluation of more than 270 billion different outcomes. Furthermore, preliminary Monte Carlo approximations performed poorly. The used neural network learned how to predict the three parameters of the GLF described in Equation 7.12. These parameters were used to define a GLF, and then the GLF was used to estimate information gain. The neural network was trained with 10,000 different priors that were sampled uniformly from a 10-D simplex because the science model consists of 10 geologic origins. Examples of GLF and information gain predictions for different priors are shown in Figure 7.10.

7.6.2 Experiments and Results

The experiments consisted of simulations and a complementary field experiment where Zoë was deployed for validation purposes.




Figure 7.10: Information gain approximation with a neural network for the TREX scenario. A non-informative is shown on the left while an informative prior appears on the right. On average, samples produce more information gain when starting with a non-informative prior.

Simulation Experiments

The exploration space was represented as an 8-connected grid with a spacing of 10 m between nodes. A slope constraint of 20° and an elevation threshold of 200 m were enforced by not including the corresponding nodes in the graph. Since it is a small and somewhat constrained environment, 50 pairs of start and end locations were randomly selected so that they were all roughly at the same distance (100-120 m). In addition, we defined a distance budget of 2 times the shortest path length to allow for exploration flexibility,

Four informative path planning algorithms were compared in this study (Figure 7.11). They were all modified so that they could handle distance budgets and end goals.

- **Direct:** A path that minimizes traverse cost using Dijkstra’s algorithm [39].
- **Random:** It sequentially samples a random neighboring location until the budget is exhausted. It serves as a baseline and corresponds to an uninformed sampling strategy. This is a *science-blind* baseline.
- **Greedy:** It makes the locally optimal choice using a one-step lookahead. This is a *myopic* exploration strategy.
- **MCTS:** A Monte Carlo tree-search planner that creates a long term strategy by planning ahead for possible observations in the future. The planner adapted the MCTS implementation by Kodgule et al. [99] using a four-step lookahead. This is a *non-myopic* path planner.

Three metrics were used to evaluate the performance of the planners in terms of science return:

- **Number of Samples:** a count of samples collected during a traverse.
- **Entropy:** a measure of uncertainty in the model. It is *directly* minimized by the planners (Equation 7.6) and is calculated without a ground truth. For normalization purposes, we compute the average with respect to the total number of points in the map.
- **Kullback-Leibler divergence (KLD):** a smooth function that measures the difference be-

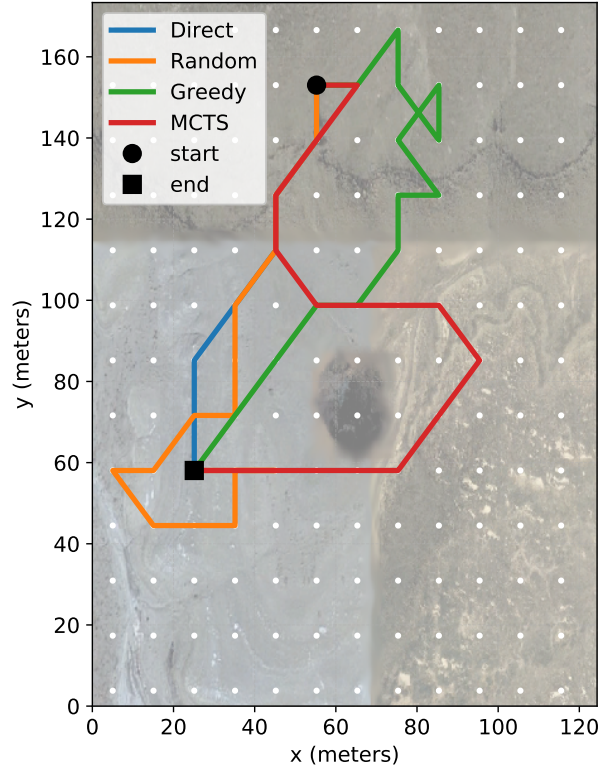


Figure 7.11: Example paths from the 4 planners in the synthetic geologic environment. Paths can differ notably from one another depending on the planning algorithm.

tween the true and predicted class probabilities (Equation 2.8). It should be *indirectly* minimized by the planners.

Additionally, the total traversed distance per traverse was calculated (in meters) as a cost metric.

The results from the simulation experiments appear in Figure 7.12. This figure shows box plots of the aforementioned performance metrics: number of samples, entropy, KLD, and distance. It also includes paired t-tests between planners. Results indicate that smart planners tend to perform better than simpler planners, but only by a modest amount in some cases. This can be attributed to the fact that the synthetic geologic environment is very simple as it only consists of 3 units. Consequently, many planners can learn most of the map very quickly and their performances converge accordingly. Furthermore, Random and Greedy do not have a statistically significant difference in terms of collected number of samples, entropy, nor KLD. Myopic planning can only decide between neighbors where most of them have the same information gain as they are usually contained inside the same unit, essentially becoming equivalent to a random strategy. In contrast, MCTS does a better job since it has a longer planning horizon that makes it more likely to consider other units. In terms of traversed distance, it is evident that the Direct planner always follows the shortest paths as permitted by the traversability map. It is interesting to observe that Greedy and MCTS basically traverse the same distance. This underscores the

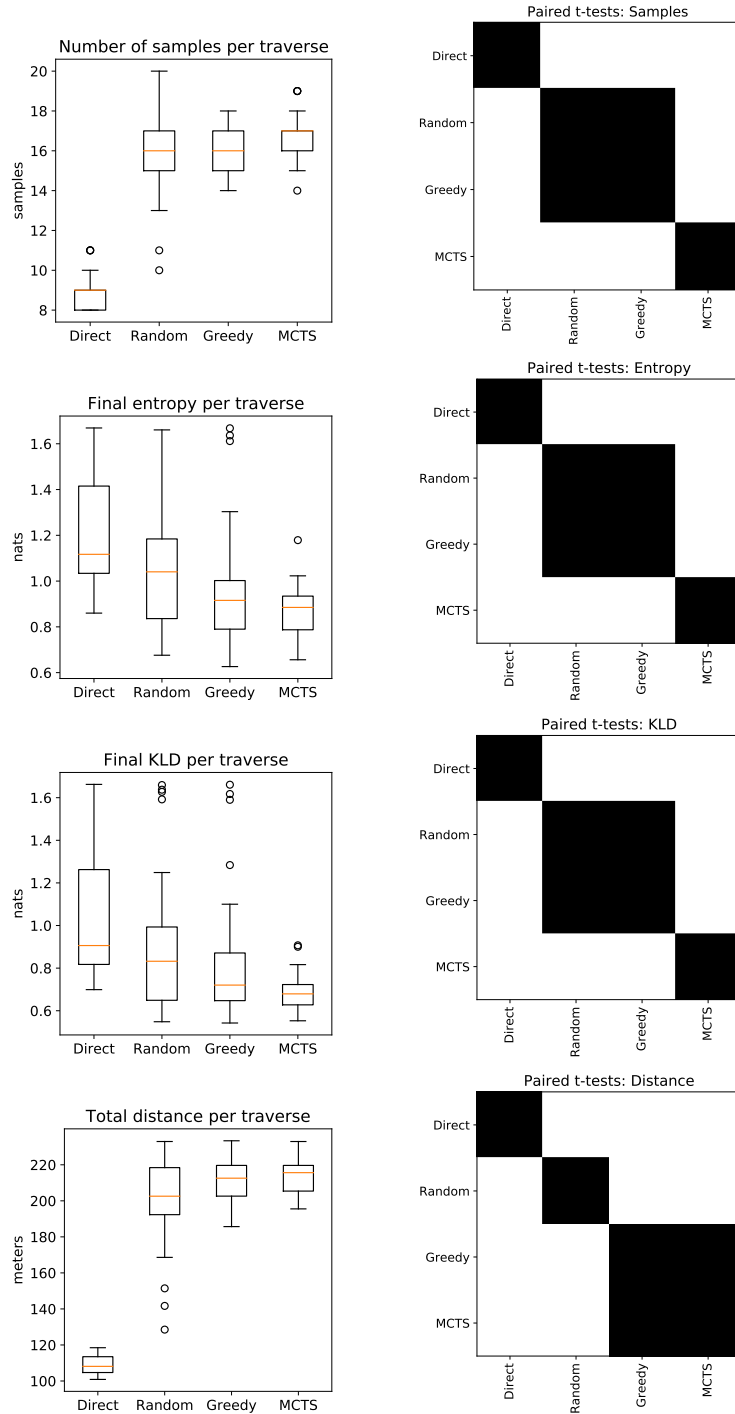


Figure 7.12: Path planners' results from the TREX simulation experiment. Left column: box plots of number of samples, entropy, Kullback-Leibler divergence (KLD), and distance per traverse. Right column: corresponding paired t-tests between planners, white indicates a statistically significant difference ($\alpha = 0.05$). Smarter planners tend to outperform simple strategies, although these differences are modest because of the relatively small exploration environment consisting of only 3 units.




Figure 7.13: Autonomous rover traverse at the Gascola site. Left: RGB drone imagery. Right: synthetic geologic environment prepared for the site.

importance of efficiency as planners that use similar resources, in this case distance, may not necessarily yield the same results in terms of science productivity.

Field Experiments

Zoë, the rover, was deployed between 26-30 July 2021 at the Gascola site near Pittsburgh, PA. More details about the TREX scenario can be found in Chapter 4. The field experiment consisted of an operations readiness test for an upcoming field campaign in Arizona. Therefore, the Gascola test had a more limited scope. The objective was to emulate data collection with the rover under more realistic conditions as opposed to the AVIRIS-NG simulations. To this end, the TREX science team prepared about 20 mineral samples that contained different fine-grained materials. These samples were placed at specific geographic coordinates according to their associated locations in the synthetic geologic environment. Whenever the rover reached a particular waypoint, ASD spectra from the corresponding sample was measured with manual assistance.

Most of the time was devoted to robot teleoperation exercises for the science team. Such efforts focused on exploring the marine and volcanic units of the synthetic geologic environment. The remaining unit (fluvial) was explored using autonomous rover science and navigation. The respective rover path is shown in Figure 7.13. Seven samples were collected during the traverse, which was generated using the MCTS planner running on board the rover. Results from

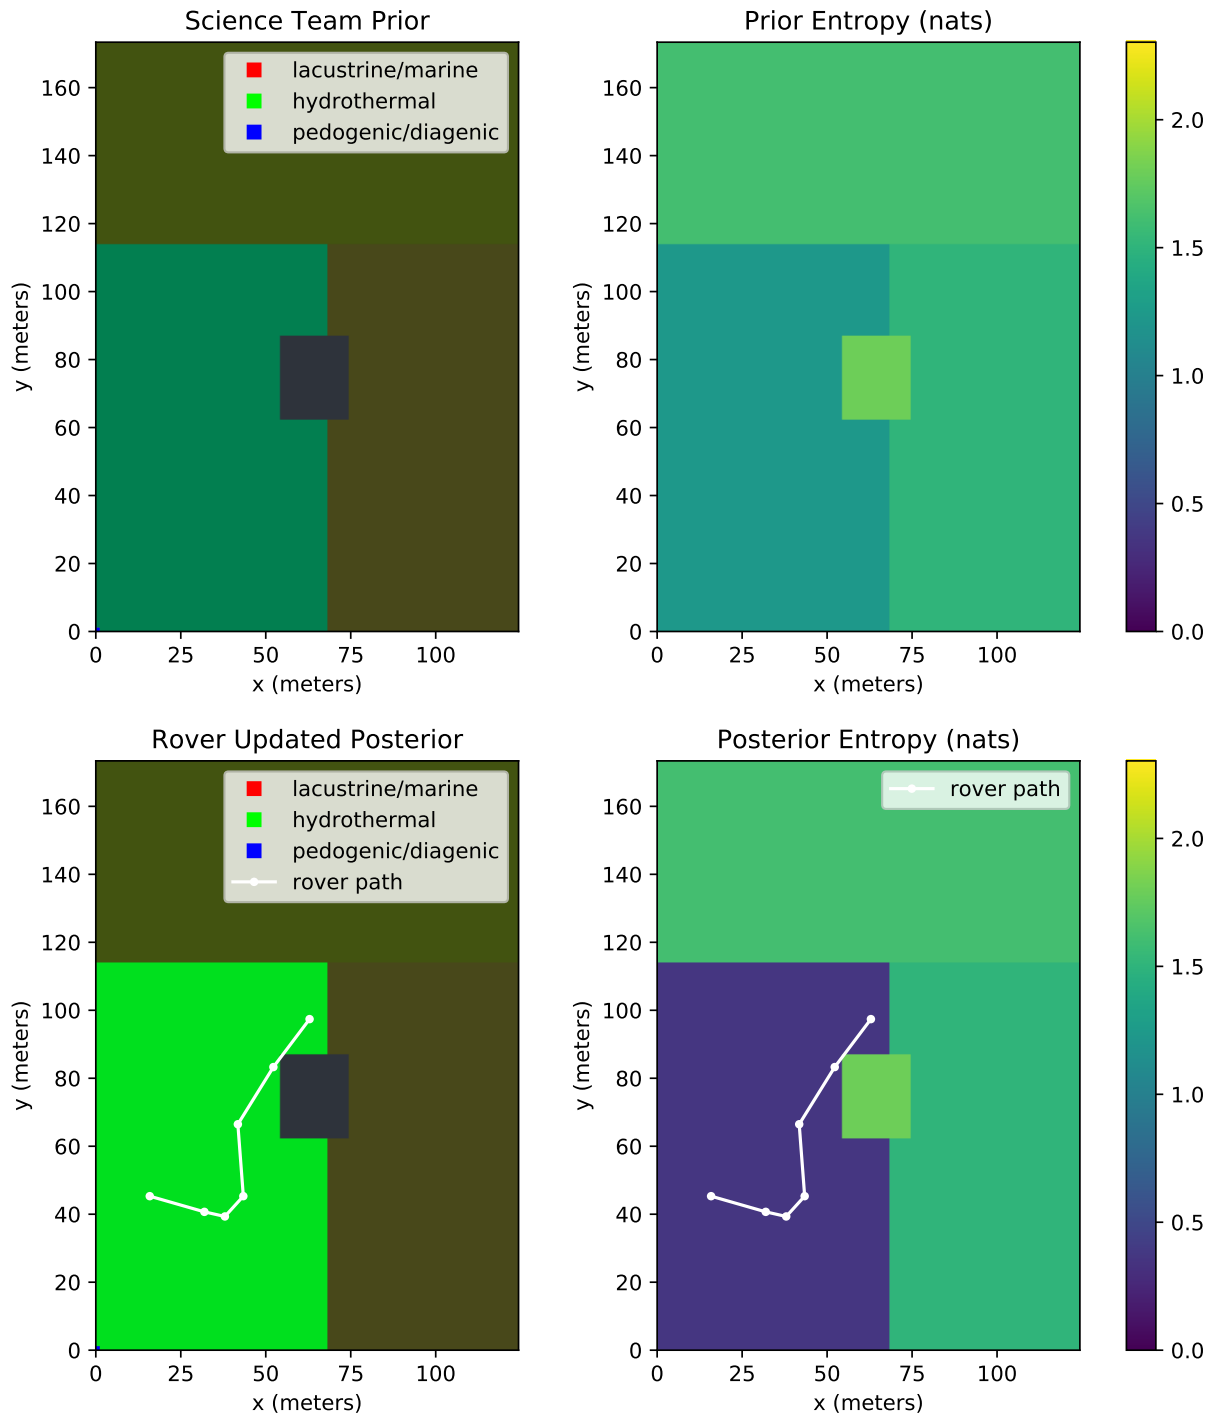


Figure 7.14: Results from the autonomous rover traverse at the Gascola site. Tetracorder detected minerals supporting the initial hypothesis that the fluvial unit may have a hydrothermal origin, which originally had a 50% confidence. The model updated accordingly and the corresponding confidence increased to 88%. At the same time, the models' entropy decreased.

this single path are limited but quite promising (Figure 7.14). First of all, this is the first time that Tetracorder [33], the state-of-the art in automated spectral composition analysis, has ever been successfully deployed on any rover. It is important to mention that the rover was not exempt from challenges that are common during field spectra collection such as calibration issues, varying light conditions due to sudden cloud coverage, and sample targeting difficulties (even though this process was assisted manually). However, Tetracorder was able to produce plausible and consistent results many times when analyzing *in situ* spectra in real time. During the autonomous traverse, Tetracorder detected minerals such as feldspars, hematites, and kaolinites. These specific minerals support the initial hypothesis that the fluvial unit may have a hydrothermal origin, which originally had a 50% confidence. The model updated accordingly and the corresponding confidence increased to 88%. At the same time, the models' entropy decreased. Overall, these results are encouraging since they support the feasibility of our approach under more realistic conditions.

7.7 Discussion and Conclusion

This chapter develops a Bayesian framework in which scientists initially describe their abstract beliefs and hypotheses, and then the state of this belief evolves as the robot makes raw measurements. This work constructs a spatial and hierarchical probabilistic structure that relates high-level science representations to measurable data. We call this model the science hypothesis map. Moreover, scientific information gain is derived from the model and efficiently computed. Finally, robotic exploration is formulated as an informative path planning problem, allowing for the maximization of scientific information gain.

We highlight several key findings of this simulation study based on the Cuprite scenario. First, the experiments demonstrate that the science hypothesis map is a sound mathematical framework for describing exploration for specific well-defined objectives. The probabilistic model can infer geologic class with a high accuracy under diverse challenging situations. In the case of robotic path planning, path planners that exploit information gain better reduce uncertainty over the investigation objectives. Not surprisingly, non-myopic planning outperforms myopic and science-blind alternatives.

Results from the TREX scenario further validate our approach. In simulations, we see similar trends in a different environment for a more general geologic model that was provided by an actual team of planetary scientists. Results from the ORT field experiment were limited in scope but are quite promising since they demonstrate the feasibility of our approach under real-life conditions with a rover. Overall, it is encouraging to observe how ideas from the science hypothesis map are starting to be adopted by researchers in the planetary science community. Additionally, this is the first time that Tetracorder [33], the state-of-the art in automated spectral composition analysis, has ever been successfully deployed on any rover.

A more surprising finding is the strength of the statistical link between measured spectra, the corresponding mineral types, and the geologic classes. This is not a foregone conclusion because geologists form this classification from many other features such as local morphology, elevation, the wide-area geographic three-dimensional structure of different strata, and domain knowledge. It is striking that reflectance spectra, drawn independently and randomly from each unit, so strongly predict geologic classes.

Chapter 8

Science and Risk-Aware Exploration

8.1 Introduction

Planetary rovers have science goals to complete, such as reaching specific locations, sampling rocks, and sending data back to Earth. Nevertheless, safe exploration is of the utmost importance. A relevant example involves the Mars Exploration Rovers (MER) Spirit and Opportunity. Although both rovers landed on Mars in early 2004, Spirit became stuck in a sand dune in 2009 and ceased communications in 2010, while careful operation allowed Opportunity to exceed its original mission plan until 2018.

Most of the work related to rover motion planning focuses on either increasing science productivity or reducing traversability risk; but not both simultaneously. In practice, these issues are sometimes addressed in conjunction by just oversimplifying one of them. This particular work proposes and describes a comprehensive technique for the holistic integration of both science and risk. Specifically, a Bayesian framework that is able to simultaneously quantify science value and risk, provide probabilistic guarantees, and design paths accordingly. In other words, a method that addresses how to reduce risk while still performing relevant science investigations.

This chapter is structured as follows. First, it discusses state-of-the art planning methods for science-aware planning. Then, it describes current approaches for risk-aware planning for rovers, mainly focusing on geometric and slip-based models and planners. Afterward, it introduces our planning framework that combines science productivity with probabilistic risk constraints. Finally, it includes results from a simulation study involving rover sample collection at Jezero Crater, Mars. These results show that risk is not only bounded but also reduced by a significant amount while slightly affecting science productivity.

8.2 Related Work

8.2.1 Science-Aware Planning

Herein we pay special attention to science-aware planners that utilize well-defined principles from information theory and statistical learning for decision-making. Bayesian experimental design [25] is a framework from which an optimal experimental design (in this case a series of measurements in a rover path) may be derived by maximizing the expected value of an information-theoretic utility function. Examples of these utility functions include Shannon entropy, mutual information, and Kullback-Leibler divergence [35]. A more detailed explanation of these utility functions can be found in Chapter 2.

There exist several planning algorithms for information gathering that use Bayesian experimental design at varying degrees, each with its own advantages and drawbacks. There are methods that assume independence between sampling locations, which is usually an oversimplification in informative exploration scenarios [27, 183]. Krause et al. [102] and Singh et al. [154] use submodular optimization techniques based on near-optimal greedy algorithms and demonstrate that they tend to work well on Gaussian processes. Thompson et al. [166] rely on a greedy graph-based approach for sampling the most representative spectra in a scene. Other approaches may be computationally intensive, but potentially closer to optimality, such as branch and bound (B&B) techniques; existing graph-based [10] and sampling-based [83] versions. There are Monte Carlo tree search (MCTS) planners that have been applied to geologic exploration scenarios since they can use long planning horizons, but results are not deterministic and may vary considerably if too few iterations are used [3, 99]. Gautam et al. [58] propose a multi-heuristic A* planner that employs an entropy reward together with a simple distance cost; they explore this trade-off using Pareto optimization via genetic algorithms. Miller et al. [121, 122] combine ergodicity theory and optimal control to derive smooth trajectories (paths) that spend a proportional amount of time in areas according to their potential information utility.

All of the previous algorithms use science-related information for information gathering, but they either ignore or use very simple cost functions and constraints. Additionally, they inherently fail to evaluate and address risk in navigation.

8.2.2 Risk-Aware Planning

The navigation systems equipped on current Mars rover missions such as MER and Mars Science Laboratory (MSL) detect and avoid local geometric hazards using the grid-based estimation of surface traversability applied to local terrain (GESTALT) algorithm [60]. For path planning, they rely on the D* algorithm [160]. Although these methods have enabled autonomous rover operations, avoiding specific geometry is not enough to guarantee safety. Both MER and MSL missions have experienced hazardous conditions due to sand dunes, which are often detected

as geometrically benign. These hazards can create adverse conditions such as wheel slip and sinkage. Excessive slip in terrain with little traction can cause rovers to become entrapped.

More recent navigation systems [52] account for slip since it is one major concerning issue during rover operations [133, 134]. Simulation studies have been conducted to estimate the teramechanical properties of Mars rovers [93, 185], but in practice it is often necessary to perform validation through empirical studies [5, 6]. Wheel slippage is usually difficult to estimate, especially for terrains with loose material. Slip for planetary rovers has been modeled as a function of two different terrain properties: slope and type (e.g. sand dunes, bedrock, etc.) [5, 89, 123]. In other words, slip has been modeled using geometric and semantic information. Especially pertinent is the Bayesian model by Cunningham et al. [37] as they train Gaussian processes to predict slip as a function of both terrain type and slope using data from the Curiosity rover.

Risk-aware planning algorithms for planetary rovers typically rely on digital elevation models (DEMs) and imagery (orbital and *in situ*) for estimating slope and terrain types, respectively [133, 134]. Inotsume et al. [90] present a Bayesian framework that is based on rapidly-exploring random trees (RRT*) [94], allowing users to define slip-based risk thresholds and generate paths accordingly. Mizuno et al. [123] also use an RRT* planner, but they model and propagate slip uncertainty with a particle filter. Ono et al. [132] first segment and classify terrain from both orbital and *in situ* imagery using the deep learning algorithm SPOC [149], and then use this labeled data together with DEMs to identify obstacles and traversability costs; paths are computed afterwards by applying a rapidly-exploring random graph (RRG) in conjunction with A*. Hedrick et al. [73] utilize a similar terrain-aware planner; however, it updates its terrain-type map given new *in situ* information by using belief propagation in Markov Random Fields, with possible replanning if needed.

All of the previous planners are inherently science-blind since they do not to evaluate scientific information and importance, let alone exploit it. They have a lot of unfulfilled potential as there is plenty of scientific value in terrain understanding, which they only use to address traversability. It has been shown that morphology is strongly correlated to mineral composition [17, 168, 173]. Mineralogy contains valuable information about water, habitability, and biosignature potential [12, 173].

8.3 Approach

This section describes a framework for planetary rovers that consists of two Bayesian models. One model addresses science information, while the other model focuses on traversability risk. They are both constructed using orbital data of Jezero Crater, the Martian landing site for the Perseverance rover. More information regarding the Mars scenario and data sets can be found in Chapter 4.

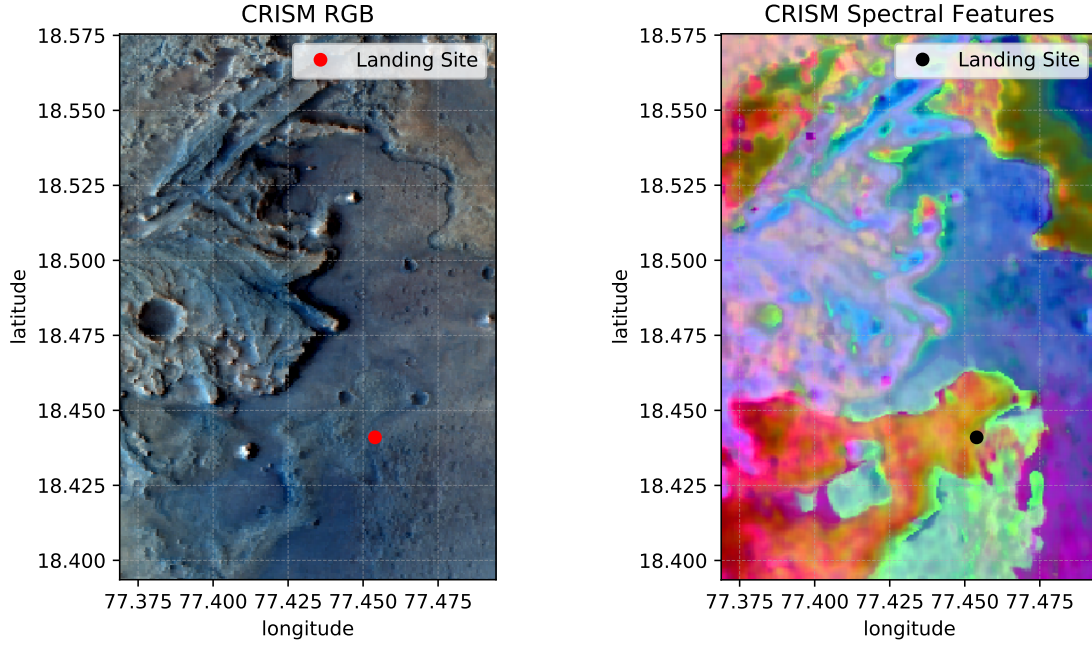


Figure 8.1: Jezero Crater orbital data used for the science model.

8.3.1 Science Model

This work emulates autonomous sample collection efforts performed by a rover at Jezero Crater. Throughout this research we have discussed various models for science. Herein we focus on the model described in Chapter 6, which combines remote and *in situ* data for efficient mapping and exploration. Similarly to current Mars investigations, we assume that there is access to high-quality orbital data. CRISM data is employed for science analysis (Figure 8.1), which has a resolution of 36 m/pixel. It serves as the source of remote data, and is also used to generate synthetic measurements from the SuperCam suite of instruments. Specifically, the infrared spectrometer sensitive to 256 channels in the 1300-2600 nm spectral range. A Variational Autoencoder (VAE) with three latent dimensions is used for feature extraction (more information can be found in Chapter 5). Hence, the science model consists of three independent Gaussian processes. For simplicity, the science model does not perform explicit composition analysis, but rather performs spectroscopic mapping.

8.3.2 Risk Model

The risk model for rover traversability is the most important element of the work described in this chapter. This model integrates geometric and semantic information in order to construct a more comprehensive representation of risk. Specifically, slope and terrain class data are used to estimate slip probability. Risk is then defined as the probability of exceeding a user-defined slip threshold. Each of these elements will be described in more detail.

Geometric and Semantic Information

Geometric information is derived from a digital elevation model (DEM) of Jezero Crater (Figure 8.2a). The DEM was generated using stereo pair images from the HiRISE orbital instrument and has a resolution of 1 m/pixel. Slope θ_0 can be easily calculated using an elevation map (Figure 8.2b). The slope at a point (x, y, z) is computed using the magnitude of the elevation gradient, which is a function of the derivatives $\frac{dz}{dx}$ and $\frac{dz}{dy}$. Specifically, we use the Sobel image filter when computing elevation derivatives, similarly to One et al. [134], as it is an isotropic gradient operator. Then, slope θ_0 is given by:

$$\theta_0(x, y) = \tan^{-1} \left(\sqrt{\left(\frac{dz}{dx}(x, y) \right)^2 + \left(\frac{dz}{dy}(x, y) \right)^2} \right). \quad (8.1)$$

Geometric data is complemented with semantic information. To this end, we rely on a geologic map of Jezero Crater especially elaborated for the Mars 2020 mission [158]. We focus on two different terrain classes: bedrock units, which consist of consolidated materials; and surficial units, which are constituted by unconsolidated materials such as sand and gravel. Bedrock units tend to be benign and traversable, whereas surficial units are more likely to create adverse conditions for a rover. The original geologic map defines 14 types of bedrock units and 7 types of surficial units. Since bedrock units may be partially or totally covered by loose materials, the geologic map has multiple regions where these two main classes overlap. For simplification purposes, our work treats all 14 bedrock units as one class. Additionally, we assume that surficial units with minor and moderate coverage can also be considered as bedrock units. This simplified geologic map is shown in Figure 8.2c.

Slip Prediction

Slip s , also known as slip ratio, is the relationship between the commanded velocity $\mathbf{v}_{ref} = (v_{ref,lon}, v_{ref,lat})$ and the actual velocity $\mathbf{v} = (v_{lon}, v_{lat})$. This work considers both longitudinal slip (in the direction of travel of the robot) and lateral slip (perpendicular to the direction of travel). Longitudinal slip s_{lon} is defined as follows:

$$s_{lon} = \begin{cases} 1 - v_{lon}/v_{ref,lon} & \text{if } v_{lon} \leq v_{ref,lon} \\ v_{ref,lon}/v_{lon} - 1 & \text{if } v_{lon} > v_{ref,lon} \end{cases} \quad (8.2)$$

Lateral slip s_{lat} is given by:

$$s_{lat} = \frac{v_{lat}}{v_{ref,lat}}. \quad (8.3)$$

A positive slip represents a state where the rover moves slower than intended. Absence of slip means that the rover is moving exactly at the desired velocity. A negative slip indicates that the

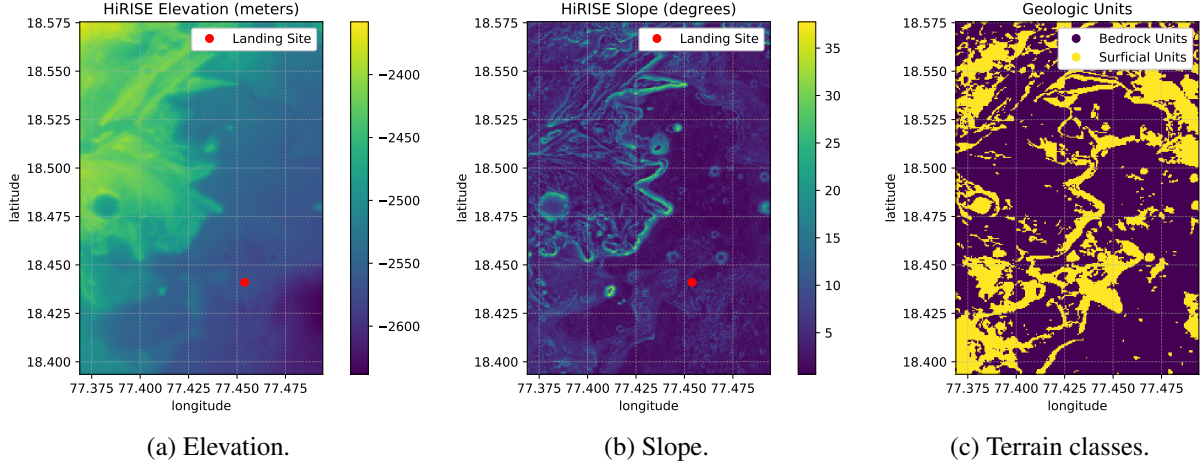


Figure 8.2: Jezero Crater data that was used to build the risk model. Geometric and semantic information are combined. Geometric information (elevation and slope) is derived from HiRISE data, while semantic information is based on the geologic map elaborated by Stack et al. [158] for the Mars 2020 mission.

rover is moving faster than the commanded velocity.

Cunningham et al. [37] use Gaussian processes (GPs) to predict slip for the Curiosity rover. Herein we operate under the simplifying assumption that the same slip prediction model can be used for other Mars rovers such as Perseverance. Cunningham et al. combine geometric and semantic information. Specifically, they use GPs to predict longitudinal and lateral slip separately, that is, they assume independence. Longitudinal slip is predicted as a function of pitch angles θ_p , while lateral slip is estimated with roll angles θ_r . Moreover, Cunningham et al. employ different GPs for each terrain class $c \in C$. That is:

$$s_{lon,c}(\theta_r) \sim \mathcal{GP}(\mu_{lon,c}(\theta_p), \sigma_{lon,c}^2(\theta_r)), \quad (8.4)$$

$$s_{lat,c}(\theta_p) \sim \mathcal{GP}(\mu_{lat,c}(\theta_r), \sigma_{lat,c}^2(\theta_p)); \quad (8.5)$$

where μ is the predicted mean and σ is the predicted variance. The learned models for two classes, bedrock and sand, are shown in Figure 8.3. In this work we assume that all surficial units have the same slip properties as sand.

Pitch θ_p and roll θ_r are calculated as a function of slope θ_0 and the rover's angle of attack α :

$$\theta_p(\theta_0, \alpha) = \sin^{-1}(\sin(\theta_0) \sin(\alpha)), \quad (8.6)$$

$$\theta_r(\theta_0, \alpha) = \sin^{-1}(\sin(\theta_0) \cos(\alpha)). \quad (8.7)$$

While slope is always nonnegative, the angle of attack defines the direction of travel (heading on slope) and whether pitch and roll angles are either positive or negative.

It is worth noting that this work focuses on global planning, not local navigation. Further-

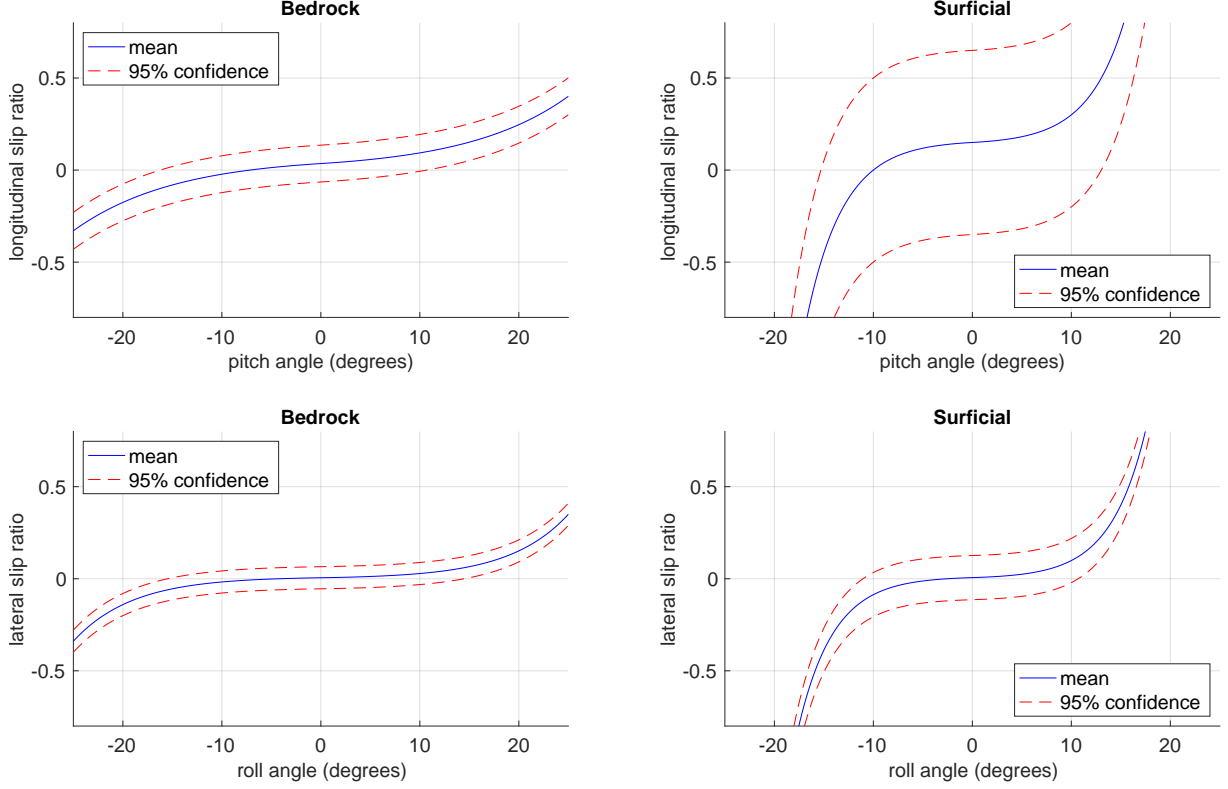


Figure 8.3: Slip estimation based on the method by Cunningham et al. [37]. Four different Gaussian processes are used to predict slip as a function of slope and terrain class. Bedrock units tend to have smaller slips with smaller variances. Longitudinal slip is often more severe than lateral slip.

more, there is a substantial difference in resolution between CRISM (36 m/pixel) and HiRISE data (1 m/pixel). Hence, a rover will most likely have different headings (and pitch and roll angles) on different slopes inside any given 36×36 m area (Figure 8.4). To account for this, we fit a categorical distribution to all possible slip scenarios (longitudinal and lateral) within each CRISM pixel (i, j) :

$$\hat{s}_{lon}(i, j) \sim \text{Cat}(\hat{\mathbf{p}}_{lon}(i, j)), \quad (8.8)$$

$$\hat{s}_{lat}(i, j) \sim \text{Cat}(\hat{\mathbf{p}}_{lat}(i, j)); \quad (8.9)$$

where the parameters $\hat{\mathbf{p}}_{lon}(i, j)$ and $\hat{\mathbf{p}}_{lat}(i, j)$ are probability vectors derived from a histogram. Such histogram is computed as follows. Possible slip values $\hat{s}_{lon}(i, j)$ and $\hat{s}_{lat}(i, j)$ are discretized between -1 and 1 using a sampling interval of 0.05, every HiRISE pixel inside a given CRISM pixel is included in the calculation (in total $36 \times 36 = 1296$), and possible angles of attack α are discretized between 0 and 360 degrees using a sampling interval of 10 degrees.




Figure 8.4: A CRISM pixel is equivalent to 36×36 HiRISE pixels. In terms of global planning, many different slope, pitch, and roll angles are possible.

Probabilistic Risk Model

This work defines risk for rover traverses in a similar fashion as Inotsume et al. [90]. We start by defining *local safety* at a pixel (i, j) as the probability of both lateral and longitudinal slips not exceeding a threshold s_{th} :

$$safety(i, j) = p(|\hat{s}_{lon}(i, j)| < s_{th}) p(|\hat{s}_{lat}(i, j)| < s_{th}). \quad (8.10)$$

Note that longitudinal and lateral slip are assumed to be independent. Furthermore, we use the absolute value of slip to avoid situations where the rover's actual velocity is either too slow or too fast with respect to the commanded velocity. We then define *local risk* as the opposite of local safety, that is:

$$risk(i, j) = 1 - safety(i, j). \quad (8.11)$$

Examples of how risk changes as a function of s_{th} are shown in Figure 8.5.

We then define the cumulative safety of a path $\mathcal{P} = (i^1, j^1), \dots, (i^k, j^k)$. For simplicity, we assume independence between points. Hence, it suffices to multiply local safeties as follows:

$$safety(\mathcal{P}) = \prod_{(i,j) \in \mathcal{P}} safety(i, j). \quad (8.12)$$

Again, we have that cumulative risk is the opposite of cumulative safety:

$$risk(\mathcal{P}) = 1 - safety(\mathcal{P}) \quad (8.13)$$

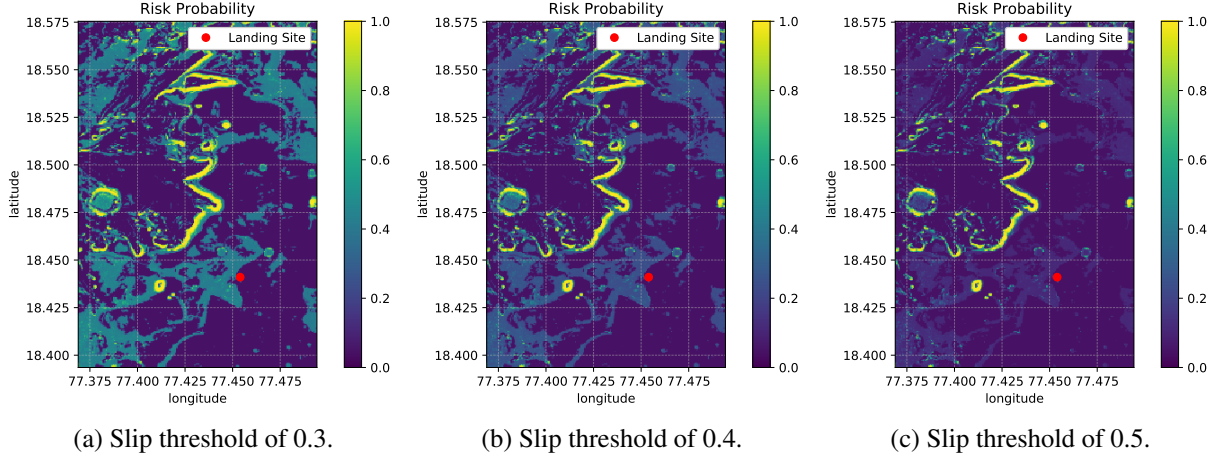


Figure 8.5: Jezero Crater risk model as a function of slip threshold.

In summary, this formulation allows for the quantification of risk for both specific locations and rover paths.

8.3.3 Path Planning Formulation

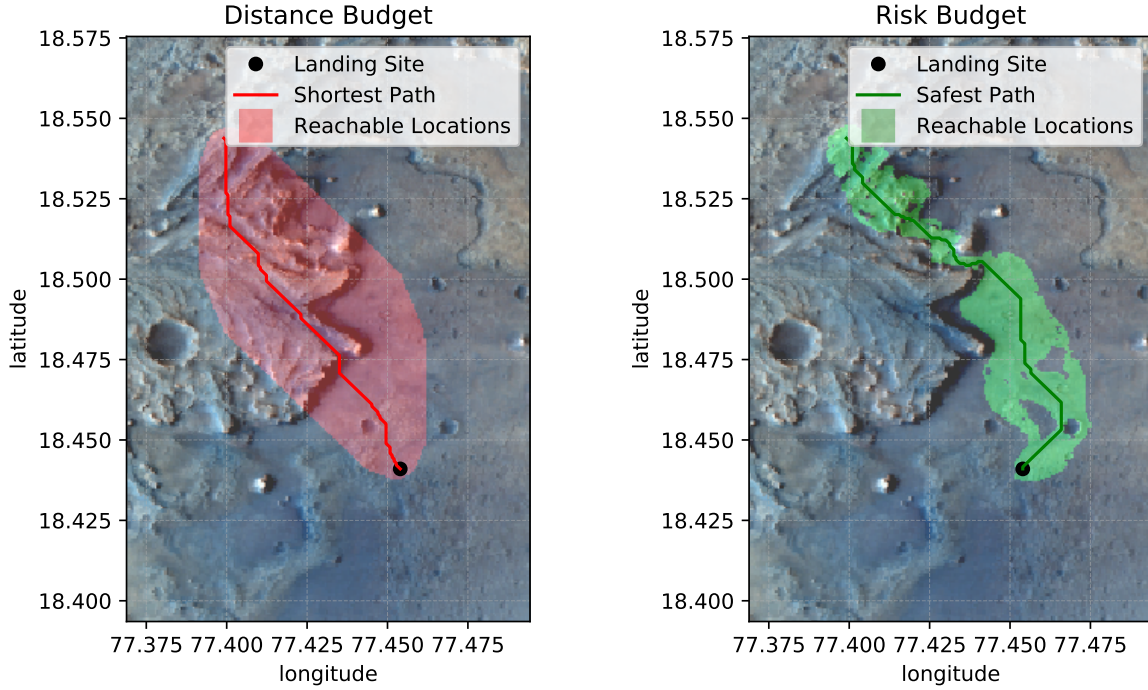
This work formulates science and risk-aware exploration as an optimization problem where the robot has to minimize model uncertainty (entropy) while respecting local and cumulative risk thresholds. That is:

$$\begin{aligned}
 \min_{\mathcal{P}} \quad & \hat{H}(\mathcal{M}|\mathcal{P}) \\
 \text{s.t.} \quad & risk(i, j) \leq \delta_{th}, \quad \forall (i, j) \in \mathcal{P} \\
 & risk(\mathcal{P}) \leq B
 \end{aligned} \tag{8.14}$$

where $\hat{H}(\mathcal{M}|\mathcal{P})$ is the posterior entropy of the science model (Chapter 6), δ_{th} is the local risk threshold, and B is the total risk budget for a rover traverse \mathcal{P} . Moreover, start and end locations are defined in advance as in many path planning problems.

Dijkstra’s algorithm [39] is employed to ensure that the rover will respect constraints and reach end goals. To this end, this problem is represented as a graph-search task. The exploration space is modeled as a graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, where \mathcal{V} are the vertices and \mathcal{E} are the edges. A trick is used so risk can be treated as an additive distance function d_{risk} . This is achieved by working in the negative log-space of safety:

$$\begin{aligned}
 d_{risk}(\mathcal{P}) &= -\log(safety(\mathcal{P})) \\
 &= -\sum_{(x,y) \in \mathcal{P}} \log(safety(x, y)).
 \end{aligned} \tag{8.15}$$



(a) Path planning using a Euclidean distance function. Here the distance budget is 1.10 times the shortest path length. Dangerous locations such as cliffs and sand dunes are reachable.

(b) Path planning using our risk framework. Here the risk budget is 1.10 times the safest path length. Paths are constrained to safer locations while providing margin for exploration.

Figure 8.6: Example of our path planning formulation.

The exploration graph is constructed with edge weights that are given by the average risk between two adjacent vertices. Additionally, vertices that violate the local risk constraint are not included in the graph. This modification allows Dijkstra’s algorithm to compute minimal risk paths instead of conventional shortest distance paths. Safest path lengths between all vertices and the end goal should be precomputed when employing a traversability budget. This way we make sure that the robotic explorer reaches the goal from any location without running out of budget. An example regarding shortest paths and traversability budgets is shown in Figure 8.6. When using a Euclidean distance function (Figure 8.6a), the robot may visit dangerous locations such as sand dunes and cliffs unless explicit local constraints are defined and enforced. In contrast, our risk framework automatically generates safer paths by avoiding such hazardous areas, but it also provides certain flexibility for assuming small risks so as to accomplish science objectives (Figure 8.6b).

8.4 Experiments and Results

8.4.1 Experiments

Two different experiments were conducted in this study. The first experiment was a sensitivity analysis with the purpose of finding an adequate slip threshold for our risk model. This was done by observing the behavior of safest paths according to each threshold. Once this threshold value was found, the second experiment compared Euclidean distance and risk-based budgets during path planning. This experiment was designed in order to analyze the benefits and drawbacks of risk-aware planning.

A subregion at Jezero Crater of size 10.8×7.2 km was used for both experiments. The exploration space was represented as an 8-connected grid using a pixel step size of 2, where each CRISM pixel has a resolution of 36 meters. A local slope constraint of 20° was enforced by not including the corresponding vertices in the graph. Many traverses were simulated in each experiment. They all started from the actual location where the Perseverance rover landed at Jezero Crater. For the first experiment, one hundred end locations were randomly selected so that they were all roughly at the same distance ($\pm 10\%$) from the start location. For the second experiment, forty end locations were randomly selected in a similar fashion, but also accounting for minimal-risk path length. Traversability budgets were defined as follows. The safest path between the start and end locations was generated; then its corresponding distance and risk path lengths were computed. The distance and risk budget were 1.25 times the shortest and safest path lengths, respectively. These values were chosen in order to have $d_{risk} < 2$, which represents an average risk probability of less than 0.1 at each visited vertex.

Four informative path planning algorithms were compared in this study (Figure 8.7). They were all modified so that they could handle budgets and end goals.

- **Direct:** A path that minimizes traverse cost using Dijkstra’s algorithm [39].
- **Random:** It sequentially samples a random neighboring location until the budget is exhausted. It serves as a baseline and corresponds to an uninformed sampling strategy. This is a *science-blind* baseline.
- **Greedy:** It makes the locally optimal choice using a one-step lookahead. This is a *myopic* exploration strategy.
- **MCTS:** A Monte Carlo tree-search planner that creates a long term strategy by planning ahead for possible observations in the future. The planner adapted the MCTS implementation by Kodgule et al. [99] using a four-step lookahead. This is a *non-myopic* path planner.

Three metrics were used to evaluate the performance of the planners in terms of science return:

- **Number of Samples:** How many samples were collected during a traverse.
- **Entropy:** It is a measure of uncertainty in the model. It is *directly* minimized by the planners

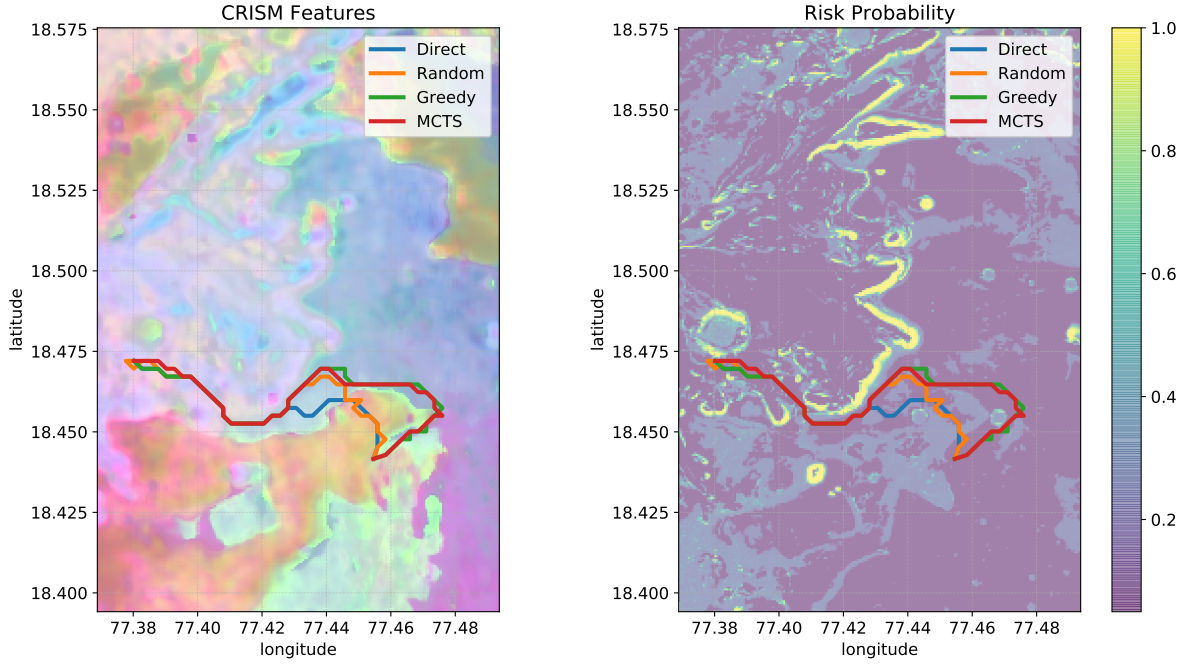


Figure 8.7: Example paths from the 4 planners at Jezero Crater.

(Equation 6.11) and is calculated without a ground truth. For normalization purposes, we compute the averages with respect to the total number of points in the map.

- **Error:** It is a measure of the spectral reconstruction error of the scene in terms of root mean squared error (RMSE). It should be *indirectly* minimized by the planners since it requires a ground truth for comparison. For normalization purposes, we compute the averages with respect to the total number of points in the map.

Additionally, three cost variables were computed:

- **Distance:** Total traversed distance, which is calculated in meters.
- **Risk-distance:** Our custom function for risk that is additive d_{risk} , which is given in Equation 8.15.
- **Risk probability:** The cumulative probability of risk on a path, which is calculated using Equation 8.13.

8.4.2 Results

The first experiment was designed as a sensitivity analysis in order to find a suitable slip threshold s_{th} for our risk model. Safe paths change when using risk models with different thresholds (Figure 8.8). In this case small slip thresholds tend to produce longer paths; and vice versa, large slip thresholds often result in more direct paths. Distance, risk, entropy, and error also change as a function of s_{th} (Figure 8.9). Results confirm that indeed traversed distances tend to increase

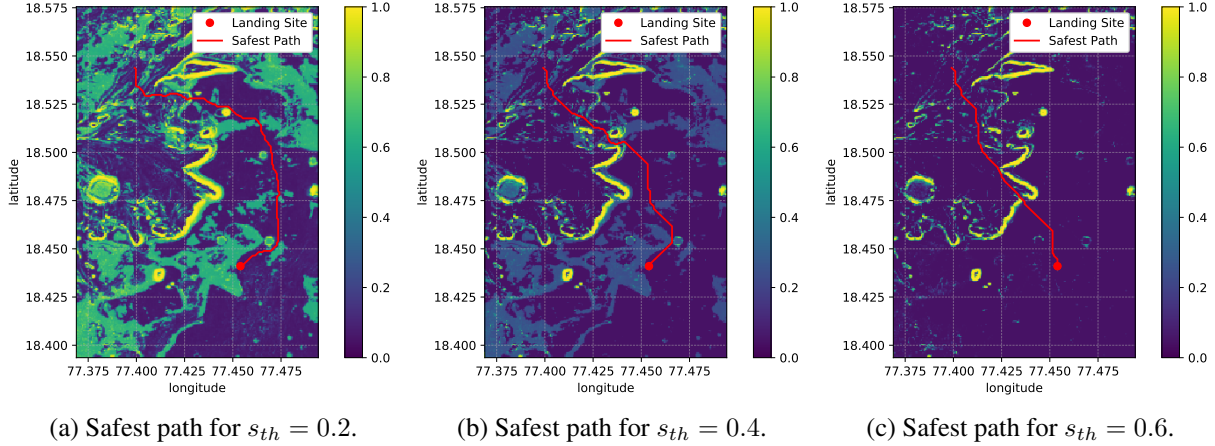


Figure 8.8: Results from the sensitivity analysis where safest paths were generated for different risk models. More conservative slip thresholds lead to longer paths.

when using larger slip thresholds. Risk as measured by d_{risk} is also reduced as a function of s_{th} ; the reason is that the perception of risk increases as the model is more averse to risk. Regarding science productivity, there is no clear trend. This is probably the consequence of a trade-off between less available sampling locations and more collected samples due to longer traverses. Overall, we observe that thresholds between 0.3 and 0.4 seem to offer a good balance between risk, total traversed distance, and science return. Hence, we selected a slip threshold of 0.35 for the risk model that was used for the second experiment.

The second experiment was conducted in order to evaluate the performance of the different path planners. The four planners used a risk budget of 1.25 times the minimal-risk path length in terms of d_{risk} (Figure 8.7). All planners avoid dangerous areas such as cliffs. They are also able to reach the end goal through a safe access route. Science-blind planners (i.e., Direct and Random) take a more direct route toward the goal, while science-aware planners (i.e., Greedy and MCTS) take an alternative route that favors spectral diversity in the collected samples. Planners may often perform worse when using a risk budget as opposed to a Euclidean distance-based budget (Figure 8.10). The reason is that sampling options are generally more constrained. Nonetheless, there is still enough margin for adaptive exploration that improves science return, as demonstrated by the superior performance of non-myopic planning over simpler methods such as myopic and blind planning. Direct paths clearly have the worst science-related performance overall because they minimize sample collection. We observe that a distance-based budget successfully bounds traversed distances (8.11). Smarter planners tend to generate riskier paths as they prioritize the collection of spectrally diverse samples regardless of the risk. In a similar fashion, smart planners that employ a risk-based budget are more likely to traverse longer distances. However, there is a critical difference when using a risk budget: d_{risk} and average $p(risk)$ are not only bounded, but also significantly reduced. This is ultimately the desired outcome: to

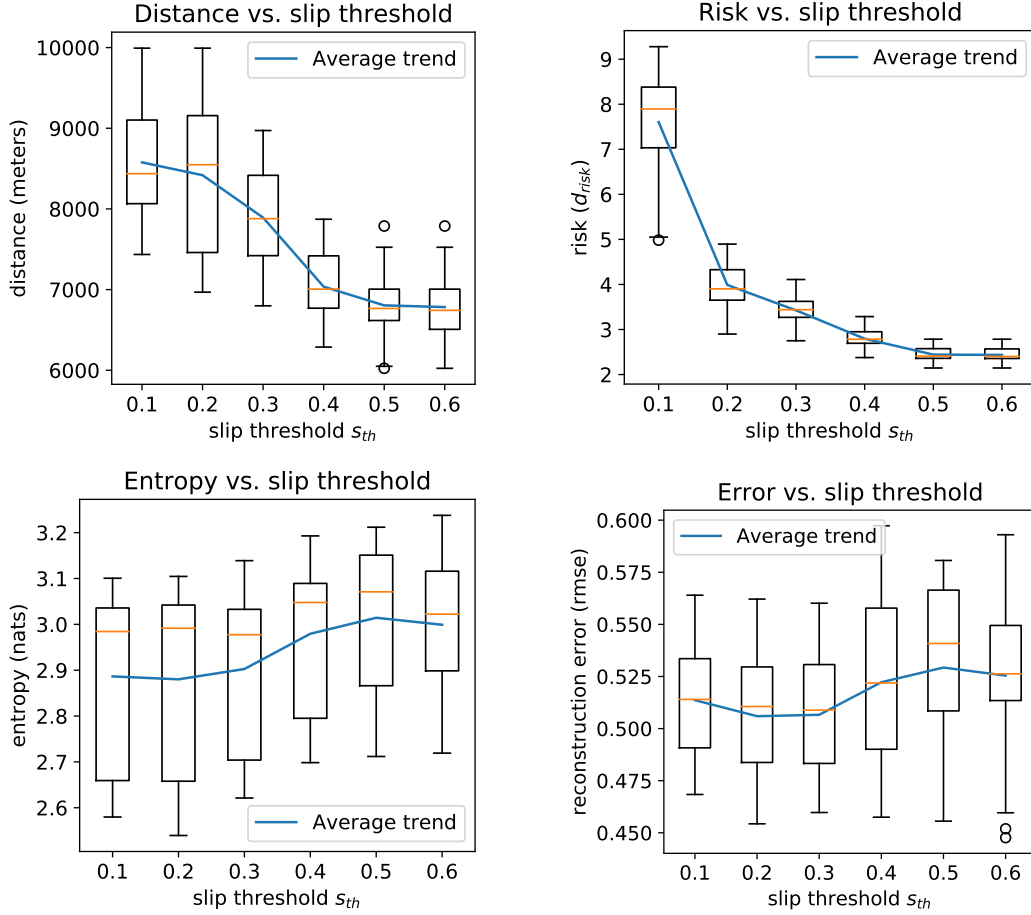


Figure 8.9: Results from the sensitivity analysis assessing the model’s risk threshold s_{th} . Distance, risk, entropy, and error change as a function of slip threshold s_{th} . Distance and risk as measured by d_{risk} tend to decrease with larger thresholds. Entropy and error do not follow a clear trend; differences in most cases are not statistically significant (p -value > 0.05).

substantially decrease risk while still enabling margin for meaningful scientific decisions. These results are further confirmed as shown in Figure 8.12. Risk-based exploration leads to longer traverses in most cases, therefore, more samples are often collected. There is no clear difference between entropy and error despite having more constrained sampling locations when using a risk-based budget. Reasons include the collection of a larger number of samples overall, as well as an efficient and representative sampling from high-risk areas – i.e., dangerous regions can be well mapped with just a few carefully selected measurements. Additionally, risk-based budgets always produce paths with a substantially lower risk.

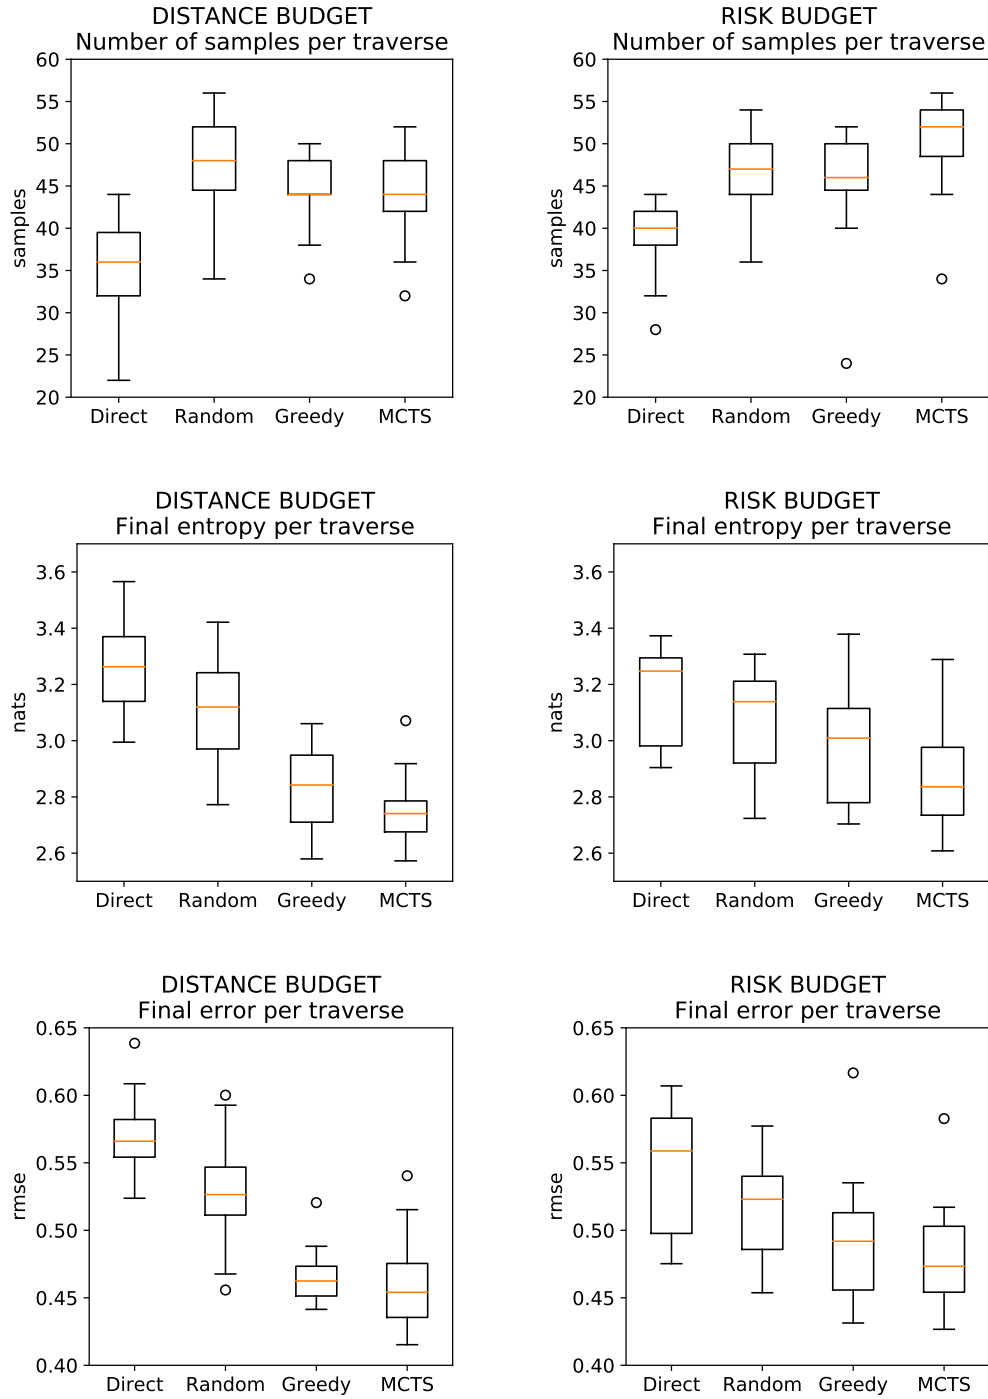


Figure 8.10: Planner's performance in terms of science productivity as measured by number of samples, entropy, and reconstruction error. Results are shown for exploration using either a distance budget or a risk budget. Smarter planners tend to perform best by collecting more samples and by further reducing entropy and error given the same allotted budget. Here all differences are statistically significant (p -value < 0.05).

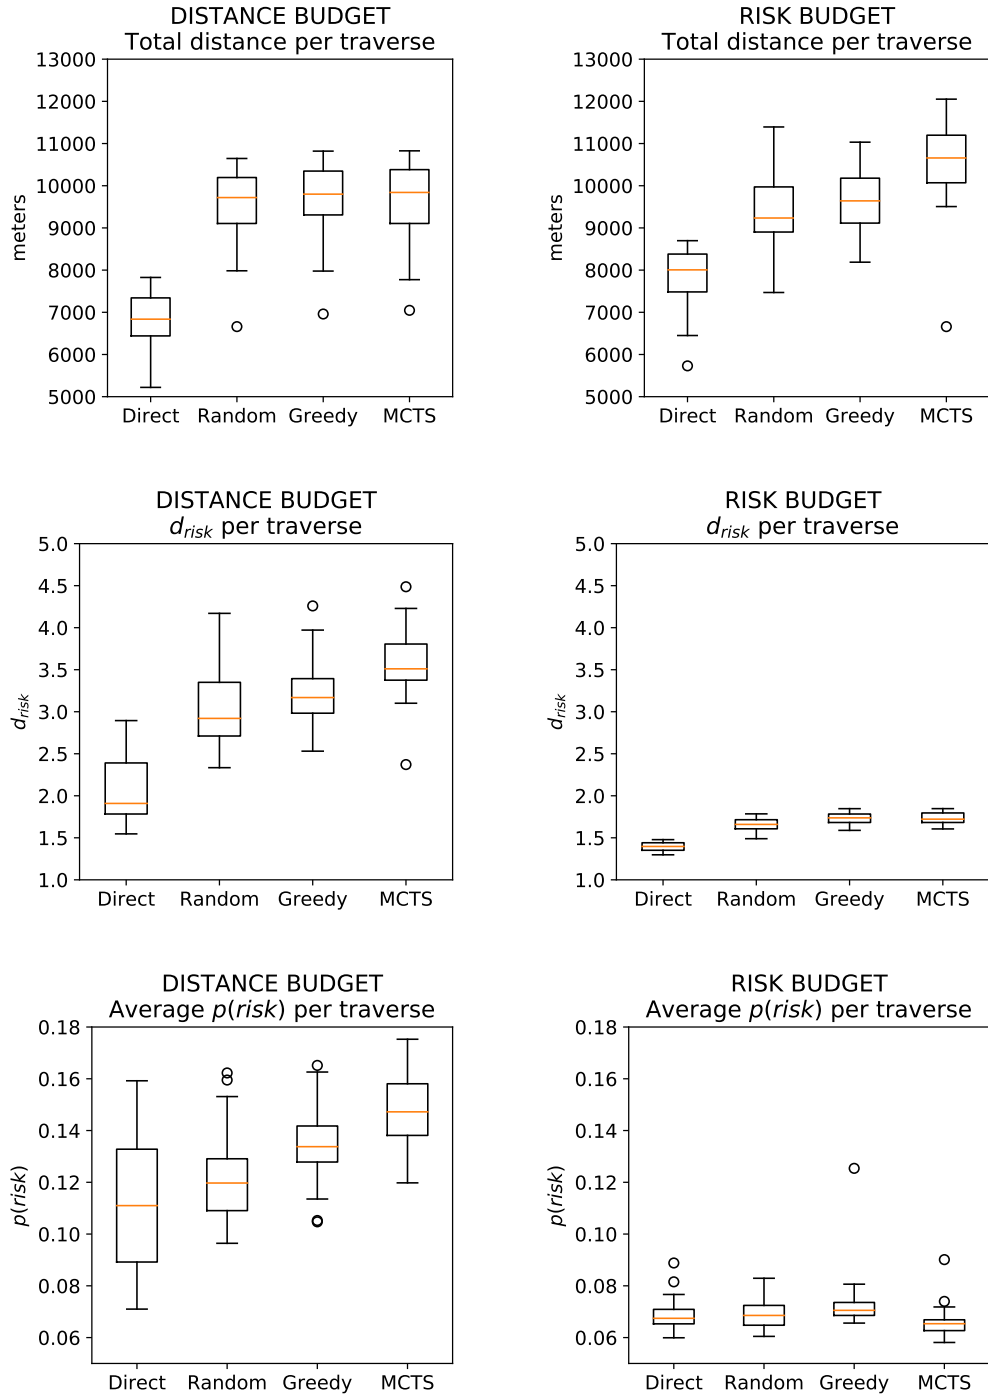


Figure 8.11: Planner’s performance in terms of traversability cost as measured by distance, risk-distance d_{risk} , and average probability of risk. Results are shown for exploration using either a distance budget or a risk budget. Smart planners tend to assume more risk when using a distance budget and often traverse longer paths when using a risk budget. Budgets are successful at bounding costs, but risk budgets are especially effective at reducing risk.

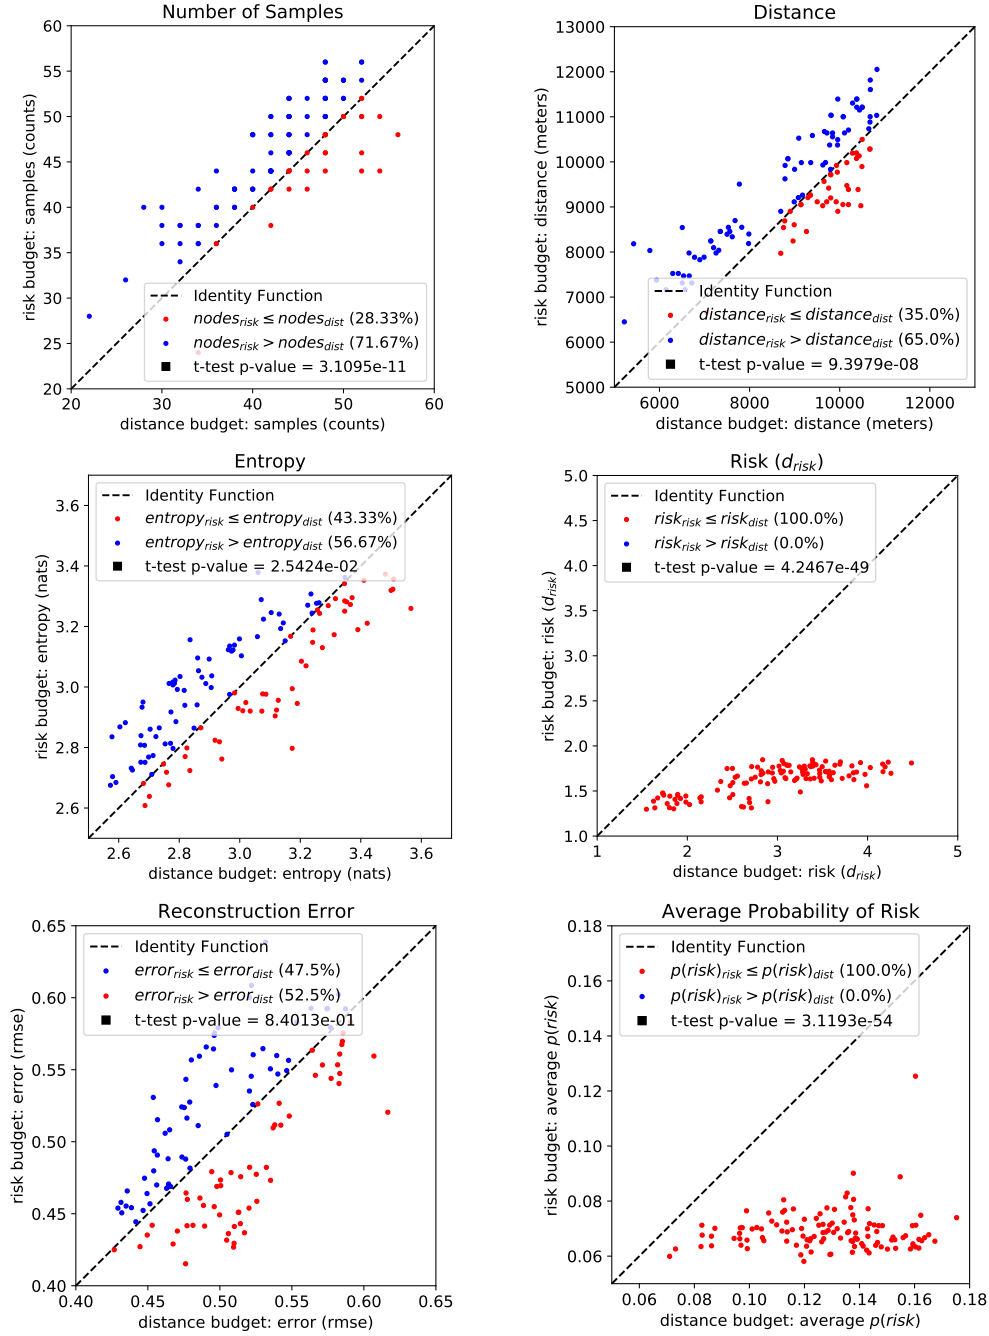


Figure 8.12: Scatter plots comparing distance against risk-based exploration in terms of science productivity (left column) and traverse costs (right column). Risk-based exploration leads to longer traverses in most cases, therefore, more samples are often collected. There is no clear difference between entropy and error despite having more constrained sampling locations when using a risk-based budget. Reasons include the collection of a larger number of samples overall, as well as an efficient and representative sampling from high-risk areas. Additionally, risk-based budgets always produce paths with a substantially lower risk.

8.5 Discussion and Conclusion

This chapter presents a Bayesian framework that addresses science productivity and traversability risk. This work focuses on Mars rover exploration, although the framework has the potential to be used for other scientific endeavors as well. Scientific information is quantified in terms of information-theoretic variables and science return is improved using Bayesian experimental design. This study specifically develops a Mars rover risk model by combining geometric and semantic information. Probabilistic slip estimation is based on the work by Cunningham et al. [37], which predicts rover wheel slippage as a function of slope and terrain class. Then, probability of risk is calculated using ideas from the work by Inotsume et al. [90] for rover path planning under slip uncertainty. Finally, path planning is formulated as a graph-based shortest path problem by deriving an equivalent risk metric that allows for the computation of minimal risk paths.

We describe a Mars simulation study that uses orbital data of Jezero Crater that is relevant for scientific and terrain analyses. Specifically, CRISM data for science analysis, HIRISE data for deriving geometric terrain information, and a geologic map produced specifically for the Mars 2020 mission [158]. Results from the simulation experiments revealed key insights into our framework. First, that total traversed distances tend to increase as a consequence of a more constrained exploration. Second, that science return is slightly affected due to less available sampling options. Third, that despite having more constraints, there is still ample margin for intelligent decision-making as demonstrated by the superior performance of non-myopic path planners over blind and greedy methods. And finally, that our framework not only bounds risk, but also reduces it by a significant amount without substantially affecting science return.

Chapter 9

Conclusion

This work has advocated a new paradigm for the next generation of science-driven robotic explorers. In this approach the robot is guided by a model that represents a state of knowledge that evolves with new information. The presented method consists of two main processes: Bayesian inference and Bayesian experimental design. Bayesian inference updates the model given new observations while Bayesian experimental design provides a probabilistic framework for optimal sampling. Overall, this research has demonstrated that Bayesian modeling can be used for science data interpretation and intelligent decision-making, and ultimately for improving science return.

9.1 Contributions

This research has made significant contributions to the field of science-driven robotic exploration. It has brought together concepts from Bayesian modeling, Bayesian experimental design, deep learning, active learning, information theory, and path planning. The key contributions of this work are as follows.

Scenarios of science-driven robotic exploration. We put together a series of relevant science investigation scenarios that can provably benefit from autonomous robotic exploration (Chapter 4). These scenarios consist of terrestrial geologic surveys at Cuprite, Nevada; Martian surface investigations at Jezero Crater and Nili Fossae; and coral reef studies at Heron Island, Australia, and Kaneohe Bay, Hawaii. The associated data sets consist of orbital, aerial, and surface data products mainly focused on spectroscopy. These scenarios permit the development and evaluation of our proposed Bayesian models via extensive simulations, and in some instances, field experiments with the rover Zoë.

Deep generative models for spectroscopic data analysis. We investigate three different deep generative models and focus on spectroscopic data interpretation (Chapter 5). The Variational Autoencoder performs spectral feature extraction. The Deep Conditional Gaussian Model

allows for probabilistic regression and spectral super-resolution. The Deep Conditional Dirichlet Model enables classification, and more generally, spectral unmixing. Thorough experimental results using synthetic and real data indicate that these methods often outperform other Bayesian and non-Bayesian approaches. But more importantly, they allow for a prompt quantification, visualization, and interpretation of learned statistical dependencies, resulting in appealing tools for scientific data analysis.

A model that combines remote and *in situ* measurements for active exploration. We introduce a spatial model that leverages contextual information from remote data to efficiently extrapolate *in situ* science features throughout an explored area (Chapter 6). The model is based on Gaussian process regression [36, 142] and builds upon the work by Thompson [162] for spatio-spectral modeling. The model is enhanced by using ideas from feature extraction in order to map and extrapolate multivariate data, which in turn is employed for map composition. Additionally, an optimal sampling framework is derived by formulating exploration as a posterior entropy minimization problem. Several sampling strategies are compared: random sampling, greedy heuristics [102], Monte Carlo tree search planning [99], and ergodic optimal control [116, 121]. These methods are evaluated in extensive simulations via the Cuprite, Mars, and Coral scenarios. Moreover, field experiments are conducted with the rover Zoë in Nevada. Results demonstrate that reduction of model uncertainty as measured by entropy leads to efficient spectroscopic and composition mapping. Results also show that science return improves when using sampling strategies that better exploit the information of the model. Additionally, important lessons were learned in terms of field robotic operations.

The science hypothesis map, a model for co-exploration. We present a Bayesian framework in which scientists can initially describe their abstract beliefs and hypotheses, and then the state of this belief evolves as the robot makes raw measurements (Chapter 7). To this end, this work constructs a spatial and hierarchical probabilistic structure that relates high-level science concepts to low-level measurable data. We call this model a *science hypothesis map*. Moreover, scientific information gain is derived from the model and is efficiently computed. Robotic exploration is formulated as a path planning problem where information gain needs to be maximized. As proof of concept, this work develops a geologic exploration model that hierarchically links geology, mineralogy, and spectroscopic measurements at Cuprite, Nevada. Results from the corresponding simulation study reveal that: a) the model can infer geology with a high accuracy under diverse initial conditions; and b) path planners that exploit information gain with farther planning horizons are better at reducing uncertainty over the investigation objectives.

A model for science and risk-aware exploration. We propose a comprehensive method for planetary rover exploration that considers both science and risk (Chapter 8). Science value is addressed using the aforementioned concepts and models. A risk model is constructed by incorporating both geometric and semantic information from the terrain. This is particularly important since rovers often experience hazardous conditions due to geometrically benign surfaces such as sand dunes. Our approach builds upon the works by Cunningham et al. [37] for probabilistic slip

estimation and Inotsume et al. [90] for rover path planning under slip uncertainty. Exploration is formulated as a graph-based orienteering problem subject to a risk budget. Moreover, a custom risk metric allows for the computation of minimal risk paths that enforce risk constraints at any location. Results from a realistic Jezero Crater simulation study indicate that risk can be bounded and substantially reduced without sacrificing too much science productivity.

9.2 Future Work

There are a number of specific improvements to this work. We next provide a summary:

- New science-driven robotic exploration scenarios could be explored in the future beyond the ones introduced in Chapter 4. Examples include agriculture [107], terrestrial biology [129], and other marine science endeavors besides coral reef studies [48, 159].
- The deep generative models presented in Chapter 5, specifically the DCGM and DCDM, could be improved upon by addressing epistemic uncertainty [95] in addition to aleatoric noise. Dropout [53] and ensemble [108] approaches may help in this endeavor. Other unimodal distributions could be explored; examples include the Poisson and Gamma distributions. The only requirement would be to have a tractable way of computing their associated log-likelihood. Finally, we would like to address probabilistic mixture models and multimodal distributions, and eventually develop architectures such as the Deep Conditional Gaussian Mixture Model. A simple approach would be to use a predefined number of mixtures; a more interesting but difficult method would learn this number of mixtures from the data.
- The active mapping model presented in Chapter 6 could be improved upon. Gaussian processes are well known for having scalability constraints, which was not a problem given our assumption of *in situ* data scarcity. However, dense *in situ* sampling will require a more scalable model. To this end, there are more scalable Gaussian processes that employ modifications based on kernel interpolations [179] or deep learning [87]. Additionally, although Gaussian processes are powerful tools for modeling epistemic uncertainty due to their ability to adapt to new data, they tend to operate with constant noise models. We would like to expand our current model so it can capture heteroscedastic uncertainties [110]. Another topic we would like to address is time modeling: geologic models may evolve very slowly and thus be considered as time-invariant for simplicity, but that is certainly not the case for coral reefs. We plan to address this by integrating temporal measurements into the Gaussian processes comprising our model, a thing that was not possible due to limited temporal sampling during PRISM data collection. Furthermore, we would like our model to incorporate data from other instruments. Spectrometers, despite their rich information and undeniable utility, are currently expensive to acquire and difficult to deploy in the field, especially on board AUVs. Also, spectral analysis usually requires a certain degree of expert knowledge together with access to special spectral libraries. Hence, we would like to adapt our models for instruments that may be easier to

deploy onboard robots, such as cameras. For instance, normal images are easier for people to understand and label; besides, widely-available computer vision and deep learning models could be used for automatic analysis. Finally, we intend to conduct experiments in actual coral reefs with the objective of helping divers or AUVs collect informative samples in real time.

- The science hypothesis map presented in Chapter 7 could be turned into a continuous model by leveraging methods such as Gaussian processes and Dirichlet distributions [120], allowing for the mapping of categorical variables in a fast and principled manner. In general, we would like to reduce the model’s heavy reliance on a spatial discretization set at the outset. Moreover, we plan to improve the fidelity and realism of both the constraints and the measurement data. Additionally, we would like to address human-scientist interactions in more detail. Finally, we are looking forward to deploying this method on the upcoming TREX campaign in the Painted Desert, Arizona.
- The science and risk-aware model presented in Chapter 8 could be expanded to integrate global and local risk planning. Specifically, we would like to build adaptive risk models that update with *in situ* data [73]. We could achieve this by leveraging contextual information from remote data in a similar fashion as our science models. Options for achieving this goal include using Gaussian processes or Markov random fields. Furthermore, we would like to develop a similar risk model for the rover Zoë.

Finally, there are several promising avenues for future work in science-driven robotic exploration. We discuss three main directions:

1. **Multimodal information models.** Our proposed models have assumed that the robot carries one instrument, and that there is essentially one source of remote data. However, this is often not true in real robotic exploration. For instance, Mars rovers have several instruments onboard such as cameras, radars, weather stations, and assorted spectrometers (e.g., x-ray, ultraviolet, infrared, etc.). The same happens with orbiters: they also carry different instruments such as laser altimeters and thermal cameras. Hence, we would like to develop new models that integrate and leverage information from multiple instruments, both remote and *in situ*. Ultimately, we would be interested in extending our current methods to produce a framework that allows scientists to compare the science value of different instruments as measured by information-theoretic variables.
2. **Multi-objective optimization models.** This work has formulated science exploration as an optimization problem where an information-theoretic function is either minimized (entropy) or maximized (information gain). It has also explored simple (e.g., number of samples, local slope) and complex constraint functions (i.e., risk budget). Nevertheless, science exploration often deals with concurrent mission objectives and multiple simultaneous constraints such as time, energy, and power consumption. Consequently, we would like to formulate multi-objective optimization models that allow for richer and more realistic exploration scenarios.
3. **Multi-agent exploration models.** This research has worked under the assumption that sci-

ence investigations are being conducted by a single robotic explorer. This may be the case for most planetary rover missions today, but this situation is rapidly changing. For instance, the Mars 2020 mission is deploying the Mars Helicopter together with the Perseverance rover. Hook et al. [84] are modeling a network of small rovers for Martian exploration missions. On Earth we are very interested in having Zoë work closely with other robots such as drones.

9.3 Closing Remarks

This is an exciting time as recent advances in hardware, sensing, and artificial intelligence have produced robotic technologies that would have seemed like science fiction a few decades ago. Examples include house-cleaning robots, accessible consumer drones, and self-driving cars. It is an undeniable fact that robots are becoming part of our daily lives.

This research addresses a much less common application of robotics: scientific exploration. Robots that perform science can be powerful tools for answering profound questions. What is beyond Earth? Is life sustainable on other planets? Are we alone in the universe?

This work has ultimately taken firm steps towards enabling the next generation of robotic explorers “to explore strange new worlds, to seek out new life and new civilizations, and to boldly go where no one has gone before!” [148].

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 and Machine Intelligence*, 34(11):2274–2281, 2012. ISSN 01628828. doi: 10.1109/TPAMI.2012.120. 3.2, 3.3
- [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. 1.1
- [3] Akash Arora, Robert Fitch, and Salah Sukkarieh. An approach to autonomous science by modeling geological knowledge in a Bayesian framework. *IEEE International Conference on Intelligent Robots and Systems*, 2017-Sept:3803–3810, 2017. ISSN 21530866. doi: 10.1109/IROS.2017.8206230. 6.2, 6.4, 7.2, 7.4, 8.2.1
- [4] David Arthur and Sergei Vassilvitskii. K-means++: The advantages of careful seeding. *Proceedings of the Annual ACM-SIAM Symposium on Discrete Algorithms*, 07-09-January-2007:1027–1035, 2007. 5.3.2, 5.5.2
- [5] R. Arvidson, D. Fuller, M. Heverly, K. Iagnemma, J. Lin, J. Matthews, R. Sletten, and N. Stein. Mars Science Laboratory Curiosity Rover Terramechanics Initial Results. In *Lunar and Planetary Science Conference*, volume 44, 2013. doi: 10.1007/s10569-009-9234-1. 8.2.2
- [6] Raymond E. Arvidson, Karl D. Iagnemma, Mark Maimone, Abigail A. Fraeman, Feng Zhou, Matthew C. Heverly, Paolo Belluta, David Rubin, Nathan T. Stein, John P. Grotzinger, and Ashwin R. Vasavada. Mars Science Laboratory Curiosity Rover Megaripple Crossings up to Sol 710 in Gale Crater. *Journal of Field Robotics*, 34(3):495–518, 2016. doi: 10.1002/rob.21647. URL <http://onlinelibrary.wiley.com/doi/10.1002/rob.21514/abstract>. 8.2.2
- [7] J. F. Bell, J. N. Maki, G. L. Mehall, M. A. Ravine, M. A. Caplinger, and Mastcam-Z Team. Mastcam-Z: Designing a Geologic, Stereoscopic, and Multispectral Pair of Zoom Cameras for the NASA Mars 2020 Rover. In *3rd International Workshop on Instrumentation for Planetary Mission*, volume 1980, page 4126, October 2016. 4.2.3

- [8] Jojanneke Berg, Van Den, 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. doi: 10.1046/j.1365-246X.2003.02048.x. URL <https://onlinelibrary.wiley.com/doi/abs/10.1046/j.1365-246X.2003.02048.x>. 7.4
- [9] José M. Bernardo. Expected Information as Expected Utility. *The Annals of Statistics*, 7(3):686—690, 1979. 7.4
- [10] Jonathan Binney and Gaurav S. Sukhatme. Branch and bound for informative path planning. In *IEEE International Conference on Robotics and Automation*, pages 2147–2154. IEEE, 2012. ISBN 9781467314039. doi: 10.1109/ICRA.2012.6224902. 6.2, 6.4, 7.4, 7.5.2, 8.2.1
- [11] Jonathan Binney, Andreas Krause, and Gaurav S. Sukhatme. Informative path planning for an autonomous underwater vehicle. *Proceedings - IEEE International Conference on Robotics and Automation*, pages 4791–4796, 2010. ISSN 10504729. doi: 10.1109/ROBOT.2010.5509714. 6.2, 7.2, 7.2
- [12] J. L. Bishop. Chapter 3 - Remote Detection of Phyllosilicates on Mars and Implications for Climate and Habitability. In N. A. Cabrol and E. A. Grin, editors, *From Habitability to Life on Mars*, pages 37 – 75. Elsevier, 2018. ISBN 978-0-12-809935-3. doi: <https://doi.org/10.1016/B978-0-12-809935-3.00003-7>. URL <http://www.sciencedirect.com/science/article/pii/B9780128099353000037>. 3.1, 8.2.2
- [13] Daniel L. Bongiorno, Mitch Bryson, Tom C.L. Bridge, Donald G. Dansereau, and Stefan B. Williams. Coregistered Hyperspectral and Stereo Image Seafloor Mapping from an Autonomous Underwater Vehicle. *Journal of Field Robotics*, 35(3):312–329, 2017. ISSN 15564967. doi: 10.1002/rob.21713. 6.2
- [14] Ricardo Augusto Borsoi, Tales Imbiriba, and José Carlos Moreira Bermudez. Deep Generative Endmember Modeling: An Application to Unsupervised Spectral Unmixing. *arXiv preprint arXiv:1902.05528*, 2019. URL <http://arxiv.org/abs/1902.05528>. 5.2
- [15] 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. 3.2
- [16] A. J. Brown, C. E. Viviano, and T. A. Goudge. Olivine-Carbonate Mineralogy of the Jezero Crater Region. *Journal of Geophysical Research: Planets*, 125(3), 2020. ISSN 21699100. doi: 10.1029/2019JE006011. 1.2
- [17] Alberto Candela, David Thompson, Eldar Noe Dobra, and David Wettergreen. Planetary robotic exploration driven by science hypotheses for geologic mapping. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 3811–3818,

- Vancouver, 2017. ISBN 9781538626825. doi: 10.1109/IROS.2017.8206231. 6.2, 8.2.2
- [18] Alberto Candela, David R Thompson, and David Wettergreen. Automatic Experimental Design Using Deep Generative Models Of Orbital Data. In *International Symposium on Artificial Intelligence, Robotics and Automation in Space (i-SAIRAS)*, 2018. 5.2, 5.4.2
 - [19] Alberto Candela, Kevin Edelson, and David Wettergreen. Mars Rover Exploration Combining Remote and In Situ Measurements for Wide-Area Mapping. In *International Symposium on Artificial Intelligence, Robotics and Automation in Space (i-SAIRAS)*, 2020. 6.6
 - [20] Alberto Candela, Suhit Kodgule, Kevin Edelson, Srinivasan Vijayarangan, David R Thompson, Eldar Noe Dobrea, and David Wettergreen. Planetary Rover Exploration Combining Remote and In Situ Measurements for Active Spectroscopic Mapping. In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 5986–5993, Paris, 2020. ISBN 9781728173955. 6.5
 - [21] Alberto Candela, David R Thompson, David Wettergreen, Kerry Cawse-Nicholson, Sven Geier, Michael L. Eastwood, and Robert O. Green. Probabilistic Super Resolution for Mineral Spectroscopy. In *AAAI Conference on Artificial Intelligence (AAAI)*, pages 13241–13247, 2020. 5.1, 5.4.1
 - [22] Alberto Candela, Kevin Edelson, Michelle M. Gierach, David R. Thompson, Gail Woodward, and David Wettergreen. Using remote sensing and in situ measurements for efficient mapping and optimal sampling of coral reefs. *Frontiers in Marine Science*, 8:1195, 2021. ISSN 2296-7745. doi: 10.3389/fmars.2021.689489. URL <https://www.frontiersin.org/article/10.3389/fmars.2021.689489>. 6.7
 - [23] Carnegie Mellon University. CMU’s MoonRanger Will Search for Water at Moon’s South Pole. <https://www.cmu.edu/news/stories/archives/2020/september/moonranger-to-search-for-water.html>, 2020. Accessed: 2020-12-01. 1.1
 - [24] A. Castano, A. Fukunaga, J. Biesiadecki, L. Neakrase, P. Whelley, R. Greeley, M. Lemon, 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. 7.2
 - [25] Kathryn Chaloner and Isabella Verdinelli. Bayesian Experimental Design: A Review. *Statistical Science*, 10(3):273—304, 1995. 2.4, 6.4, 8.2.1
 - [26] 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. 3.2

- [27] Chandra Chekuri and Martin Pál. A recursive greedy algorithm for walks in directed graphs. *Proceedings - Annual IEEE Symposium on Foundations of Computer Science, FOCS*, 2005:245–253, 2005. ISSN 02725428. doi: 10.1109/SFCS.2005.9. 6.4, 7.5.2, 8.2.1
- [28] Steve Chien, Robert Sherwood, Daniel Tran, Benjamin Cichy, Gregg R Rabideau, Rebecca Castano, Ashley Gerard Davies, and Darrell Boyer. Using autonomy flight software to improve science return on earth observing one. *Journal of Aerospace Computing, Information, and Communication*, 2(April):196–216, 2005. ISSN 1542-9423. doi: 10.2514/1.12923. 7.2
- [29] F. Chollet, J.J. Allaire, et al. R interface to keras. <https://github.com/rstudio/keras>, 2017. 5.4.2, 5.5.2
- [30] Philip R Christensen, Bruce M Jakosky, Hugh H Kieffer, Michael C Malin, Harry Y Mcsween, Kenneth Nealson, Greg L Mehall, Steven H Silverman, Steven Ferry, Michael Caplinger, and Michael Ravine. THE THERMAL EMISSION IMAGING SYSTEM (THEMIS) FOR THE MARS 2001 ODYSSEY MISSION. *Space Science Reviews*, 110: 85–130, 2004. 3.1
- [31] R.N. Clark et al. Usgs digital spectral library splib06a. *U.S. Geological Survey, Digital Data Series 231*, 2007. 7.2, 7.6.1
- [32] Roger N. Clark. Spectroscopy of rocks and minerals, and principles of spectroscopy. In Andrew N. Rencz and Robert A. Ryerson, editors, *Manual of Remote Sensing, Volume 3, Remote Sensing for the Earth Sciences*, chapter 1, pages 3—58. John Wiley and Sons, New York, 3rd edition, 1999. ISBN 978-0-471-29405-4. 1.4, 3.1, 3.2, 3.2, 5.1, 5.4.2, 7.3.3
- [33] Roger N. Clark, Gregg A. Swayze, K. Eric Livo, Raymond F. Kokaly, Steve J. Sutley, 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*, 108(12):1–44, 2003. ISSN 0148-0227. doi: 10.1029/2002JE001847. URL <http://doi.wiley.com/10.1029/2002JE001847>. 3.2, 5.4.2, 7.6.1, 7.6.2, 7.7
- [34] Robert Costanza, Rudolf de Groot, Paul Sutton, Sander van der Ploeg, Sharolyn J. Anderson, Ida Kubiszewski, Stephen Farber, and R. Kerry Turner. Changes in the global value of ecosystem services. *Global Environmental Change*, 26(1):152–158, 2014. ISSN 09593780. doi: 10.1016/j.gloenvcha.2014.04.002. URL <http://dx.doi.org/10.1016/j.gloenvcha.2014.04.002>. 4.3
- [35] Thomas M. Cover and Joy A. Thomas. *Elements of Information Theory*. Wiley- Interscience, Hoboken, NJ, 2nd edition, 2006. 2.3, 2.3, 5.4.1, 6.4, 6.4, 8.2.1
- [36] Noel A.C. Cressie. *Statistics for Spatial Data*. John Wiley & Sons, Inc., 1993. ISBN 9780471002550. doi: 10.1002/9781119115151. 6.2, 9.1

- [37] Chris Cunningham, Masahiro Ono, Issa Nesnas, Jeng Yen, and William L. Whittaker. Locally-adaptive slip prediction for planetary rovers using Gaussian processes. In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 5487–5494. IEEE, 2017. ISBN 9781509046331. doi: 10.1109/ICRA.2017.7989646. 8.2.2, 8.3.2, 8.3, 8.5, 9.1
- [38] Ryan Dahl, Mohammad Norouzi, and Jonathon Shlens. Pixel Recursive Super Resolution. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 5449–5458, 2017. ISBN 9781538610329. doi: 10.1109/ICCV.2017.581. 5.2, 5.4.1
- [39] E. W. Dijkstra. A Note on Two Problems in Connexion with Graphs. *Numerische Mathematik*, 1:269–271, 1959. 7.5.2, 7.6.2, 8.3.3, 8.4.1
- [40] 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. 3.2
- [41] Morris L. Eaton. *Multivariate Statistics: a Vector Space Approach*. John Wiley & Sons, 1st edition, 1983. ISBN 978-0471027768. 5.4.2
- [42] Kevin Edelson. Ergodic trajectory optimization for information gathering. Master’s thesis, Carnegie Mellon University, Pittsburgh, PA, October 2020. 6.7.1
- [43] Tara Estlin, Daniel Gaines, Benjamin Bornstein, Steve Schaffer, Vandi Tompkins, David R. Thompson, Alphan Altinok, Robert C Anderson, Michael Burl, Rebecca Castaño, Diana Blaney, Lauren De Flores, Tony Nelson, and Roger Wiens. Automated Targeting for the MSL Rover ChemCam Spectrometer. *International Symposium on Artificial Intelligence, Robotics and Automation in Space (i-SAIRAS)*, 2014. 1.1
- [44] 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 Transactions on Intelligent Systems and Technology*, 3(3):1–19, 2012. ISSN 21576904. doi: 10.1145/2168752.2168764. URL <http://dl.acm.org/citation.cfm?doid=2168752.2168764>. 7.1, 7.2
- [45] 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>. 3.2
- [46] Genevieve Flaspohler, Victoria Preston, Anna P. M. Michel, Yogesh Girdhar, and Nicholas Roy. Information-Guided Robotic Maximum Seek-and-Sample in Partially Observable Continuous Environments. *IEEE Robotics and Automation Letters*, 4(4):3782–3789,

2019. ISSN 2377-3766. doi: 10.1109/lra.2019.2929997. 6.2

- [47] Greydon Foil and David S. Wettergreen. Efficiently Sampling from Underlying Physical Models. In *Robotics: Science and Systems*, Pittsburgh, PA, 2018. URL https://www.ri.cmu.edu/pub/_files/2016/10/g{_}foil{_}robotics{_}2016.pdf. 6.2
- [48] Trygve Olav Fossum, John Ryan, Tapan Mukerji, Jo Eidsvik, Thom Maughan, Martin Ludvigsen, and Kanna Rajan. Compact models for adaptive sampling in marine robotics. *International Journal of Robotics Research*, 39(1):127–142, 2020. ISSN 17413176. doi: 10.1177/0278364919884141. 6.2, 9.2
- [49] H. Fujisada, F. Sakuma, A. Ono, and M. Kudoh. Design and preflight performance of aster instrument protoflight model. *IEEE Transactions on Geoscience and Remote Sensing*, 36(4):1152–1160, 1998. ISSN 01962892. doi: 10.1109/36.701022. 4.2.1
- [50] Daniel Gaines, Gary Doran, Heather Justice, Gregg Rabideau, Steve Schaffer, Vandana Verma, Kiri Wagstaff, Ashwin Vasavada, William Huffman, Robert Anderson, Ryan Mackey, and Tara Estlin. Productivity Challenges for Mars Rover Operations: A Case Study of Mars Science Laboratory Operations Executive Summary. Technical report, Technical Report D-97908, Jet Propulsion Laboratory, 2016. URL https://ai.jpl.nasa.gov/public/papers/gaines{_}report{_}roverProductivity.pdf. 1.1, 6.1
- [51] Daniel Gaines, Joseph Russino, Gary Doran, Ryan Mackey, Michael Paton, Brandon Rothrock, Steve Schaffer, Ali-Akbar Agha-Mohammadi, Chet Joswig, Heather Justice, Ksenia Kolcio, Jacek Sawoniewicz, Vincent Wong, Kathryn Yu, Gregg Rabideau, Robert Anderson, and Ashwin Vasavada. Self-Reliant Rover Design for Increasing Mission Productivity. In *International Symposium on Artificial Intelligence, Robotics and Automation in Space (i-SAIRAS)*, 2018. URL https://ai.jpl.nasa.gov/public/papers/gaines{_}planrob2018{_}reliant.pdf. 1.1
- [52] Daniel Gaines, Gary Doran, Michael Paton, Brandon Rothrock, Joseph Russino, Ryan Mackey, Robert Anderson, Raymond Francis, Chet Joswig, Heather Justice, Ksenia Kolcio, Gregg Rabideau, Steve Schaffer, Jacek Sawoniewicz, Ashwin Vasavada, Vincent Wong, Kathryn Yu, and Ali-akbar Agha-mohammadi. Self-reliant rovers for increased mission productivity. *Journal of Field Robotics*, 37(7):1171–1196, 2020. doi: <https://doi.org/10.1002/rob.21979>. URL <https://onlinelibrary.wiley.com/doi/abs/10.1002/rob.21979>. 8.2.2
- [53] Yarin Gal and Zoubin Ghahramani. Dropout as a Bayesian approximation: Representing model uncertainty in deep learning. *33rd International Conference on Machine Learning, ICML 2016*, 3:1651–1660, 2016. 5.2, 9.2
- [54] Marc Gallant, Alex Ellery, and Joshua Marshall. Science-influenced mobile robot guid-

ance using Bayesian Networks. In *Canadian Conference on Electrical and Computer Engineering*, pages 1135–1139. IEEE, 2011. ISBN 9781424497898. doi: 10.1109/CCECE.2011.6030639. 7.2

- [55] Yang Gao and Steve Chien. Review on space robotics: Toward top-level science through space exploration. *Science Robotics*, 2(7), 2017. ISSN 24709476. doi: 10.1126/scirobotics.aan5074. 1.1
- [56] J. P. Gattuso, A. Magnan, R. Billé, W. W.L. Cheung, E. L. Howes, F. Joos, D. Allemand, L. Bopp, S. R. Cooley, C. M. Eakin, O. Hoegh-Guldberg, R. P. Kelly, H. O. Pörtner, A. D. Rogers, J. M. Baxter, D. Laffoley, D. Osborn, A. Rankovic, J. Rochette, U. R. Sumaila, S. Treyer, and C. Turley. Contrasting futures for ocean and society from different anthropogenic CO₂ emissions scenarios. *Science*, 349(6243), 2015. ISSN 10959203. doi: 10.1126/science.aac4722. 4.3
- [57] Leon Gatys, Alexander Ecker, and Matthias Bethge. A Neural Algorithm of Artistic Style. *Journal of Vision*, 16(12):326, 201. ISSN 1534-7362. doi: 10.1167/16.12.326. 5.2
- [58] Shivam Gautam, Bishwamoy Sinha Roy, Alberto Candela, David S. Wettergreen, Bishwamoy Sinha Roy, Alberto Candela, and David S. Wettergreen. Science-aware exploration using entropy-based planning. *IEEE International Conference on Intelligent Robots and Systems*, pages 3819–3825, 2017. ISSN 21530866. doi: 10.1109/IROS.2017.8206232. URL http://ri.cmu.edu/wp-content/uploads/2017/08/IROS17{}_1472{}_FI.pdf<http://ieeexplore.ieee.org/document/8206232/>. 8.2.1
- [59] 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. 7.2
- [60] Steven B. Goldberg, Mark W. Maimone, and Lany Matthies. Stereo vision and rover navigation software for planetary exploration. In *IEEE Aerospace Conference Proceedings*, volume 5, pages 2025–2036, 2002. ISBN 078037231X. doi: 10.1109/AERO.2002.1035370. 8.2.2
- [61] Ian J Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative Adversarial Nets. In Z. Ghahramani, M. Welling, C. Cortes, N. D. Lawrence, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems 27*, pages 2672–2680. Curran Associates, Inc., 2014. URL <http://papers.nips.cc/paper/5423-generative-adversarial-nets.pdf>. 5.2
- [62] Ian J Goodfellow, Yoshua Bengio, and Aaron Courville. *Deep Learning*. MIT Press, 2016. 5.2, 5.3.2
- [63] A.A. Green, M. Berman, P. Switzer, and M.D. Craig. A transformation for ordering multispectral data in terms of image quality with implications for noise removal. *IEEE Trans-*

actions on Geoscience and Remote Sensing, 26(1):65–74, 1988. doi: 10.1109/36.3001.5.4.2

- [64] E. P. Green, P. J. Mumby, A. J. Edwards, and C. D. Clark. A review of remote sensing for the assessment and management of tropical coastal resources. *Coastal Management*, 24(1):1–40, 1996. doi: 10.1080/08920759609362279. URL <https://doi.org/10.1080/08920759609362279>. 4.3, 6.1
- [65] R. Green. The earth surface mineral dust source investigation (EMIT) using new imaging spectroscopy measurements from space. In *International Geoscience and Remote Sensing Symposium (IGARSS)*, 2019. 5.6
- [66] Michael Gross. The past and future habitability of planet Mars. *Current Biology*, 24(5):R175–R178, 2014. ISSN 09609822. doi: 10.1016/j.cub.2014.02.029. URL <http://dx.doi.org/10.1016/j.cub.2014.02.029>. 1.1
- [67] Jacques Hadamard. Résolution d’une question relative aux déterminants. *Bulletin des Sciences Mathématiques*, 17:240–246, 1893. 6.4
- [68] 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. *IEEE Aerospace Conference Proceedings*, pages 1–7, 2011. ISSN 1095323X. doi: 10.1109/AERO.2011.5747395. 4.2.1, 7.5.1
- [69] Bruce Hapke. Bidirectional reflectance spectroscopy: 1. theory. *Journal of Geophysical Research: Solid Earth*, 86(B4):3039–3054, 1981. doi: <https://doi.org/10.1029/JB086iB04p03039>. URL <https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1029/JB086iB04p03039>. 3.2
- [70] Bruce Hapke. *Theory of Reflectance and Emittance Spectroscopy*. Cambridge University Press, second edition edition, 2012. 3.2
- [71] David Haussler and Manfred Opper. General bounds on the mutual information between a parameter and n conditionally independent observations. In *Proceedings of the 8th Annual Conference on Computational Learning Theory, COLT 1995*, volume 1995-Janua, pages 402–411, 1995. ISBN 0897917235. doi: 10.1145/225298.225347. 7.4
- [72] John D. Hedley, Chris M. Roelfsema, Iliana Chollett, Alastair R. Harborne, Scott F. Heron, Scarla J. Weeks, William J. Skirving, Alan E. Strong, C. Mark Eakin, Tyler R.L. Christensen, Victor Ticzon, Sonia Bejarano, and Peter J. Mumby. Remote sensing of coral reefs for monitoring and management: A review. *Remote Sensing*, 8(2), 2016. ISSN 20724292. doi: 10.3390/rs8020118. 4.3, 6.1
- [73] Gabrielle Hedrick, Nicholas Ohi, and Yu Gu. Terrain-Aware Path Planning and Map Update for Mars Sample Return Mission. *IEEE Robotics and Automation Letters*, 5(4): 5181–5188, 2020. ISSN 23773766. doi: 10.1109/LRA.2020.3005123. 8.2.2, 9.2

- [74] Alejandro Hidalgo-Paniagua, Miguel A. Vega-Rodríguez, Joaquín Ferruz, and Nieves Pavón. Solving the multi-objective path planning problem in mobile robotics with a firefly-based approach. *Soft Computing*, 21(4):949–964, 2017. ISSN 14337479. doi: 10.1007/s00500-015-1825-z. 6.7.1
- [75] G. E. Hinton. Deep belief networks. *Scholarpedia*, 4(5):5947, 2009. doi: 10.4249/scholarpedia.5947. revision #91189. 5.2
- [76] Eric J. Hochberg and Marlin J. Atkinson. Capabilities of remote sensors to classify coral, algae, and sand as pure and mixed spectra. *Remote Sensing of Environment*, 85(2):174–189, 2003. ISSN 00344257. doi: 10.1016/S0034-4257(02)00202-X. 4.3
- [77] Eric J. Hochberg and Michelle M. Gierach. Missing the reef for the corals: Unexpected trends between coral reef condition and the environment at the ecosystem scale. *Frontiers in Marine Science*, 8:1191, 2021. ISSN 2296-7745. doi: 10.3389/fmars.2021.727038. URL <https://www.frontiersin.org/article/10.3389/fmars.2021.727038>. 4.3, 6.1
- [78] Eric J. Hochberg, Marlin J. Atkinson, and Serge Andréfouët. Spectral reflectance of coral reef bottom-types worldwide and implications for coral reef remote sensing. *Remote Sensing of Environment*, 85(2):159–173, 2003. ISSN 00344257. doi: 10.1016/S0034-4257(02)00201-8. 4.3, 4.3
- [79] 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. 1.1, 7.1
- [80] O. Hoegh-Guldberg, P. J. Mumby, A. J. Hooten, R. S. Steneck, P. Greenfield, E. Gomez, C. D. Harvell, P. F. Sale, A. J. Edwards, K. Caldeira, N. Knowlton, C. M. Eakin, R. Iglesias-Prieto, N. Muthiga, R. H. Bradbury, A. Dubi, and M. E. Hatzitolos. Coral Reefs Under Rapid Climate Change and Ocean Acidification. *Science*, 318(5857):1737–1742, 2007. ISSN 00368075. doi: 10.1126/science.1152509. 4.3
- [81] Ove Hoegh-Guldberg, Elvira S. Poloczanska, William Skirving, and Sophie Dove. Coral reef ecosystems under climate change and ocean acidification. *Frontiers in Marine Science*, 4(MAY), 2017. ISSN 22967745. doi: 10.3389/fmars.2017.00158. 4.3
- [82] Joseph P. Hoffbeck and David A. Landgrebe. Covariance matrix estimation and classification with limited training data. *IEEE Transactions on Pattern Analysis and Machine*

Intelligence, 18(7):763–767, 1996. ISSN 01628828. doi: 10.1109/34.506799. 7.5.1

- [83] Geoffrey A. Hollinger and Gaurav S. Sukhatme. Sampling-based robotic information gathering algorithms. *International Journal of Robotics Research*, 33(9):1271–1287, 2014. ISSN 17413176. doi: 10.1177/0278364914533443. 6.4, 7.4, 8.2.1
- [84] Joshua Vander Hook, Tiago Vaquero, Federico Rossi, Martina Troesch, Marc Sanchez Net, Joshua Schoolcraft, Jean Pierre De La Croix, and Steve Chien. Mars on-site shared analytics information and computing. In *Proceedings International Conference on Automated Planning and Scheduling, ICAPS*, pages 707–715, 2019. ISBN 9781577358077. 3
- [85] Briony H.N. Horgan, Ryan B. Anderson, Gilles Dromart, Elena S. Amador, and Melissa S. Rice. The mineral diversity of Jezero crater: Evidence for possible lacustrine carbonates on Mars. *Icarus*, 339:113526, 2020. 4.2.3
- [86] Zhiting Hu, Zichao Yang, Ruslan Salakhutdinov, and Eric P. Xing. On Unifying Deep Generative Models. In *International Conference on Learning Representations*, 2018. 5.2
- [87] Wenbing Huang, Deli Zhao, Fuchun Sun, Huaping Liu, and Edward Chang. Scalable Gaussian process regression using deep neural networks. In *IJCAI International Joint Conference on Artificial Intelligence*, pages 3576–3582, 2015. ISBN 9781577357384. 9.2
- [88] Bernard E Hubbard, James K Crowley, and David R Zimbelman. Comparative alteration mineral mapping using visible to shortwave infrared (0.4–2.4 μm) Hyperion, ALI, and ASTER imagery. *IEEE Transactions on Geoscience and Remote Sensing*, 41(6): 1401–1410, 2003. 5.1
- [89] Hiroaki Inotsume, Colin Creager, David Wettergreen, and William Whittaker. Finding Routes for Efficient and Successful Slope Ascent for Exploration Rovers. In *International Symposium on Artificial Intelligence, Robotics and Automation in Space (i-SAIRAS)*, 2016. 8.2.2
- [90] Hiroaki Inotsume, Takashi Kubota, and David Wettergreen. Robust Path Planning for Slope Traversing under Uncertainty in Slip Prediction. *IEEE Robotics and Automation Letters*, 5(2):3390–3397, 2020. ISSN 23773766. doi: 10.1109/LRA.2020.2975756. 8.2.2, 8.3.2, 8.5, 9.1
- [91] R. Jaumann et al. The high-resolution stereo camera (HRSC) experiment on Mars Express: Instrument aspects and experiment conduct from interplanetary cruise through the nominal mission. *Planetary and Space Science*, 55(7-8):928–952, 2007. 4.2.3
- [92] Jet Propulsion Laboratory, California Institute of Technology. Aquatic Rover Goes for a Drive Under the Ice. <https://www.jpl.nasa.gov/news/news.php?feature=7543>, 2019. Accessed: 2020-12-01. 1.1

- [93] Jerome B. Johnson, Anton V. Kulchitsky, Paul Duvoy, Karl Iagnemma, Carmine Senatore, Raymond E. Arvidson, and Jeffery Moore. Discrete element method simulations of Mars Exploration Rover wheel performance. *Journal of Terramechanics*, 62:31–40, 2015. ISSN 00224898. doi: 10.1016/j.jterra.2015.02.004. 8.2.2
- [94] Sertac Karaman and Emilio Frazzoli. Sampling-based algorithms for optimal motion planning. *International Journal of Robotics Research*, 30(7):846–894, 2011. ISSN 02783649. doi: 10.1177/0278364911406761. 8.2.2
- [95] Alex Kendall and Yarin Gal. What uncertainties do we need in Bayesian deep learning for computer vision? In *Advances in Neural Information Processing Systems*, pages 5575–5585, 2017. 2.2, 5.2, 9.2
- [96] Diederik P. Kingma and Jimmy Ba. Adam: A Method for Stochastic Optimization. *arXiv preprint arXiv:1412.6980*, pages 1–15, 2014. ISSN 09252312. doi: 10.1063/1.4902458. URL <http://arxiv.org/abs/1412.6980>. 5.4.2, 5.5.2
- [97] Diederik P. Kingma and Max Welling. Auto-Encoding Variational Bayes. *arXiv preprint arXiv:1312.6114*, 2013. URL <http://arxiv.org/abs/1312.6114>. 5.2, 5.3.1, 5.4.1, 6.3.1
- [98] Armen Der Kiureghian and Ove Ditlevsen. Aleatory or epistemic? Does it matter? *Structural Safety*, 31(2):105–112, 2009. ISSN 01674730. doi: 10.1016/j.strusafe.2008.06.020. URL <http://dx.doi.org/10.1016/j.strusafe.2008.06.020>. 2.2
- [99] Suhit Kodgule, Alberto Candela, and David Wettergreen. Non-myopic Planetary Exploration Combining in Situ and Remote Measurements. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 536–543, 2019. ISBN 9781728140049. doi: 10.1109/IROS40897.2019.8967769. 6.4, 6.5.1, 6.6.1, 6.7.1, 7.6.2, 8.2.1, 8.4.1, 9.1
- [100] Raymond F. Kokaly, Roger N. Clark, Gregg A. Swayze, K. Eric Livo, Todd M. Hoefen, Neil C. Pearson, Richard A. Wise, William M. Benzel, Heather A. Lowers, Rhonda L. Driscoll, and Anna J. Klein. USGS Spectral Library Version 7. *U.S. Geological Survey Data Series 1035*, 2017. doi:10.3133/ds1035. 6.3.3, 6.5.1
- [101] A. Krause, C. Guestrin, A. Gupta, and J. Kleinberg. Near-optimal sensor placements: maximizing information while minimizing communication cost. *2006 5th International Conference on Information Processing in Sensor Networks*, pages 2–10, 2006. ISSN 14523981. doi: 10.20964/2017.12.03. URL <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=1662434>. 7.4
- [102] Andreas Krause, Ajit Paul Singh, and Carlos Guestrin. Near-Optimal Sensor Placements in Gaussian Processes: Theory, Efficient Algorithms and Empirical Studies. *Journal of Machine Learning Research*, 9:235–284, 2008. doi: 10.1145/1102351.1102385. URL <http://portal.acm.org/citation.cfm?doid=1102351.1102385>. 6.2,

6.4, 6.4, 6.7.1, 8.2.1, 9.1

- [103] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E. Hinton. Handbook of approximation algorithms and metaheuristics. In *Advances in Neural Information Processing Systems*, 2012. ISBN 9781420010749. doi: 10.1201/9781420010749. 5.2
- [104] 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](http://dx.doi.org/10.1016/0034-4257(93)90013-N). URL <http://www.sciencedirect.com/science/article/pii/003442579390013N>. Airborne Imaging Spectrometry. 3.2, 3.2
- [105] 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>. 3.2
- [106] D.A. Kuchler, R.T. Biña, and D. van R. Claasen. Status of High-Technology Remote Sensing for Mapping and Monitoring Coral Reef Environments. In *Sixth International Coral Reef Symposium*, pages 97–101, 1988. 4.3, 6.1
- [107] Sumit Kumar, Wenhao Luo, George Kantor, and Katia Sycara. Active Learning with Gaussian Processes for High Throughput Phenotyping. In *Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems*, pages 2078—2080, Montreal QC, Canada, 2019. International Foundation for Autonomous Agents and Multiagent Systems. URL <http://dl.acm.org/citation.cfm?id=3306127.3332016>. 6.2, 6.4, 9.2
- [108] Balaji Lakshminarayanan, Alexander Pritzel, and Charles Blundell. Simple and scalable predictive uncertainty estimation using deep ensembles. *Advances in Neural Information Processing Systems*, 2017-December(Nips):6403–6414, 2017. ISSN 10495258. 5.2, 9.2
- [109] David A Landgrebe. *Signal Theory Methods in Multispectral Remote Sensing*. Wiley, 2003. 3.2, 3.2
- [110] Quoc V. Le, Alex J. Smola, and Stéphane Canu. Heteroscedastic Gaussian process regression. *Proceedings of the 22nd International Conference on Machine Learning*, pages 489–496, 2005. doi: 10.1145/1102351.1102413. 9.2
- [111] J.K. Lein. *Environmental Sensing: Analytical Techniques for Earth Observation*. Springer, 1st edition, 2012. 3.1, 3.2, 3.2
- [112] Shuai Li, Paul G. Lucey, Ralph E. Milliken, Paul O. Hayne, Elizabeth Fisher, Jean-Pierre Williams, Dana M. Hurley, and Richard C. Elphic. Direct evidence of surface exposed water ice in the lunar polar regions. *Proceedings of the National Academy of Sciences*, 115(36), 2018. ISSN 0027-8424. doi: 10.1073/pnas.1802345115. URL <http://www.>

pnas.org/lookup/doi/10.1073/pnas.1802345115. 3.1

- [113] D.V. Lindley. On the measure of information provided by an experiment. *The Annals of Mathematical Statistics*, 27(4):986–1005, 1956. 2.4, 7.4
- [114] Andrey Malinin and Mark Gales. Predictive Uncertainty Estimation via Prior Networks. In *32nd Conference on Neural Information Processing Systems (NeurIPS)*, 2018. URL <http://arxiv.org/abs/1802.10501>. 5.2
- [115] Stéphane Maritorena, André Morel, and Bernard Gentili. Diffuse reflectance of oceanic shallow waters: Influence of water depth and bottom albedo. *Limnology and Oceanography*, 39(7):1689–1703, 1994. ISSN 19395590. doi: 10.4319/lo.1994.39.7.1689. 4.3
- [116] George Mathew and Igor Mezi. Metrics for ergodicity and design of ergodic dynamics for multi-agent systems. *Physica D: Nonlinear Phenomena*, 240(4-5):432–442, 2011. ISSN 01672789. doi: 10.1016/j.physd.2010.10.010. URL <http://dx.doi.org/10.1016/j.physd.2010.10.010>. 6.7.1, 9.1
- [117] 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. 3.2
- [118] Alfred S. McEwen, Eric M. Eliason, James W. Bergstrom, Nathan T. Bridges, Candice J. Hansen, W. Alan Delamere, John A. Grant, Virginia C. Gulick, Kenneth E. Herkenhoff, Laszlo Keszthelyi, Randolph L. Kirk, Michael T. Mellon, Steven W. Squyres, Nicolas Thomas, and Catherine M. Weitz. Mars reconnaissance orbiter’s high resolution imaging science experiment (HiRISE). *Journal of Geophysical Research E: Planets*, 112(5):1–40, 2007. ISSN 01480227. doi: 10.1029/2005JE002605. 4.2.3
- [119] Andrew McHutchon and Carl Edward Rasmussen. Gaussian Process training with input noise. In *Advances in Neural Information Processing Systems*, pages 1–9, 2011. ISBN 9781618395993. 6.3, 6.3.2
- [120] Dimitrios Milios, Raffaello Camoriano, Pietro Michiardi, Lorenzo Rosasco, and Maurizio Filippone. Dirichlet-based Gaussian processes for large-scale calibrated classification. In *Advances in Neural Information Processing Systems*, pages 6005–6015, 2018. 9.2
- [121] Lauren M. Miller and Todd D. Murphey. Trajectory optimization for continuous ergodic exploration on the motion group SE(2). In *American Control Conference*, pages 4196–4201. IEEE, 2013. ISBN 978-1-4799-0178-4. doi: 10.1109/ACC.2013.6580484. 6.7.1, 8.2.1, 9.1
- [122] Lauren M. Miller, Yonatan Silverman, Malcolm A. MacIver, and Todd D. Murphey. Ergodic Exploration of Distributed Information. *IEEE Transactions on Robotics*, 32(1): 36–52, 2016. doi: 10.1109/TRO.2015.2500441. 8.2.1
- [123] Michihiro Mizuno and Takashi Kubota. A New Path Planning Architecture to Con-

- sider Motion Uncertainty in Natural Environment. In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 2182–2188, 2020. ISBN 9781728173955. 8.2.2
- [124] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A. Rusu, Joel Veness, Marc G. Bellemare, Alex Graves, Martin Riedmiller, Andreas K. Fidjeland, Georg Ostrovski, Stig Petersen, Charles Beattie, Amir Sadik, Ioannis Antonoglou, Helen King, Dharmashan Kumaran, Daan Wierstra, Shane Legg, and Demis Hassabis. Human-level control through deep reinforcement learning. *Nature*, 518(7540):529–533, 2015. ISSN 14764687. doi: 10.1038/nature14236. URL <http://dx.doi.org/10.1038/nature14236>. 5.2
- [125] Fredrik Moberg and Carl Folke. Ecological goods and services of coral reef ecosystems. *Ecological Economics*, 29(2):215–233, 1999. ISSN 09218009. doi: 10.1016/S0921-8009(99)00009-9. 4.3
- [126] José Luis Morales and Jorge Nocedal. Remark on ”algorithm 778: L-BFGS-B: Fortran subroutines for large-scale bound constrained optimization”. *ACM Transactions on Mathematical Software*, 38(1):2–5, 2011. ISSN 00983500. doi: 10.1145/2049662.2049669. 6.3.2
- [127] Peter J. Mumby, William Skirving, Alan E. Strong, John T. Hardy, Ellsworth F. LeDrew, Eric J. Hochberg, Rick P. Stumpf, and Laura T. David. Remote sensing of coral reefs and their physical environment. *Marine Pollution Bulletin*, 48(3-4):219–228, 2004. ISSN 0025326X. doi: 10.1016/j.marpolbul.2003.10.031. 4.3, 6.1
- [128] Scott Murchie, R. Arvidson, P. Bedini, K. Beisser, J. P. Bibring, J. Bishop, J. Boldt, P. Cavender, T. Choo, R. T. Clancy, E. H. Darlington, D. Des Marais, R. Espiritu, D. Fort, R. Green, E. Guinness, J. Hayes, C. Hash, K. Heffernan, J. Hemmler, G. Heyler, D. Humm, J. Hutcheson, N. Izenberg, R. Lee, J. Lees, D. Lohr, E. Malaret, T. Martin, J. A. McGovern, P. McGuire, R. Morris, J. Mustard, S. Pelkey, E. Rhodes, M. Robinson, T. Roush, E. Schaefer, G. Seagrave, F. Seelos, P. Silverglate, S. Slavney, M. Smith, W. J. Shyong, K. Strohbehn, H. Taylor, P. Thompson, B. Tossman, M. Wirzburger, and M. Wolff. Compact Connaissance Imaging Spectrometer for Mars (CRISM) on Mars Reconnaissance Orbiter (MRO). *Journal of Geophysical Research E: Planets*, 112(5):1–57, 2007. ISSN 01480227. doi: 10.1029/2006JE002682. 3.1, 4.2.3
- [129] National Academies of Sciences, Engineering, and Medicine. *Thriving on Our Changing Planet: A Decadal Strategy for Earth Observation from Space*. The National Academies Press, Washington, DC, 2018. ISBN 978-0-309-46757-5. doi: 10.17226/24938. 5.6, 9.2
- [130] National Aeronautics and Space Administration. Aquatic Rover Goes for a Drive Under the Ice. <https://www.nasa.gov/viper>, 2020. Accessed: 2020-12-01. 1.1
- [131] E. Z. Noe Dobrea, M. E. Banks, A. Candela, R. C. Clark, D. R. Gaylord, A. Hendrix, G. Holsclaw, M. D. Lane, M. Osterloo, T. H. Prettyman, S. Vijayarangan, R. Clegg-

- Watkins, D. Wettergreen, and S. P. Wright. Tools for autonomous rover science: Preparing for analog field exercises. In *52th Lunar and Planetary Science Conference*, 2021. 4.2.2
- [132] Masahiro Ono, Thoams J. Fuchs, Amanda Steffy, Mark Maimone, and Jeng Yen. Risk-aware planetary rover operation: Autonomous terrain classification and path planning. In *IEEE Aerospace Conference Proceedings*, pages 1–10. IEEE, 2015. ISBN 9781479953790. doi: 10.1109/AERO.2015.7119022. 8.2.2
- [133] Masahiro Ono, Brandon Rothrock, Eduardo Almeida, Adnan Ansar, Richard Otero, Andres Huertas, and Matthew Heverly. Data-driven surface traversability analysis for Mars 2020 landing site selection. In *IEEE Aerospace Conference Proceedings*, pages 1–12. IEEE, 2016. ISBN 9781467376761. doi: 10.1109/AERO.2016.7500597. 8.2.2
- [134] Masahiro Ono, Matthew Heverly, Brandon Rothrock, Eduardo Almeida, Fred Calef, Tariq Soliman, Nathan Williams, Hallie Gengl, Takuto Ishimatsu, Austin Nicholas, Erisa Stilley, Kyohei Otsu, Robert Lange, and Sarah M. Milkovich. Mars 2020 Site-Specific Mission Performance Analysis: Part 2. Surface Traversability. In *2018 AIAA SPACE and Astronautics Forum and Exposition*, pages 1–14, 2018. doi: 10.2514/6.2018-5419. URL <https://arc.aiaa.org/doi/abs/10.2514/6.2018-5419>. 8.2.2, 8.3.2
- [135] John M. Pandolfi, Sean R. Connolly, Dustin J. Marshall, and Anne L. Cohen. Projecting Coral Reef Futures Under Global Warming and Ocean Acidification. *Science*, 333(6041): 418–422, 2011. ISSN 00368075. doi: 10.1126/science.1204794. 4.3
- [136] Mario Parente, Ian Gemp, and Ishan Durugkar. Unmixing in the presence of nuisances with deep generative models. In *International Geoscience and Remote Sensing Symposium (IGARSS)*, pages 5189–5192, 2017. ISBN 9781509049516. doi: 10.1109/IGARSS.2017.8128170. 5.2
- [137] L. Pedersen, M. Wagner, D. Apostolopoulos, and W. Whittaker. Autonomous robotic meteorite identification in Antarctica. In *IEEE International Conference on Robotics and Automation*, volume 4, pages 4158–4165, 2001. ISBN 0780365763. doi: 10.1109/ROBOT.2001.933268. 7.2
- [138] K. Petersen. *Ergodic Theory*. Cambridge University Press, 1984. 6.7.1
- [139] 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. 3.2
- [140] Morgan Quigley, Brian Gerkey, Ken Conley, Josh Faust, Tully Foote, Jeremy Leibs, Eric Berger, Rob Wheeler, and Andrew Ng. ROS: an open-source Robot Operating System. In *ICRA Workshop on Open Source Software*, 2009. 4.1
- [141] Dushyant Rao, Mark De Deuge, Navid Nourani-Vatani, Stefan B. Williams, and Oscar

- Pizarro. Multimodal learning and inference from visual and remotely sensed data. *International Journal of Robotics Research*, 36(1):24–43, 2017. ISSN 17413176. doi: 10.1177/0278364916679892. 6.2
- [142] Carl Edward Rasmussen and Christopher K. I. Williams. *Gaussian Processes for Machine Learning*. The MIT Press, Cambridge, Massachusetts, 2006. ISBN 026218253X. doi: 10.1142/S0129065704001899. URL www.GaussianProcess.org/gpml. 6.2, 6.3.2, 6.3.2, 9.1
- [143] 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.
- [144] Danilo Jimenez Rezende and Shakir Mohamed. Variational Inference with Normalizing Flows. In *Proceedings of the 32nd International Conference on Machine Learning*, volume 37, pages 1530—1538. JMLR.org, 2015. URL <http://dl.acm.org/citation.cfm?id=3045118.3045281>. 5.2
- [145] Danilo Jimenez Rezende, Shakir Mohamed, and Daan Wierstra. Stochastic Backpropagation and Approximate Inference in Deep Generative Models. In *Proceedings of the 31st International Conference on Machine Learning (ICML)*, pages 1278—1286, 2014. URL <http://arxiv.org/abs/1401.4082>. 5.2, 5.3.1, 5.4.1, 6.3.1
- [146] F. J. Richards. A Flexible Growth Function for Empirical Use. *Journal of Experimental Botany*, 10(29):290–300, 1959. ISSN 00220957. doi: 10.1093/jxb/10.2.290. 7.4
- [147] 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.
- [148] Gene Roddenberry. *Star Trek: The Next Generation*. Paramount Domestic Television, 1987. 9.3
- [149] Brandon Rothrock, Jeremie Papon, Ryan Kennedy, Masahiro Ono, Matthew Heverly, and Chris Cunningham. SPOC: Deep Learning-based Terrain Classification for Mars Rover Missions. In *AIAA SPACE*, 2016. doi: 10.2514/6.2016-5539. 8.2.2
- [150] Silvia Guillén Ruiz, Luis V. Calderita, Alejandro Hidalgo-Paniagua, and Juan P. Bandera Rubio. Measuring smoothness as a factor for efficient and socially accepted robot motion. *Sensors (Switzerland)*, 20(23):1–20, 2020. ISSN 14248220. doi: 10.3390/s20236822. 6.7.1
- [151] David E. Rumelhart, Geoffrey E. Hinton, and Ronald J. Williams. Learning Representations by Back-Propagating Errors. *Nature*, 323(6088):533–536, 1986. doi: 10.1038/323533a0. 5.2
- [152] Kenneth J. 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. ISSN 10618600. doi: 10.1198/1061860032012. 7.4
- [153] Jackson Shields, Oscar Pizarro, and Stefan B. Williams. Towards Adaptive Benthic Habitat Mapping. In *IEEE International Conference on Robotics and Automation*, pages 9263–9270, 2020. ISBN 9781728173955. doi: 10.1109/ICRA40945.2020.9196811. 6.2
- [154] Amarjeet Singh, Andreas Krause, and William J. Kaiser. Nonmyopic adaptive informative path planning for multiple robots. *IJCAI International Joint Conference on Artificial Intelligence*, pages 1843–1850, 2009. ISSN 10450823. doi: 10.1002/rcs. 7.4, 7.5.2, 8.2.1
- [155] Geoffrey M. Smith and Edward J. Milton. The use of the empirical line method to calibrate remotely sensed data to reflectance. *International Journal of Remote Sensing*, 20(13): 2653–2662, 1999. ISSN 13665901. doi: 10.1080/014311699211994. 4.2.3
- [156] Noah A. Smith and Roy W. Tromble. Sampling uniformly from the Unit Simplex. *Johns Hopkins University, Tech. Rep.*, 29, 2004. 7.4
- [157] S.V. Smith and R.W. Buddemeier. Global Change and Coral Reef Ecosystems. *Annual Review of Ecology and Systematics*, 23(1992):89–118, 1992. 4.3
- [158] Kathryn M. Stack, Nathan R. Williams, Fred Calef, Vivian Z. Sun, Kenneth H. Williford, Kenneth A. Farley, Sigurd Eide, David Flannery, Cory Hughes, Samantha R. Jacob, Linda C. Kah, Forrest Meyen, Antonio Molina, Cathy Quantin Nataf, Melissa Rice, Patrick Russell, Eva Scheller, Christina H. Seeger, William J. Abbey, Jacob B. Adler, Hans Amundsen, Ryan B. Anderson, Stanley M. Angel, Gorka Arana, James Atkins, Megan Barrington, Tor Berger, Rose Borden, Beau Boring, Adrian Brown, Brandi L. Carrier, Pamela Conrad, Henning Dypvik, Sarah A. Fagents, Zachary E. Gallegos, Brad Garczynski, Keenan Golder, Felipe Gomez, Yulia Goreva, Sanjeev Gupta, Svein Erik Hamran, Taryn Hicks, Eric D. Hinterman, Briony N. Horgan, Joel Hurowitz, Jeffrey R. Johnson, Jeremie Lasue, Rachel E. Kronyak, Yang Liu, Juan Manuel Madariaga, Nicolas Mangold, John McClean, Noah Miklusicak, Daniel Nunes, Corrine Rojas, Kirby Runyon, Nicole Schmitz, Noel Scudder, Emily Shaver, Jason SooHoo, Russell Spaulding, Evan Stanish, Leslie K. Tamppari, Michael M. Tice, Nathalie Turenne, Peter A. Willis, and R. Aileen Yingst. Photogeologic Map of the Perseverance Rover Field Site in Jezero Crater Constructed by the Mars 2020 Science Team. *Space Science Reviews*, 216(127), 2020. ISSN 15729672. doi: 10.1007/s11214-020-00739-x. URL <http://dx.doi.org/10.1007/s11214-020-00739-x>. 4.2.3, 8.3.2, 8.2, 8.5
- [159] Paul Stankiewicz, Yew T. Tan, and Marin Kobilarov. Adaptive sampling with an autonomous underwater vehicle in static marine environments. *Journal of Field Robotics*, 38(4):572–597, 2021. ISSN 15564967. doi: 10.1002/rob.22005. 6.2, 9.2
- [160] Anthony Stentz. Optimal and Efficient Path Planning for Partially-Known Environments. In *IEEE International Conference on Robotics and Automation (ICRA)*, volume 4, pages

3310–3317, 1994. doi: 10.1109/ROBOT.1994.351061. 8.2.2

- [161] 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, Roger E Stoffregen, and Roger P Ashley. Mapping Advanced Argillic Alteration at Cuprite, Nevada, Using Imaging Spectroscopy. *Economic Geology*, 109(5):1179–1221, 2014. doi: 10.2113/econgeo.109.5.1179. 4.2.1, 5.4.2, 7.5.1, 7.4
- [162] David R. Thompson. *Intelligent Mapping for Autonomous Robotic Survey*. PhD thesis, Robotics Institute, Carnegie Mellon University, Pittsburgh, PA, August 2008. 1.2, 1.4, 6.2, 6.3.2, 6.3.2, 9.1
- [163] David R. Thompson, Lukas Mandrake, Martha S. Gilmore, and Rebecca Castaño. Super-pixel Endmember Detection. *IEEE Transactions on Geoscience and Remote Sensing*, 48(11):4023–4033, 2010. 3.2
- [164] David R. Thompson, David S. Wettergreen, and Francisco J. Calderón P. Autonomous Science during Large-scale Robotic Survey. *Journal of Field Robotics*, 28(4):542–564, 2011. ISSN 1556-4967. doi: 10.1002/rob. URL <http://dx.doi.org/10.1002/rob.20391>. 7.1
- [165] 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 hyspirc preparatory campaign. *Remote Sensing of Environment*, 167:64–77, 2015. ISSN 0034-4257. doi: <https://doi.org/10.1016/j.rse.2015.02.010>. URL <https://www.sciencedirect.com/science/article/pii/S0034425715000607>. Special Issue on the Hyperspectral Infrared Imager (HyspIRI). 4.2.1
- [166] David R. Thompson, David Wettergreen, Greydon Foil, Michael Furlong, and Anatha Ravi Kiran. Spatio-spectral exploration combining in situ and remote measurements. In *Twenty-Ninth AAAI Conference on Artificial Intelligence*, pages 3679–3685, 2015. 6.2, 7.5.1, 8.2.1
- [167] David R. Thompson, Eric J. Hochberg, Gregory P. Asner, Robert O. Green, David E. Knapp, Bo Cai Gao, Rodrigo Garcia, Michelle Gierach, Zhongping Lee, Stephane Martorena, and Ronald Fick. Airborne mapping of benthic reflectance spectra with Bayesian linear mixtures. *Remote Sensing of Environment*, 200(August):18–30, 2017. ISSN 00344257. doi: 10.1016/j.rse.2017.07.030. URL <http://dx.doi.org/10.1016/j.rse.2017.07.030>. 4.3, 4.9, 4.3, 4.10, 6.27a
- [168] David R. Thompson, Alberto Candela, David S. Wettergreen, Eldar Noe Dobrea, Gregg A Swayze, Roger N. Clark, and Rebecca Greenberger. Spatial Spectroscopic Models for Remote Exploration. *Astrobiology*, 18(7):934–954, 2018. ISSN 0261-2194. doi: 10.1089/ast.2017.1782. 4.2.1, 6.2, 8.2.2

- [169] David R. Thompson, Amy Braverman, Philip G. Brodrick, Alberto Candela, Nimrod Carmon, Roger N. Clark, David Connelly, Robert O. Green, Raymond F. Kokaly, Longlei Li, Natalie Mahowald, Ronald L. Miller, Gregory S. Okin, Thomas H. Painter, Gregg A. Swayze, Michael Turmon, Jouni Susilouto, and David S. Wettergreen. Quantifying uncertainty for remote spectroscopy of surface composition. *Remote Sensing of Environment*, 247:111898, 2020. ISSN 00344257. doi: 10.1016/j.rse.2020.111898. URL <https://doi.org/10.1016/j.rse.2020.111898>. 6.5
- [170] Sebastian Thrun, Wolfram Burgard, and Dieter Fox. *Probabilistic Robotics*. The MIT Press, Cambridge, Massachusetts, 1st edition, 2005. ISBN 978-0262201629. 6.3.4, 6.5
- [171] T.N. Tran, R. Wehrens, and L.M.C Buydens. Clustering multispectral images: a tutorial. *Chemometrics and Intelligent Laboratory Systems*, 77:3–17, 2005. 3.2
- [172] Aäron van den Oord, Nal Kalchbrenner, Oriol Vinyals, Lasse Espeholt, Alex Graves, and Koray Kavukcuoglu. Conditional Image Generation with PixelCNN Decoders. In *Proceedings of the 30th International Conference on Neural Information Processing Systems*, pages 4797—4805. Curran Associates Inc., 2016. ISBN 978-1-5108-3881-9. URL <http://dl.acm.org/citation.cfm?id=3157382.3157633>. 5.2
- [173] Christina E. Viviano-Beck, Frank P. Seelos, Scott L. Murchie, Eliezer G. Kahn, Kimberley D. Seelos, Howard W. Taylor, Kelly Taylor, Bethany L. Ehlmann, Sandra M. Wiseman, John F. Mustard, and M. Frank Morgan. Revised crism spectral parameters and summary products based on the currently detected mineral diversity on mars. *Journal of Geophysical Research: Planets*, 119(6):1403–1431, 2014. doi: 10.1002/2014JE004627. URL <https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1002/2014JE004627>. 1.2, 4.2.3, 4.2.3, 6.3.3, 8.2.2
- [174] Yannan Wang, Jun Du, Li Rong Dai, and Chin Hui Lee. A maximum likelihood approach to deep neural network based nonlinear spectral mapping for single-channel speech separation. In *Proceedings of the Annual Conference of the International Speech Communication Association, INTERSPEECH*, pages 1178–1182, 2017. doi: 10.21437/Interspeech.2017-830. 5.2, 5.4.1
- [175] David Wettergreen, Nathalie Cabrol, Vijayakumar Baskaran, Francisco Calderón, Stuart Heys, Dominic Jonak, Allan Luders, David Pane, Trey Smith, James Teza, Paul Tompkins, Daniel Villa, Chris Williams, and Michael Wagner. Second Experiments in the Robotic Investigation of Life in the Atacama Desert of Chile. In *International Symposium on Artificial Intelligence, Robotics and Automation in Space (i-SAIRAS)*, 2005. 4.1
- [176] David Wettergreen, Nathalie Cabrol, James Teza, Paul Tompkins, Chris Urmson, Vandí Verma, Michael Wagner, and William Whittaker. First experiments in the robotic investigation of life in the Atacama Desert of Chile. In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 873–878, 2005. ISBN 078038914X. doi:

- [177] Roger C. Wiens, Sylvestre Maurice, Scott H. Robinson, Anthony E. Nelson, Philippe Cais, Pernelle Bernardi, Raymond T. Newell, Sam Clegg, Shiv K. Sharma, Steven Storms, Jonathan Deming, Darrel Beckman, Ann M. Ollila, Olivier Gasnault, Ryan B. Anderson, Yves André, S. Michael Angel, Gorka Arana, Elizabeth Auden, Pierre Beck, Joseph Becker, Karim Benzerara, Sylvain Bernard, Olivier Beyssac, Louis Borges, Bruno Bousquet, Kerry Boyd, Michael Caffrey, Jeffrey Carlson, Kepa Castro, Jorden Celis, Baptiste Chide, Kevin Clark, Edward Cloutis, Elizabeth C. Cordoba, Agnes Cousin, Magdalena Dale, Lauren Deflores, Dorothea Delapp, Muriel Deleuze, Matthew Dirmyer, Christophe Donny, Gilles Dromart, M. George Duran, Miles Egan, Joan Ervin, Cecile Fabre, Amaury Fau, Woodward Fischer, Olivier Forni, Thierry Fouchet, Reuben Fresquez, Jens Frydenvang, Denine Gasway, Ivair Gontijo, John Grotzinger, Xavier Jacob, Sophie Jacquinod, Jeffrey R. Johnson, Roberta A. Klisiewicz, James Lake, Nina Lanza, Javier Laserna, Jeremie Lasue, Stéphane Le Mouélic, Carey Legett, Richard Leveille, Eric Lewin, Guillermo Lopez-Reyes, Ralph Lorenz, Eric Lorigny, Steven P. Love, Briana Lucero, Juan Manuel Madariaga, Morten Madsen, Soren Madsen, Nicolas Mangold, Jose Antonio Manrique, J. P. Martinez, Jesus Martinez-Frias, Kevin P. McCabe, Timothy H. McConnochie, Justin M. McGlown, Scott M. McLennan, Nouredine Melikechi, Pierre Yves Meslin, John M. Michel, David Mimoun, Anupam Misra, Gilles Montagnac, Franck Montmessin, Valerie Mousset, Naomi Murdoch, Horton Newsum, Logan A. Ott, Zachary R. Ousnamer, Laurent Pares, Yann Parot, Rafal Pawluczyk, C. Glen Peterson, Paolo Pilleri, Patrick Pinet, Gabriel Pont, Francois Poulet, Cheryl Provost, Benjamin Quertier, Heather Quinn, William Rapin, Jean Michel Reess, Amy H. Regan, Adriana L. Reyes-Newell, Philip J. Romano, Clement Royer, Fernando Rull, Benigno Sandoval, Joseph H. Sarrao, Violaine Sautter, Marcel J. Schoppers, Susanne Schröder, Daniel Seitz, Terra Shepherd, Pablo Sobron, Bruno Dubois, Vishnu Sridhar, Michael J. Toplis, Imanol Torre-Fdez, Ian A. Trettel, Mark Underwood, Andres Valdez, Jacob Valdez, Dawn Venhaus, and Peter Willis. The SuperCam Instrument Suite on the NASA Mars 2020 Rover: Body Unit and Combined System Tests. *Space Science Reviews*, 217(1), 2021. ISSN 15729672. doi: 10.1007/s11214-020-00777-5. URL <http://dx.doi.org/10.1007/s11214-020-00777-5>. 4.2.3
- [178] Peter M. Williams. Using Neural Networks to Model Conditional Multivariate Densities. *Neural Computation*, 8(4):843–854, 1996. ISSN 08997667. doi: 10.1162/neco.1996.8.4.843. 5.2
- [179] Andrew Gordon Wilson and Hannes Nickisch. Kernel interpolation for scalable structured Gaussian processes (KISS-GP). In *32nd International Conference on Machine Learning, ICML 2015*, volume 3, pages 1775–1784, 2015. ISBN 9781510810587. 9.2
- [180] Mark Woods, Andy Shaw, Dave Barnes, Dave Price, Derek Long, and Derek Pullan.

- Autonomous Science for an ExoMars RoverLike Mission. *Journal of Field Robotics*, 26(4):358—390, 2009. doi: <https://doi.org/10.1002/rob.20289>. 1.1, 7.2
- [181] Y. Wu, Y. Burda, R. Salakhutdinov, and R. Grosse. On the Quantitative Analysis of Decoder-Based Generative Models. In *International Conference on Learning Representations*, 2017. 5.2
- [182] Takaki Yamada, Adam Prügel-Bennett, and Blair Thornton. Learning features from geo-referenced seafloor imagery with location guided autoencoders. *Journal of Field Robotics*, 38(1):52–67, 2021. ISSN 15564967. doi: 10.1002/rob.21961. 6.3.1
- [183] Namik Kemal Yilmaz, Constantinos Evangelinos, Pierre F.J. Lermusiaux, and Nicholas 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, 2008. 6.4, 7.4, 8.2.1
- [184] Tingbin Zhang, Guihua Yi, Hongmei Li, Ziyi Wang, Juxing Tang, Kanghui Zhong, Yubin Li, Qin Wang, and Xiaojuan Bie. Integrating data of ASTER and Landsat-8 OLI (AO) for hydrothermal alteration mineral mapping in duolong porphyry cu-au deposit, Tibetan Plateau, China. *Remote Sensing*, 8(11), 2016. ISSN 20724292. doi: 10.3390/rs8110890. 5.1
- [185] Feng Zhou, Raymond E. Arvidson, Keith Bennett, Brian Trease, Randel Lindemann, Paolo Bellutta, Karl Iagnemma, and Carmine Senatore. Simulations of Mars Rover Traverses. *Journal of Field Robotics*, 31(1):141–160, 2014. doi: 10.1002/rob.21483. 8.2.2
- [186] Richard W. Zurek and Suzanne E. Smrekar. An overview of the Mars Reconnaissance Orbiter (MRO) science mission. *Journal of Geophysical Research E: Planets*, 112(5): 1–22, 2007. ISSN 01480227. doi: 10.1029/2006JE002701. 6.1