

# Robotic Rock Classification and Autonomous Exploration

Liam Pedersen



## Acknowledgements

A tremendous debt of gratitude is due to the many people who assisted or contributed in different ways to this research. The Robotic Antarctic Meteorites Search project was a team effort, involving many people, all of who deserve credit for the project's success.

Firstly, thanks go to my thesis committee: Martial Hebert, Red Whittaker, Andrew Moore, Reid Simmons and Peter Cheeseman, for their time and effort.

Martial Hebert, my academic advisor, has carefully nurtured this thesis and steadily kept me on track over the last few years. In spite of a busy schedule, he has always selflessly managed to make time to help make this research happen, write references, and generally smooth the way. All this in spite of not being directly involved with, or receiving direct credit for the successes of the Robotic Antarctic Meteorite Search project.

Dimitrius Apostolopoulos, as project manager for the Robotic Antarctic Meteorite Search project tirelessly and consistently backed my research. Without his leadership and tenacity, the project would not have been completed and this research would not have been possible.

William "Red" Whittaker, as director of the Field Robotics Center, has supported this project and my research throughout. His tenacity and experience at handling funding agencies ensured essential financial support, and his counsel has been invaluable both in the preparation of this document and in furthering this research.

Dr. Ted Roush was happy for me to spend three summers at NASA's Ames Research Center working with him on rock classification and spectroscopy. His continuous generosity and support have helped me tremendously with this research. He spent substantial time explaining fundamental geology and spectroscopy, and made available, at an inconvenience to himself, a marvelous spectrometer with which to study rocks in Antarctica.

Dr. Bill Cassidy's unstinting help and commitment to the cause of the Robotic Antarctic Meteorite Search were invaluable. I first met Bill at the beginning of the project, when I naïvely asked how to recognize meteorites. Bill nodded wisely and, every bit the wise veteran Antarctic explorer, told me "First you find a few, then you know what to look for". Soon followed by "meteorites are where you find them". Thus began a fruitful partnership, during which Bill patiently explained a lot of basic geology to me, loaned me his personal collection of meteorites, identified Antarctic rocks and was a paragon of dignity and humor under extreme conditions on the Antarctic ice sheets.

Dr Gunter Faure of Ohio State University generously made his collection of Antarctic rocks from Elephant Moraine available for this research. Without this assistance it would not have been possible to prepare for the Elephant Moraine expedition.

Dr Marilyn Lindstrom and the staff of the Antarctic meteorite curation facility at NASA's Johnson Space Center gave me invaluable access to the Antarctic meteorite collection. The meteorites are kept in a clean room and are handled according to strict protocols. Lindstrom and others spent many hours with me in the clean room, patiently preparing the meteorites and helping me acquire spectral measurements crucial for this research.

Pascal Lee, Mathew Deans, Kurt Schwehr, Michael Parris, Mark Sibenac, Stewart Moorehead and Alex Foessel spent significant time in the field with me, measuring and collecting rocks on the Antarctic ice sheet. In spite of extreme conditions they always maintained a sense of humor, patience and professionalism. It was a pleasure and a privilege to work with them.

Michael Wagner and Ben Shamah kept the Nomad robot running in the field during the final trials. Mike was primarily responsible for the software operating Nomad, and spent considerable time with me to integrate my software with the vehicle. Ben designed and built the manipulator arm on Nomad, worked ceaselessly to build the mechanical parts and managed the logistics of the final operations in Antarctica. Their dedication and professionalism are impressive.

Gretchen Benedix helped to identify and characterize many of the rock samples that were collected. Ralph Harvey and John Schutt of ANSMET respectively gave me a lot of advice about Antarctica and kept us out of trouble in the field. John is the gold standard to which all meteorite hunters (human or robotic) aspire. His ability to spot an extra-terrestrial rock at a thousand paces is impressive.

Dot Marsh worked behind the scenes in so many ways to smooth things over and grease the wheels of academic progress. Without her help, consideration, cheerfulness and occasional prod in the back to wake me up things would have been very different. Finally, last, but by no means least, I must thank my family and friends for their continued encouragement and support throughout. My mother was always available and understanding. Chris Scarpinatto would periodically drag me away from the computer to make sure I was eating properly, had enough vegetables and could decompress a little. My neighbors Nick and Jo-Anne Redondo would send their children around with food for a certain thin (possibly starving) graduate student next door.

# Contents

Chapter 1.	Introduction.....	6
1.1.	Robotic Antarctic Meteorite Search (RAMS) project.....	7
1.1.1.	Meteorites and the Antarctic environment.....	8
1.2.	Prior work.....	11
1.3.	Contributions.....	12
1.4.	Thesis outline.....	12
1.5.	References.....	12

## Part I.....

15

### The Bayesian Paradigm for Robotic Rock Classification.....

15

Chapter 2.	Robotic Rock Classification.....	16
2.1.	Issues in robotic rock and meteorite classification.....	16
2.2.	Bayes network generative model.....	20
2.2.1.	Notation.....	20
2.2.2.	Generative and diagnostic statistical models.....	21
2.2.3.	Bayes networks.....	23
2.3.	Other classification systems.....	24
2.4.	Conclusions.....	26
2.5.	References.....	27

Chapter 3.	Bayes network classifier topology.....	28
3.1.	Naïve Bayes architecture.....	28
3.2.	Naïve Bayes network with intermediate nodes.....	29
3.3.	Bayes network with weak data dependencies and intermediate node.....	31
3.4.	The rock and meteorite Bayes network classifier.....	32
3.5.	References.....	33

Chapter 4.	Bayes Network Training.....	35
4.1.	Bayesian parameter learning.....	35
4.1.1.	Learning probabilities from uncertain data.....	38
4.2.	Quantization of continuous parameters and the “Curse of dimensionality”.....	39
4.2.1.	Sensor noise models.....	41
4.3.	Laboratory training results.....	41
4.4.	References.....	43

## Part II.....

44

### Sensors and Sensor Deployment.....

44

Chapter 5.	Spectroscopy.....	45
5.1.	Principles of reflectance spectroscopy.....	46
5.2.	Autonomous spectroscopy in the field.....	47
5.3.	Spectrometer noise model.....	50
5.4.	Spectral feature extraction.....	52
5.4.1.	Spectral peak and trough detection.....	52
5.4.2.	Spectral red, green and blue indices.....	56
5.5.	Conclusions.....	57
5.6.	References.....	57

Chapter 6.	Color Image Processing.....	59
6.1.	Object/background segmentation.....	59
6.1.1.	Markov chain segmentation model.....	60
6.2.	Image features extraction.....	64
6.3.	Nomad high-resolution camera.....	70
6.4.	Conclusions.....	74
6.5.	References.....	75

Chapter 7.	Electromagnetic sensors.....	77
------------	------------------------------	----

7.1.	Metal detectors.....	78
7.2.	Magnetic field measurement.....	80
7.2.1.	Magnetic field of meteorites.....	83
7.3.	Atacama robotic sensor deployment.....	84
7.3.1.	Atacama signal processing.....	84
7.3.2.	Atacama results.....	85
7.3.3.	Atacama lessons.....	90
7.4.	References.....	91
<b>Part III</b> .....		92
Experimental Results.....		92
Chapter 8.	Antarctica 2000 Nomad Field Trials.....	93
8.1.	First autonomous robotic meteorite find.....	95
8.2.	Overall classification performance.....	102
8.3.	Spectrometer performance.....	118
8.4.	Camera performance.....	119
8.5.	References.....	120
<b>Part IV</b> .....		122
Extensions of Bayes Network Generative Model.....		122
Chapter 9.	Autonomous Profiling of Unknown Environments.....	123
9.1.	Introduction.....	123
9.2.	Robotic learning of the environment.....	124
9.2.1.	Representing the rock priors over an area.....	125
9.2.2.	Pseudo Bayes network generative model.....	126
9.2.3.	Geographical models.....	126
9.2.4.	Properties of $M_k(\bar{\theta}_y; x, y)$ .....	127
9.2.5.	Learning from uncertain observations.....	130
9.2.6.	Multiple rock finds.....	130
9.3.	Experimental results.....	131
9.4.	Conclusions.....	136
9.5.	References.....	137
Chapter 10.	Active Sensor Selection.....	139
10.1.	Information gain as measure of sensor utility.....	140
10.1.1.	Computing Information Gain for sensor selection.....	142
10.1.2.	Monte-Carlo computation of information gain.....	143
10.1.3.	Experimental results.....	144
10.1.4.	Problems with Entropy and Information Gain.....	145
10.1.5.	“Re-centered” entropy.....	146
10.1.6.	Experimental results with re-centered entropy.....	148
10.2.	Conclusions.....	150
10.3.	References.....	150
Chapter 11.	Accomplishment and conclusions.....	152
Lessons – “Perspectives and Wisdom”.....		154
Other applications and extensions.....		156
References.....		158
Master Reference.....		159

## Introduction

*This thesis is about the application of Bayesian statistics to the problem of autonomous geological exploration with a robotic vehicle. It concentrates on the sub-problem of classifying rock types while addressing the issues associated with operating onboard a mobile robot, and argues that the Bayesian statistical paradigm, using a Bayes network generative statistical model of the geological environment, is particularly suited to this task. This paradigm is extended in a natural way to solve the more general robotic problems of autonomously profiling an area and allocating scarce sensor resources.*

Modern unmanned planetary rovers suffer the serious shortcomings from their inability to autonomously interpret the scientific data they are sent to acquire, and make real-time decisions on how to explore their environment. Because of bandwidth and storage limitations, they return only a fraction of the data they are capable of acquiring. Nor can they respond to observations that require immediate follow up to be scientifically useful.

The typical operating mode for a robotic explorer, such as the Mars Pathfinder and Sojourner duo, is for mission control to upload a simple set of instructions for the robot, usually of the form “go to such and such a rock and take a close-up picture”. The robot attempts to execute these instructions and acquire the desired sensor data. When communications are again established with mission control, the robot will signal the success or failure on the assigned task, and download the sensor data it acquired. Scientists will ponder this data and, in conjunction with the engineers running the robot, issue further commands to the robot. Operating in this manner, the robot will not investigate more than one rock per command cycle. In the case of Mars, communications are possible only once every 25 hours, for most of which Sojourner was idle. A more “scientifically intelligent” robot might investigate many rocks and return a summary to mission control, along with representative sensor data.

Future missions, such as a robot to search for life in the oceans believed to lie under the ice cap of Europa, one of the moons of Jupiter, are even more constrained. Communications will be extremely limited and the vehicle will have to conduct investigations in an unknown and dynamic environment. The robot will need to autonomously recognize scientifically interesting occurrences and conduct follow up observations if it deems them necessary. In this case, the ability for the vehicle to conduct an autonomous scientific investigation is not an option but a necessity.

A first step in the development of an autonomous science capability for a robotic explorer is a system allowing it to autonomously recognize a restricted class of scientifically interesting objects (such as rocks, meteorites or fossils) and survey an area, learning what objects are present and permitting the identification of anomalies.

This thesis proposes a Bayesian statistical framework for the autonomous classification of interesting objects from onboard a mobile robotic vehicle. Because of a focus on geological exploration, the application is the classification of rocks. However, there are robotic considerations that take this beyond standard classification techniques and make it applicable to more general tasks involving the classification of objects from a vehicle in a poorly characterized environment.

Major considerations are the need to use of multiple sensors and the ability of a robotic vehicle to acquire data from different locations. Needless sensor use must be curtailed if possible, such as when an object is

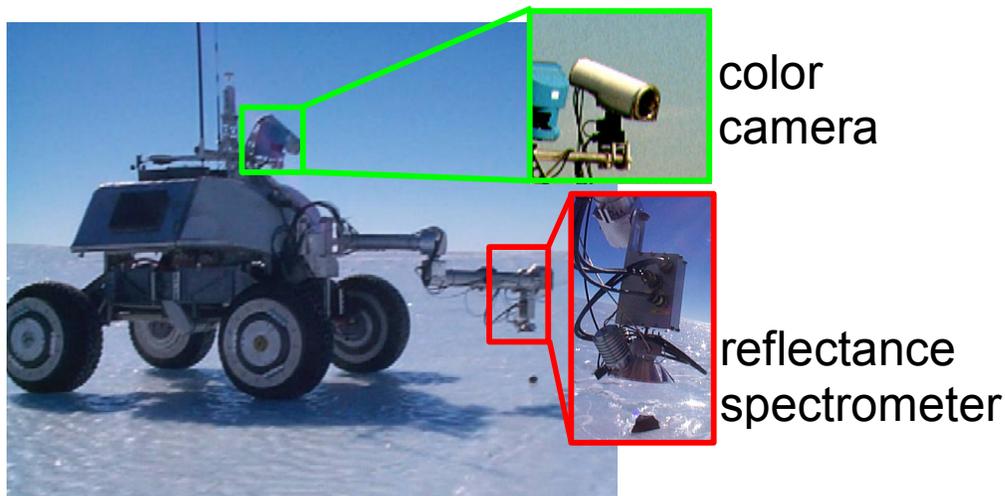
sufficiently well identified given sensor data acquired so far. Furthermore, by investigating rocks in many locations, the robot has the opportunity to profile the environment. Different rock samples are statistically dependent on each other. These dependencies can be exploited to substantially improve classification accuracy.

The robotic rock classification system developed for this thesis is based upon a Bayes network generative model of the statistical relationships between select rock classes and sensor data. It addresses the issues of classifying rocks and working from a robot within a principled statistical framework, and extends in a natural manner to autonomously profiling the geology of an area, exploiting the statistical dependencies between rock samples and making decisions about which of the available onboard sensors to use.

The classification system has been implemented onboard the Nomad robot developed at Carnegie Mellon University, and applied to the task of recognizing meteorites amongst terrestrial rocks in Antarctica. In January, 2000 A.D. Nomad was deployed to Antarctica where it made the first autonomous robotic identification of a meteorite.

### 1.1. Robotic Antarctic Meteorite Search (RAMS) project

Carnegie Mellon University has developed a robot vehicle, Nomad (Figure 1.1), to autonomously search Antarctica for meteorites [1.1]. The eventual goal beyond Antarctic meteorite search is to develop technologies for extended robotic exploration of remote areas, including planetary surfaces. These include the autonomous onboard interpretation of scientific data [1.2][1.4] so that a robot may efficiently explore an area without having to frequently contact mission control. Doing so allows more data to be obtained than can be returned, and reduces the number of personnel needed to monitor the mission. Autonomous rock classification, such as meteorite search, falls under this rubric and faces many similar robotic problems.



**Figure 1.1** *The Nomad robot examining potential meteorite with color camera and visible light spectrometer.*

The meteorite identification system onboard Nomad consists the Bayes network rock classifier of this thesis, appropriate sensors for identifying rocks, especially extra-terrestrial ones (meteorite), and a planning system to integrate these with the rest of the robot.

The task of the onboard rock classifier is to determine the likelihood of a rock sample being a meteorite given the available sensor data. It must accept data from various sensors, not necessarily all at once. The data is noisy and may be systematically corrupted by both the environment and the manner in which the robot takes measurements. Furthermore, the classifier must coincide with robot operations, operating in real time and informing the planning system of the classification results and whether more data is needed.

The foremost problem in classification is to determine features that can distinguish different classes of objects. Appropriate sensors must then be chosen to observe these features. However, when operating from a mobile robot there are severe constraints on what sensors can be used. They must be rugged enough to operate in the robot environment and conform to the robot payload mass, power and volume constraints. More importantly, they must allow autonomous operation, or at least minimal human intervention, in the field. This excludes complex sensors that require careful preparation of rock samples or excessive deployment precision.

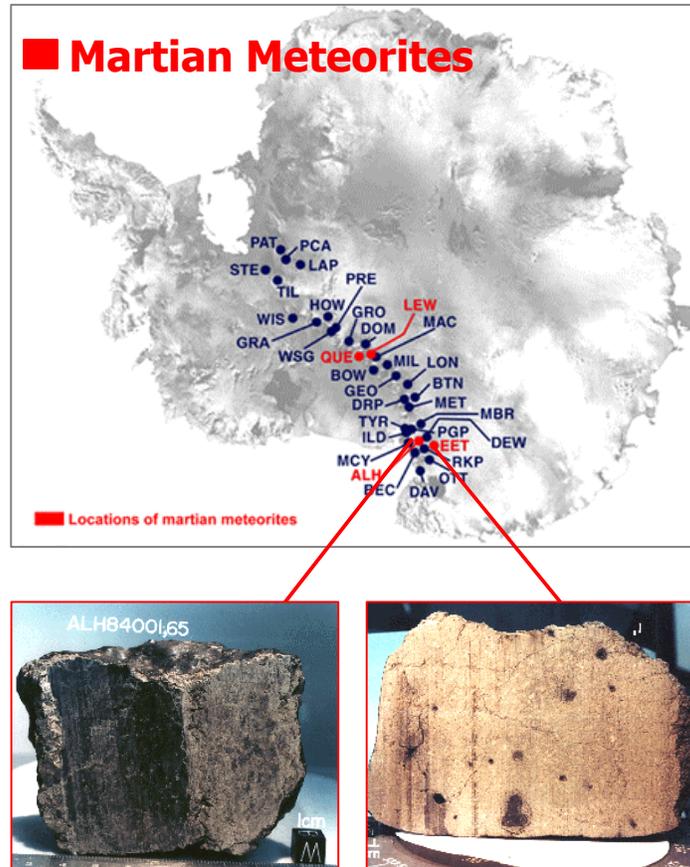
Several sensor modalities for discriminating meteorites from rocks have been explored, including metal detectors and magnetometers [1.3], visible and near IR reflectance spectroscopy and high-resolution color imagery [1.4][1.5][1.6]. These are good for reasonable discrimination, and are rugged enough for use in Antarctica.

Nomad is equipped with a color camera on a pan-tilt platform and an arm mounted reflectance spectrometer (Figure 1.1). The color camera is both to search the ice sheet for rocks and to provide close-up images of them for initial classification. The spectrometer is used for final verification. A metal detector was originally installed but was subsequently damaged.

Nomad's planning system, referred to as the *Science Autonomy System* (SAS) [1.7] interfaces with the robot's sensors, the classifier and navigation systems. It is responsible for acquiring needed scientific data, controlling sensors and routing data to the classifier

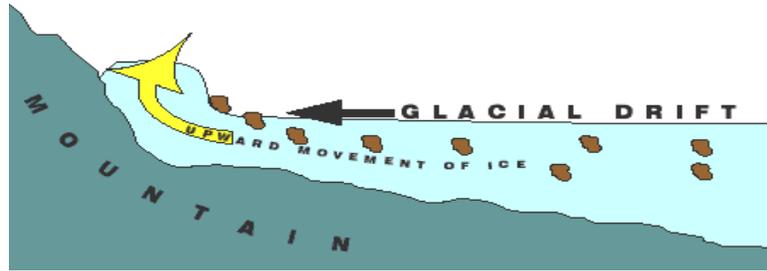
### **1.1.1. Meteorites and the Antarctic environment**

Antarctica is an excellent place to find meteorites (Figure 1.2). Although they fall uniformly over the Earth, those that fall in Antarctica are preserved for longer and easier to find. The pristine, dry and cold Antarctic environment ensures that meteorites deposited there survive for thousands of years and remain relatively uncontaminated, maximizing their scientific value.



**Figure 1.2** Locations in Antarctica where meteorites are abundant. To date, most meteorites ever collected have been found in Antarctica. The cold pristine environment preserves them for thousands of years and the movement of the ice sheets concentrates them in certain area where they can be more easily found. Of particular interest are meteorites originally from Mars and the Moon. One Martian meteorite, ALH84001, is postulated to contain signs of fossilized life.

Most of Antarctica is completely covered by vast ice sheets that slowly glaciare outward from the continental interior until they reach the coast and break off as icebergs. Mountains or other rock formations may obstruct the flow of ice, forcing the ice sheet upwards (Figure 1.3). As the it is pushed up, the relentless katabatic winds of Antarctica ablate it away, stranding buried objects, including meteorites, on the surface. This way, meteorites that have landed over a large area of ice over thousands of years get concentrated in areas adjacent to the Trans-Antarctic mountain range.



**Figure 1.3** Mechanism by which meteorites are believed to be concentrated in certain areas of Antarctica. Meteorites land and are eventually buried in the vast continental ice sheets. As these glaciers move outwards they may push against mountains, forcing the ice upwards, where it is ablated away by the wind, stranding meteorites and other buried rocks on the ice surface.

Meteorites are of interest to space scientists because, with the exception of the Apollo lunar samples, they are the sole source of extraterrestrial material and provide a window on the early evolution of the solar system. The identification of Martian and lunar meteorite samples, and the controversial evidence of fossil bacteria in the former underscore the importance of systematically retrieving as many samples as possible.

Currently, Antarctic meteorite samples are collected by human searchers, either on foot, or on snowmobiles, who systematically search an area and retrieve samples according to strict protocols [1.8]. In certain blue ice fields there are no terrestrial rocks whatsoever. The only rocks visible there arrived from the heavens. Other meteorite stranding surfaces occur around moraines, where the ice flow also brings terrestrial rocks to the surface. In these cases searching is harder as the searchers must also examine many terrestrial rocks (Figure 1.4).

Human teams have searched Antarctica for meteorites only for the last 20 years, yet have recovered more meteorites there than from anywhere else throughout history. Furthermore, since access is limited, Antarctica is the only place where systematic and coordinated meteorite searches are possible. This is important for gaining accurate knowledge of the abundances of the different extraterrestrial objects impacting Earth. Elsewhere on Earth, many meteorites are collected by many private individuals and sold to collectors, removing these valuable specimens from circulation and compromising their scientific value.



**Figure 1.4** *Human search for Antarctic meteorites. In some blue ice fields (top) the only rocks around are those that fell from the sky, so searchers can comb the area on snowmobiles. In moraines, places where the ice flow causes rocks to be brought up to the surface, there are many terrestrial rocks complicating the search for meteorites (bottom).*

## 1.2. Prior work

There is *no* prior work explicitly dealing with the robotic classification of terrestrial rocks and meteorites. The Carnegie Mellon University Robotic Antarctic Meteorite Search project was the first of its kind. Only recently have robot vehicles become sufficiently capable that such a task could even be considered.

The concept of onboard interpretation of scientific data by a robot is also relatively recent. The Mars Pathfinder mission in 1997 highlighted the need for this type of technology, which showed the limitations of the traditional approach to exploring remote environments by returning all scientific data back to Earth for interpretation. [1.2] have been developing tools for the basic interpretation of geological data, such as finding sedimentary layers in images and identifying carbonates.

There is significant interest amongst the geological community in the automated classification of reflectance spectra from different rocks and minerals [1.6][1.9][1.10][1.11][1.12][1.13][1.14][1.15][1.16][1.17][1.18]. Most work deals with the offline classification of laboratory spectra from

pure minerals in powder form. They do not deal with the robotic issues, or address the difficulties of operating in the field.

Classification and Bayesian statistics are mature fields, with extensive literature. This thesis does not purport to advance the theory of classification or Bayesian statistics. Rather, it is about their application to a *robotic* problem that is the concern here.

### **1.3. Contributions**

The applicability of this work goes beyond robotic meteorite search and rock classification. It establishes a framework for recognizing objects and interpreting the environment from a mobile robot. Landmine clearance, agricultural soil profiling, airborne search and rescue, and mineral prospecting are all applications that might be done with a mobile robot, or team of robots, equipped with multiple sensors and sent out to search, patrol or explore an area. Their task is to recognize objects of interest in poorly characterized areas whilst minimizing sensor use and extracting the maximum information from each measurement.

### **1.4. Thesis outline**

There are five parts to this thesis. Part I (chapters 2,3 and 4) deals with the choice of classification system, how it was built and trained from example data.

Part II (chapters 5, 6, and 7) describe the sensors used on Nomad for meteorite and rock classification, how they are robotically deployed in the field and how feature vectors suitable for classifying are extracted from them.

Part III (chapter 8) is the experimental results from the deployment of Nomad in Antarctica in January 2000, when it found its first meteorite.

Part IV (chapters 9 and 10) show the extension of the Bayes network generative model classifier architecture to the problems of autonomously profiling the environment and active sensor selection.

### **1.5. References**

- [1.1] Apostolopoulos, D., Wagner, M., Whittaker, W., "Technology and Field Demonstration Results in the Robotic Search for Antarctic Meteorites", *Field and Service Robotics*, Pittsburgh, USA, 1999.
- [1.2] Roush, T., V. Gulick, V., Morris, R., Gazis, P., Benedix, G., Glymour, C., Ramsey, J., Pedersen, L., Ruzon, M., Buntine, W., and Oliver, J., "Autonomous Science Decision Making for Mars Sample Return", in *Lunar and Planetary Science Conference XXX*, Houston, TX, 1999.
- [1.3] Pedersen, L., "Robotic Deployment of Electro-Magnetic Sensors for Meteorite Search", *IEEE International Conference on Robotics and Automation*, Leuven, Belgium, 1998.
- [1.4] Pedersen, L., Apostolopoulos, D., Whittaker, W. "Bayes Networks on Ice: Robotic Search for Antarctic Meteorites", *Neural Information Processing Symposium*, Denver, Colorado, 2000.

- [1.5] Pedersen, L., Apostolopoulos, D., Whittaker, W., Roush, T. and G. Benedix “Sensing and data classification for robotic meteorite search” Proceedings of *SPIE Photonics East* Conference. Boston, USA, 1998.
- [1.6] Pedersen, L., Wagner, M., Apostolopoulos, D., Whittaker, W., “Autonomous Robotic Meteorite Identification in Antarctica”, in *IEEE International Conference on Robotics and Automation*, Seoul, Korea, 2001.
- [1.7] Wagner, M., Apostolopoulos, D., Shillcutt, K., Shamah, B., Whittaker, W., “The Science Autonomy System of the Nomad Robot”, in *IEEE International Conference on Robotics and Automation*, Seoul, Korea, 2001.
- [1.8] Cassidy, W., Harvey, R., Schutt, J., Delisle, G., and Yanai, K., “The Meteorite Collection Sites of Antarctica”, *Meteoritics*, 27(5), pp 490 – 525, 1992.
- [1.9] Glymour, C., Ramsey, J., Roush, T., Scheines, R., Richardson, T., Bailer, H., “Automated Identification of Mineral Composition from Reflectance Spectra”, *Draft interim report*, July 24, 1998.
- [1.10] Glymour, C., Ramsey, J., Roush, T., “Comparison of Human and Machine Identification of Mineral Classes from Reflection Spectra”, *Draft interim report*, January, 1999.
- [1.11] G. Healey, D. Slater, “Image Understanding Research at UC Irvine: Automated Invariant Recognition in Hyperspectral Imagery”.
- [1.12] Clark, R.N. “Spectroscopy and Principles of Spectroscopy”, Manual of Remote Sensing, A. Rencz, Editor, John Wiley and Sons, Inc. 1999.
- [1.13] Clark, R.N. and Swayze, G.A., “Mapping Minerals, Amorphous Materials, Environmental Materials, Vegetation, Water, Ice and Snow, and Other Materials: The USGS Tricorder Algorithm”, Summaries of the Fifth Annual JPL Airborne Earth Science Workshop, January 23- 26, R.O. Green, Ed., JPL Publication 95-1, p. 39-40, 1995.
- [1.14] Clark, R.N., A.J. Gallagher, and G.A. Swayze, “Material Absorption Band Depth Mapping of Imaging Spectrometer Data Using a Complete Band Shape Least-Squares Fit with Library Reference Spectra”, Proceedings of the Second Airborne Visible/Infrared Imaging Spectrometer (AVIRIS) Workshop. JPL Publication 90-54, 176-186, 1990.
- [1.15] Clark, R.N., G.A. Swayze, A. Gallagher, N. Gorelick, and F. Kruse, “Mapping with Imaging Spectrometer Data Using the Complete Band Shape Least-Squares Algorithm Simultaneously Fit to Multiple Spectral Features from Multiple Materials”, Proceedings of the Third Airborne Visible/Infrared Imaging Spectrometer (AVIRIS) Workshop, JPL Publication 91-28, 2-3, 1991.
- [1.16] Clark, R.N., G.A. Swayze, C. Koch, A. Gallagher, and C. Ager, “Mapping Vegetation Types with the Multiple Spectral Feature Mapping Algorithm in both Emission and Absorption”, Summaries of the Third Annual JPL Airborne Geosciences Workshop, Volume 1: AVIRIS Workshop. JPL Publication 92-14, 60-62, 1992.
- [1.17] Clark, R.N., T.V.V. King, and N.S. Gorelick, “Automatic Continuum Analysis of Reflectance Spectra”, Proceedings of the Third Airborne Imaging Spectrometer Data Analysis Workshop, JPL Publication 87-30, pp 138-142, 1987.

[1.18] Clark, R.N., T.V.V. King, M. Klejwa, G. Swayze, and N. Vergo, "High Spectral Resolution Reflectance Spectroscopy of Minerals", *Journal Geophysical Research*, 95, pp12653-12680, 1990.

# **Part I**

## **The Bayesian Paradigm for Robotic Rock Classification**

## Robotic Rock Classification

In this chapter, we will justify the choice of a Bayes network generative statistical model for robotic rock classification, with an emphasis where appropriate on meteorite recognition. Autonomously classifying rocks poses several difficulties that take it beyond standard classification problems. Working from onboard a mobile robotic vehicle adds further uniquely robotic constraints and capabilities. Many rock types are inherently hard to distinguish, and no single sensor is sufficient by itself for the task. Therefore, information from multiple sensors must be combined. However, there is a cost associated with getting sensor data. The robot must maneuver to deploy instruments to a rock sample, taking both time and energy. It is therefore incumbent that unnecessary sensor use be curtailed.

A Bayes network generative statistical model of the relationships between meteorites, terrestrial rock types and sensor data is the classification method of choice. It addresses the robotic rock classification issues in a natural manner, unlike other classification systems such as neural networks or support vector machines.

The Bayesian paradigm expounded in this chapter is useful beyond the problem of distinguishing meteorites from rocks. It is appropriate in any application involving robotic search and object identification, such as landmine clearance, mineral prospecting, or agricultural soil profiling. All these applications face similar problems to those of robotic meteorite search.

### 2.1. Issues in robotic rock and meteorite classification

We now examine each of the issues affecting the design of a robot terrestrial rock and meteorite classifier:

- *Classifier must accept incomplete data and revise its conclusions as more data becomes available.*

Multiple sensors are needed to distinguish all terrestrial rocks types from all different meteorites. Sometimes a particular sensor can discriminate between the two, but it is common for that data from additional sensors may be required. However, acquiring sensor data from a mobile robotic platform is expensive and time consuming. Therefore, it is desirable to initially classify samples using data from long range sensors such as a color camera before committing resources (time and energy) to maneuver the robot so as to get data from short range sensors such as a spectrometer or metal detector. A classifier that requires that all sensor data be acquired before venturing a hypothesis is unacceptable.

The classifier must compound evidence for hypotheses as it becomes available. A corollary of this is that the classifier accepts *prior* evidence from other sources, such as knowledge of the likelihood of different rock types in a location (which may be very different from their relative proportions in the training data).

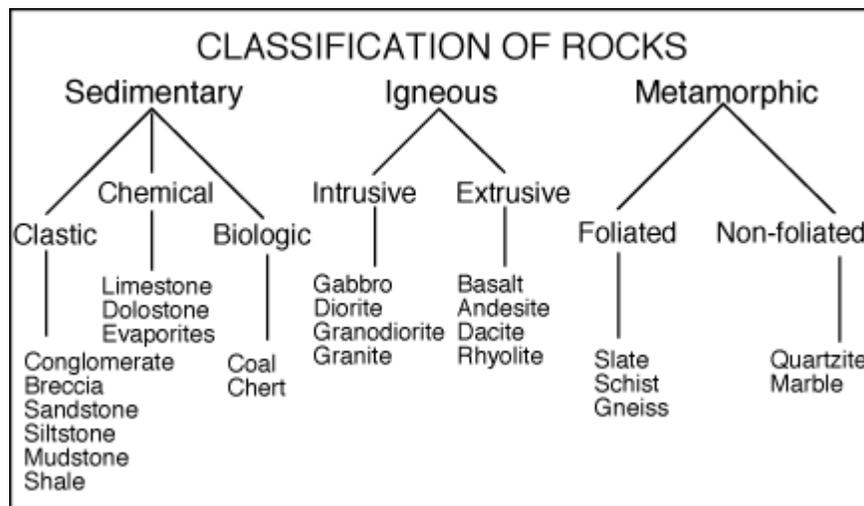
- *Uncertainty of rock classifications*

Geological rock classifications can be quite ambiguous, depending on a rock's mineral composition, its geological context, how it was formed and whether it was altered during its history (Figure 2.1). Rocks are conglomerates of different mineral crystals. The relative proportions of these minerals can vary continuously, and there is no sharp dividing line between certain rock types characterized by differing amounts of these minerals (Figure 2.2). Other differences between rock types arise from different cooling rates when the rocks formed from solidifying magma, leading to different crystal sizes. For example, Basalts and Gabbros (commonly called Dolerite) are two such rock types, consisting of identical minerals

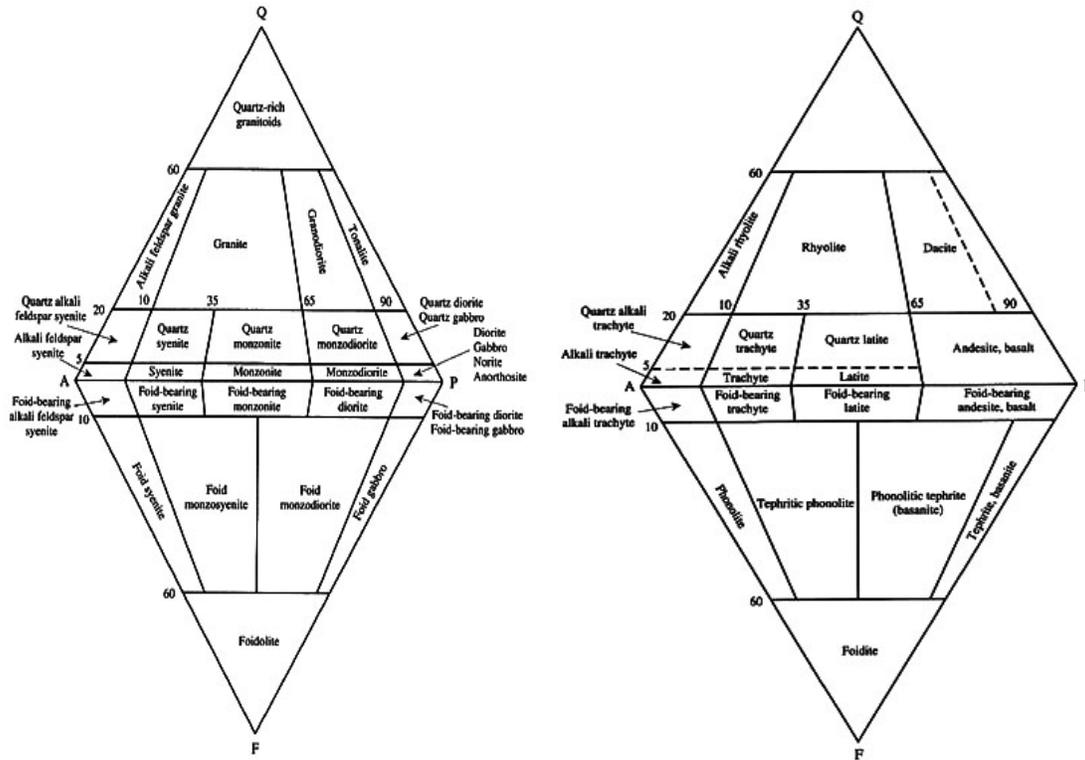
yet having different crystal grain sizes because of different cooling rates when they formed. Telling the difference between them in the field is very difficult. Clearly, the distinction between many rock types is fuzzy, and intermediate rock types abound.

Furthermore, it is common to qualify a rock classification with additional observations, such as that it exhibits traits from another rock type, is contaminated with something, has been thermally or chemically altered, or contains inclusions of another rock or mineral.

The differences between meteorites and terrestrial rocks are not always clear-cut either. Most meteorites are believed to have formed from solidifying melt on a parent body, and are therefore igneous in origin. They have minerals in common with terrestrial igneous rocks, especially Basalts. However, the rare and valuable SNC meteorites, including some meteorites from Mars, have a different origin and are more similar to some terrestrial sedimentary or metamorphic rocks.



**Figure 2.1** Basic geological classification of rocks. Igneous rocks are solidified magma and are characterized mainly by the relative proportions of 3 mineral classes (silica, plagioclase and feldspars). If the magma is extruded to the surface and cools quickly, little crystallization takes place, giving rise to fine grained extrusive igneous rocks. Intrusive igneous rocks are mineralogically identical, but come from slowly cooled magma and therefore larger crystal grains. Consequently, there is no proper dividing line between the different igneous rocks. Classes actually represent different areas of a continuum. Sedimentary rocks are formed by the accumulation and ultimate solidification of sediment, usually derived from the weathering of other rocks. They come in innumerable varieties, depending on the sediment source, the sedimentation process, what contaminants are present. Metamorphic rocks are pressure-cooked sedimentary or igneous rocks, and tend to have properties intermediate between the two (depending on how rare or well done they are).



**Figure 2.2** Classification of igneous rocks according to the relative amounts of their main constituent minerals (Alkali Feldspars, Plagioclase or Quartz) and whether they were formed in a plutonic (left) or volcanic (right) process. Note the lack of clear boundaries, and the distinctions based on the rate at which the original magma cooled (plutonic or volcanic). (Reproduced from [2.1]).

To be effective, a robotic rock classifier must be able to maintain multiple hypotheses and handle these ambiguous classifications. Even in the simplified problem of simply distinguishing terrestrial rocks from meteorites, as significant classes of meteorites are very similar to certain terrestrial rocks.

- *Learning from examples*

The field observable properties that distinguish different rock types from each other and from meteorites are often vague and not explicitly understood by human experts. For example, Cassidy [ref] states that in order to identify a meteorite “first you find a few, then you know what to look for”. The classifier must be able to learn from examples.

Unfortunately, obtaining complete training *field* data is expensive, especially if conditions experienced by the robot are to be replicated. A classifier that requires vast (i.e. several hundred or more) training samples in order to learn how to discriminate different rocks and meteorites is unacceptable.

Furthermore, meteorites and certain terrestrial rock types are rare. A statistically representative sample of the rocks in an area would need to be very large in order to contain a useful number of these samples. As large sample sets are not available, the classifier must tolerate biased sample sets.

- *High dimensionality*

Rock classification from multiple sensors involves a multitude of features. For example, a 15-element feature vector may conservatively describe a reflectance spectrum. Without additional constraints on the relationships between sensor data and rock classes the problem is intrinsically of high dimension. The infamous “curse of dimensionality” states that the number of training examples needed to learn the relationships between classification’s and observed data increases exponentially with dimension. The scarcity of training data makes this unacceptable.

Imposing constraints, usually embedded in the classifier form or in the choice of features describing sensor observations, serves to reduce the intrinsic dimensionality of the problem. It is desirable that the rock classifier has a structure that relates easily to the problem domain, making it easier to impose constraints due to geology or sensor physics.

- *Sensor data and rock features are dependent*

There are significant statistical dependencies between different sensor readings, as well as features extracted from those readings. For example, the albedo of a rock sample measured with a spectrometer is linked to the average pixel brightness of the same rock in a color image. Complicating things further, these quantities are not even conditionally independent given the rock type.

- *Classifier must tolerate sensor noise.*

Taking sensor measurements in the field entails significant random errors because of the uncontrolled environmental conditions. For example, the visible surface of a rock will be at a random orientation with respect to the light sources, therefore a measurement of the surface albedo will have similar random fluctuations. Additionally there are systematic effects that must be considered. Rocks may be covered with a “desert varnish”, preventing direct sensing of a fresh surface. Snow and water are sometimes present and obscure otherwise observable features.

Autonomous deployment of sensors by the robot can give results significantly different from those obtained by a human operating the same sensor equipment. These effects must be accounted for.

The classifier must be robust to sensor noise, and perhaps even recognize certain systematic errors if and when they occur. For this purpose, a classifier that can evaluate new sensor data based on its present knowledge of the world would be suitable. In other words, a *reversible* classifier is desired.

- *Geographically dependent rock and meteorite statistics*

The likelihood’s of finding different meteorite and terrestrial rock types in an area depend on where that area is. For example, the Patriot Hills region of Antarctica contained mostly marbles and limestones, but no meteorites. On the other hand, Elephant Moraine contains many meteorites and is dominated by dolerite and basalt rock types. A classifier optimized to work in one area will not work as well in an area characterized by different rocks.

Furthermore, it is difficult to know *a priori* the chances of finding different rock types or meteorites in an area. Geological maps, if available, only indicate the dominant rock type over a very large area. It does not indicate all rock types present, their relative proportions, or any small-scale variations.

A mobile robot explorer can examine many samples over a large area. The rock and meteorite statistics will likely change as the robot moves. Autonomously learning the local rock and meteorite distributions and exploiting the correlations between nearby rocks can significantly boost classification accuracy.

## 2.2. Bayes network generative model

The restrictions imposed by terrestrial rocks and meteorite classification from a mobile robotic vehicle argue for a classification system based on a generative statistical model. Classification ambiguity is handled by directly computing the class probabilities of a rock sample, and new evidence from sensors can be compounded in a natural way using Bayes rule.

A generative statistical model can be built using a Bayes network. This is an augmented graphical model of the relationships between classifications and observations. A Bayes network model can be learned from example data. Its structure can be intuitively related to the specific problem domain, thus constraining the intrinsic dimensionality of the problem. Dependencies between sensor observations can be accounted for.

### 2.2.1. Notation

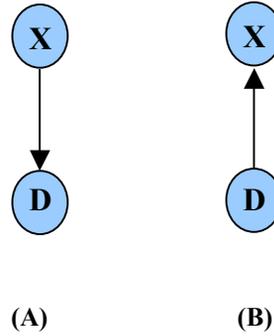
Let the random variable  $\mathbf{X}$ , taking its values from a discrete set  $\{c_1, c_2, \dots, c_M\}$ , represent possible classifications for a rock sample under investigation by the robot. In the pure meteorite recognition problem these classes would be  $\{meteorite, \neg meteorite\}$ . Measured sensor data for the sample is in the form of an  $n$ -dimensional random vector  $\mathbf{D} = [d_1, d_2, \dots, d_n]^T$ . Given such a measurement  $\mathbf{D}$  we wish to determine the most likely classification  $\mathbf{X}$ .

In general, boldface Latin letters (e.g.  $\mathbf{X}$ ,  $\mathbf{Y}$ ,  $\mathbf{a}$ ,  $\mathbf{b}$ , etcetera) will represent random variables (discrete or continuous, and possibly multi-dimensional) and italicized lower case Latin letters (e.g.  $i$ ,  $j$ ,  $k$ ,  $l$ ) will represent discrete states of a random variable. If the random variable is continuous, then these represent the state associated with some predetermined interval on the range of the random variable.

$P(\mathbf{Y} = i)$  is the probability that  $\mathbf{Y}$  is in the state  $i$ . Sometimes  $P_i$  is used if  $\mathbf{Y}$  can be assumed from the context. Similarly,  $P(\mathbf{Y})$  is used if  $i$  does not need to be explicitly stated. For continuous variables  $\mathbf{y}$  the probability *density*  $p(\mathbf{y})$  is used (unless  $\mathbf{y}$  has been quantized).

$P(\mathbf{X}|\mathbf{D})$  is the conditional probability that  $\mathbf{X}$  is in some state given data measurement  $\mathbf{D}$ . If both  $\mathbf{X}$  and  $\mathbf{D}$  are discrete variables, then their conditional relationship is given by the conditional probability matrix (CPM)  $P(\mathbf{X}=i | \mathbf{D}=j)$ .

## 2.2.2. Generative and diagnostic statistical models



**Figure 2.3** “Generative” (A) and “diagnostic” (B) belief networks for classifying the random variable  $\mathbf{X} \in \{c_1, c_2, \dots, c_M\}$  given measurement data  $\mathbf{D}$  [2.1]. The generative model specifies the distribution of measurement data from the classification, whereas the diagnostic model directly models class probabilities given a measurement. To classify with the generative model the marginal probabilities  $P(\mathbf{X})$  and conditional probabilities  $P(\mathbf{D}|\mathbf{X})$  must be specified, whereas in the diagnostic model only  $P(\mathbf{X}|\mathbf{D})$  need be specified.

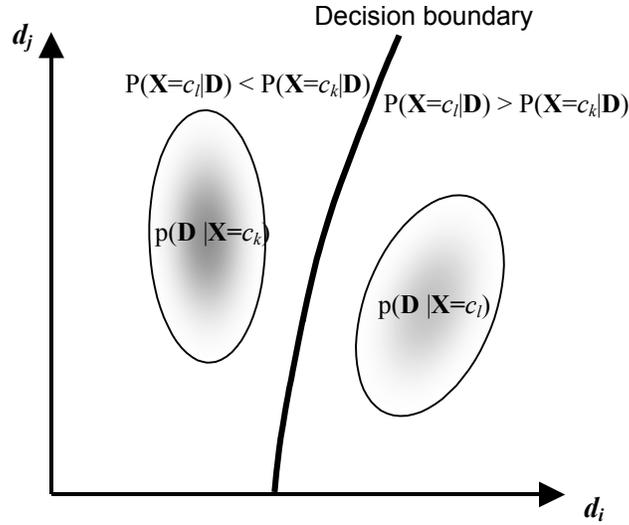
Consider the classification problem in which data are labeled by a random variable  $\mathbf{X}$ , taking its values from a discrete set  $\{c_1, c_2, \dots, c_M\}$ . The measured data are in the form of an  $n$ -dimensional random vector  $\mathbf{D} = [d_1, d_2, \dots, d_n]^T$ . Given a measurement  $\mathbf{D}$  we wish to determine the classification  $\mathbf{X}$ .

A generative statistical model (Figure 2.3(A)) [2.1] for the joint probabilities of  $\mathbf{X}$  and  $\mathbf{D}$  is specified by the marginal probabilities  $P(\mathbf{X})$  and conditional probabilities  $P(\mathbf{D}|\mathbf{X})$ . The posterior probability  $P(\mathbf{X}=c_i|\mathbf{D})$  is computed using Bayes rule:

$$(Eqn 2-1) \quad P(\mathbf{X} = c_i | \mathbf{D}) = \frac{p(\mathbf{D} | \mathbf{X} = c_i)P(\mathbf{X} = c_i)}{\sum_{k=1}^M p(\mathbf{D} | \mathbf{X} = c_k)P(\mathbf{X} = c_k)}$$

The most likely classification of the data is the class  $c_i$  with the greatest posterior probability  $P(\mathbf{X}=c_i|\mathbf{D})$  (Figure 2.4). A diagnostic model would directly specify  $P(\mathbf{X}=c_i|\mathbf{D})$ .

In fact, determining the optimal decision boundary (Figure 2.4) does not require that the posterior be computed explicitly. Some classifiers, such as neural networks, support vector machines, and decision trees attempt to learn it directly from training data. Nonetheless, there are advantages to specifying the posterior. It provides a measure of the classification confidence and handles ambiguities amongst multiple classifications.



**Figure 2.4** Classification decision boundary for discriminating between two classes based on the posterior density  $P(\mathbf{X}|\mathbf{D})$ . For simple classification it is sufficient to specify the decision boundary only. However, knowing the full joint density  $P(\mathbf{X}, \mathbf{D})$  allows evidence for a classification to be incrementally compounded in a natural manner.

With marginal density  $P(\mathbf{D})$  specified, the diagnostic model and the generative model are equivalent parameterizations of the joint density  $P(\mathbf{X}, \mathbf{D})$  and lead to equivalent classifications. Nevertheless, they are not equivalent, and have complementary advantages and disadvantages [2.1][2.3]. Generative models are more “modular” than diagnostic models. The class conditional probabilities  $P(\mathbf{D}|\mathbf{X}=c_k)$  are likely to be local, characteristic functions of the objects being classified, invariant to the nature and number of the other classes [2.1][2.4]. Furthermore, generative models correspond more to the notion of cause and effect, with the rock type  $\mathbf{X}$  being the root cause of the observed sensor data  $\mathbf{D}$ .

A major statistical advantage of a generative model is that if the model specification is nearly correct then estimation will be more efficient than for a diagnostic model. That is, fewer examples are needed to learn the joint density [2.1]. Recall that not many training examples are available for rock and meteorite classification.

On the other hand, diagnostic models are more robust for classification. [2.1] shows that whole families of class conditional densities give rise to similar diagnostic model discriminant functions. In other words, it is not necessary to completely specify a statistical distribution in order to determine a decision surface (Figure 2.4).

The prime advantage of an explicit generative statistical model for the purpose of robotic rock and meteorite classification is the natural way in which evidence from sensor observations can be incrementally compounded. From Bayes rule it follows that

$$(Eqn 2-2) \quad P(\mathbf{X} = c_j | \mathbf{d}_{i+1}, \mathbf{d}_1, \dots, \mathbf{d}_i) = \frac{P(\mathbf{d}_{i+1} | \mathbf{X} = c_j, \mathbf{d}_1, \dots, \mathbf{d}_i) P(\mathbf{X} = c_j | \mathbf{d}_1, \dots, \mathbf{d}_i)}{\sum_{k=1}^M P(\mathbf{d}_{i+1} | \mathbf{X} = c_k, \mathbf{d}_1, \dots, \mathbf{d}_i) P(\mathbf{X} = c_k | \mathbf{d}_1, \dots, \mathbf{d}_i)}$$

Thus, as measurements or observations  $\{\mathbf{d}_1, \mathbf{d}_2, \dots\}$  are made, the posterior probability  $P(\mathbf{X}|\dots)$  is refined, based on its values from the previous measurements, and updated with the latest measurement. This

requires a model of the conditionals  $P(\mathbf{d}_{i+1}|\mathbf{X}, \mathbf{d}_1, \mathbf{d}_2, \dots, \mathbf{d}_i)$  for every  $i \in \{1, 2, \dots, n\}$ . In some cases, the measurements are conditionally independent of each other given the classification  $\mathbf{X}$ . Therefore,

$$(Eqn\ 2-3) \quad P(\mathbf{d}_{i+1}|\mathbf{X}, \mathbf{d}_1, \mathbf{d}_2, \dots, \mathbf{d}_i) = P(\mathbf{d}_{i+1}|\mathbf{X})$$

Therefore:

$$(Eqn\ 2-4) \quad P(\mathbf{X} = c_j | \mathbf{d}_{i+1}, \mathbf{d}_1, \dots, \mathbf{d}_i) = \frac{P(\mathbf{d}_{i+1} | \mathbf{X} = c_j)P(\mathbf{X} = c_j | \mathbf{d}_1, \dots, \mathbf{d}_i)}{\sum_{k=1}^M p(\mathbf{d}_{i+1} | \mathbf{X} = c_k)P(\mathbf{X} = c_k | \mathbf{d}_1, \dots, \mathbf{d}_i)}$$

This is known as a “naïve” Bayes model. It is very efficient, requiring the specification of  $P(\mathbf{X})$ ,  $P(\mathbf{d}_1|\mathbf{X})$ , ...,  $P(\mathbf{d}_n|\mathbf{X})$ . Unfortunately, most “real world” problems cannot be so simplified, as measurements may affect one another. This can become intractable unless statistical independencies can be exploited. Bayes networks (next section) do just that.

Note that the generative model separates out the class priors  $P(\mathbf{X})$  from the conditionals  $P(\mathbf{D}|\mathbf{X})$ , unlike diagnostic models in which the priors are tangled up with the conditionals. This is extremely important for rock classification. The priors indicate what rocks are likely to be found in an area. Setting them explicitly allows the model to be optimized for a particular geographic area. Furthermore, the conditionals can be independently learned from for each rock type [2.5]. Therefore, an unbiased sample set of all rock types and meteorites is *not* required for training purposes.

Generative models have sometimes been criticized for the explicit use of prior probabilities  $P(\mathbf{X})$ . They are often poorly characterized and therefore set arbitrarily. However, it should be noted that *all* classification methods make implicit use of these priors. They are usually hidden in the choice of the training data set. Unless this is an unbiased set, the implicitly learned priors will themselves be arbitrary

### 2.2.3. Bayes networks

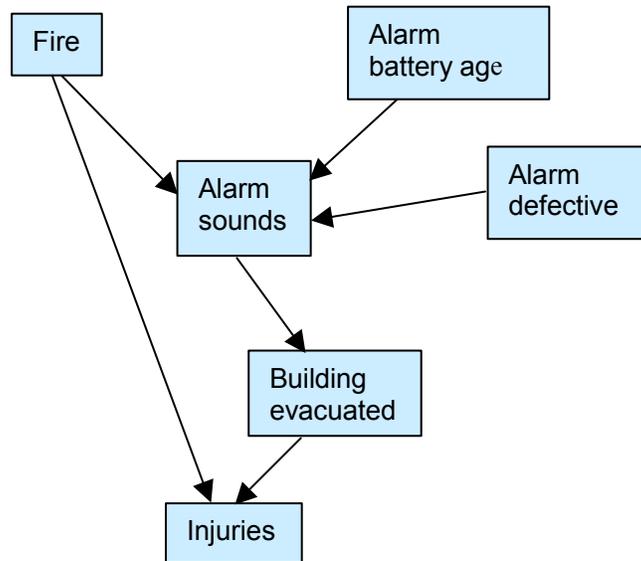
A *Bayes network* [2.4][2.5][2.6][2.7][2.8][2.9] (or *probabilistic influence diagram*) is a network consisting of a directed acyclic graph (DAG), where each node represents a (random) variable in the model. Each variable is *conditionally independent* given values for the variables represented by its parent nodes, and stored with each node is a conditional probability matrix (CPM) describing the probability distribution of its variable given the parent variables. Nodes with no parents store the marginal distributions of Any such network fully describes a joint probability distribution on its associated variables, and conversely, any distribution can be described by a network (not necessarily unique)

Taking advantage of conditional independencies amongst the variables, a Bayes network is both a compact and efficient representation of their joint distribution. Moreover, there exist effective algorithms [2.7][2.4][2.6][2.5] to compute posterior conditional probabilities of any variable given any other variable(s) in the model. This is essentially the task of classification, where we wish to compute the probability distribution of classes given observations.

The advantage of a Bayes network for rock and meteorite classification is that they can efficiently represent the generative statistical model of rock classifications and sensor data. Statistical dependencies between sensor data and rock classifications are explicitly accounted for by the network structure.

If the network structure is known, then it is relatively easy to determine the conditional densities associated with each node from example data [2.5]. The network structure constrains the dependencies between variables, therefore limiting the intrinsic dimensionality of the problem.

Consult the references at the end of the chapter for a detailed discussion on Bayes network theory and how to construct and evaluate them. The subsequent chapters will detail the design of a Bayes network for rock and meteorite identification, and how the conditional densities are learned.



**Figure 2.5** Example Bayes network representing joint distribution of random variables, and their conditional dependencies. Given that a fire occurred and the building was evacuated, the chance of injuries is independent of whether the fire alarm was activated.

### 2.3. Other classification systems

The Bayes network generative model classification paradigm can handle all the issues in robotic rock classification. Other classification systems can handle some of them, but none of them can handle them all and in as natural a manner. The following is a review of the alternatives:

- *Rule based systems*

Rule based systems are acceptable for cases where a human expert can elucidate all the rules allowing one to distinguish between objects. This is not the case for rocks and meteorites, the distinctions must be automatically learned from examples. Further, rule based systems do not handle ambiguity, vague distinctions, noise and uncertainty. They are clearly unacceptable for this task.

Augmented rule based systems attempt to handle uncertainty and ambiguity. However, they still do not allow learning, and make unrealistic assumptions about the independence of data.

- *Predicate logic*

Predicate logic has the same failings as rule based systems: an inability to learn from examples or handle uncertainty.

- *Decision trees*

Decision trees can learn to classify objects from training examples, and can tolerate vague class distinctions. For simple classification tasks they are very effective. However, they do not express uncertainty or express multiple hypotheses. More importantly, they cannot compound evidence as it becomes available, making them unacceptable for classification from a robotic vehicle.

Furthermore, Decision Tree's cannot be incrementally trained. For a robot classification system the initial training data might not accurately reflect what happens in the field. In this case the classifier would need to incorporate additional training examples as they become available.

- *Memory based classification*

Memory based classifier's include k-Nearest Neighbor (kNN) classifiers and Parzen estimation. Learning is trivial with these systems, although run time performance degrades with increasingly big training sets. Unfortunately, they do not enforce any physical constraints that might limit the dimensionality of the problem, making them vulnerable to the curse of dimensionality and limiting their ability to generalize from limited training data.

Although they can be extended to handle ambiguity and uncertainty, they do not do so naturally, nor do they permit the compounding of evidence as it becomes available.

- *Connectionist Models*

These include artificial neural networks, Hopfield networks, Boltzmann machines, perceptrons and others. Neural networks, particularly backpropagation feedforward networks, have proven very popular for classification tasks. They are very effective at learning from data, and can be incrementally retrained as more example data becomes available. Furthermore, depending on their structure, they can generalize from limited training data, as long as the problem is not intrinsically of a high dimension. They can be effective at learning vaguely defined distinctions.

Unfortunately, building a neural network appropriate to a problem is often akin to black magic. The optimal network structure for a particular task does not immediately relate to the task domain.

The main drawback for conventional neural networks in the robotic rock classification task is their inability to compound evidence or accept prior information.

In fact, under certain conditions, a feedforward neural network with logistic transfer functions is able to learn the posterior distribution  $P(X|D)$  from a training data set [2.1][2.3]. In this case they are an efficient representation for a diagnostic statistical model of classifications and measured data, with the associated advantages and disadvantages.

- *Dempster-Shafer theory*

Dempster-Shafer theory [2.10][2.11] is a technique for combining different propositions, assigning to each an interval [*Belief, Plausibility*] in which the degree of belief must lie. This interval measures not only the belief in a proposition ("rock is a meteorite"), but also the amount of information available to support our belief. This contrasts with Bayesian statistical classifiers, which do not distinguish between uncertainty due to nothing being known about something, and uncertainty due to the random nature of something. For example, consider a coin toss. A priori one could assume a 0.5 probability of heads. After 1000 coin

tosses it is observed that approximately half the time heads is observed, confirming the 0.5 probability of heads the next time the coin is tossed. In the first case the 0.5 probability is a guess, in the second case, it is a well-substantiated probability. A standard Bayesian classifier does not make this distinction. This difference can matter when deciding if more evidence needs to be collected.

Unfortunately, although elegant and able to compound evidence, Dempster-Shafer theory lacks the statistical rigor of Bayesian statistics. The beliefs attached to different propositions are *ad hoc*, and cannot be learned from training data. This is a fatal flaw making Dempster-Shafer theory unsuitable for robotic rock classification.

Furthermore, Dempster-Shafer theory does not extend to continuous variables. Although this can be overcome by quantizing them, it is only practical if the relationships between them can be automatically learned.

- *Fuzzy logic*

Fuzzy logic [2.13] is an extension of Boolean logic based systems, able to express uncertainty and ambiguity. Evidence for a hypothesis can be compounded. However, it is very ad hoc in some respects, lacks the rigor of Bayesian statistics, and is unable to learn relationships.

- *Support vector machines*

Support vector machines (SVM's) [2.14] are classifiers based on explicitly learning a decision surface from data. They have gained popularity because of their effectiveness at learning and generalizing well from limited training data.

However, like neural networks, they do not incorporate explicit priors, and are not appropriate for compounding evidence as it becomes available. Furthermore, they do not handle ambiguity amongst multiple possible classifications or express uncertainty.

## 2.4. Conclusions

The viability of a Bayes network based generative statistical model for rock classification has been established. Classification uncertainty and ambiguity is handled by computing explicit probabilities for each possible rock class, rather than just the most likely. Evidence from sensor readings can be incrementally compounded via Bayesian update. Prior rock and meteorite likelihood's in an area are explicitly accounted. The structure of a Bayes network reflects the statistical relationships between rock samples and sensor measurements. These can be deduced from knowledge of geology and the physics of the sensors. In this way, relationships between variables are constrained by domain specific knowledge, and the intrinsic dimensionality controlled. Given a suitable network structure, the full statistical model can be learned from the limited and statistically biased training data available.

It will be shown in subsequent chapters how the Bayes network based generative model can be extended to allow autonomous profiling of the environment to learn rock probabilities in an area and exploit correlations amongst the rock samples in an area. Furthermore, in addition to been able to handle incremental data, a Bayes network model can be used to determine which unused sensors are likely to be useful to classify a sample, enabling active sensor selection.

## 2.5. References

- [2.1] Dietrich, R.V. and Skinner, B.J., *Rocks and Rock Minerals*, John Wiley & Sons, Inc, 1979.
- [2.2] M.I. Jordan, “Why the Logistic Function? A Tutorial Discussion on Probabilities and Neural Networks”, *Computational Cognitive Science Technical Report 9503*, August 1995
- [2.3] M.I. Jordan and C.M. Bishop, “Neural Networks”, *CRC Handbook of Computer Science*, A. Tucker (ed.), CRC Press, In press 1996
- [2.4] Pearl, Judea (1988) *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*, Morgan Kaufmann, San Mateo, CA. 2nd edition 1991.
- [2.5] Spiegelhalter, David J., A. Philip Dawid, Steffen L. Lauritzen and Robert G. Cowell, “Bayesian analysis in expert systems”, *Statistical Science*, 8(3), 219-283. 1993.
- [2.6] Ross D. Shachter, “Probabilistic inference and influence diagrams”, *Operations Research*, vol. 36, No. 4. July – August 1988.
- [2.7] Ross D. Shachter, “Evaluating influence diagrams”, *Operations Research*, vol. 34, No. 6. November-December 1986
- [2.8] Charniak, Eugene “Bayesian networks without tears” in *AI Magazine*, 12(4), 50-63, Winter 1991.
- [2.9] Jensen, Finn V., *An Introduction to Bayesian Networks*, Springer-Verlag, 1996.
- [2.10] G. Shafer, *A Mathematical Theory of Evidence*, Princeton, NJ: Princeton University Press, 1976.
- [2.11] A.P. Dempster, “A Generalization of Bayesian Inference”, *Journal of the Royal Statistical Society, Series, B* 30:pp205-247, 1968.
- [2.12] E. Rich and K. Knight, *Artificial Intelligence*, McGraw-Hill, 1991.
- [2.13] L. Zadeh, “A Theory of Approximate Reasoning”, *Machine Intelligence 9*, ed. J. Hayes, D. Michie, and L.I. Mikulich, pp. 149-194, New York: Halstead Press, 1979.
- [2.14] Christopher J.C. Burges, “A Tutorial on Support Vector Machines for Pattern Recognition”, *Knowledge Discovery and Data Mining*, 2(2), 1999

## Bayes network classifier topology

The previous chapter established the need for a Bayes network generative model of the relationships between rock and meteorite classifications. A Bayes network model is defined by (i) a network topology and (ii) the conditional probabilities associated with each node (random variable) in the network giving the conditional distribution of that random variable given the parent variables.

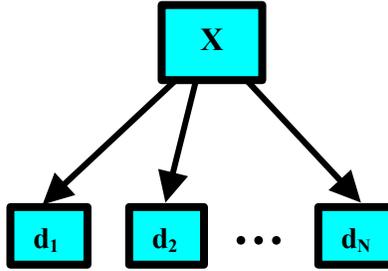
This chapter describes the development of a suitable Bayes network topology for classifying meteorites and terrestrial rocks using features extracted from visible reflectance spectra, color images and metal detector data. These features and sensors are motivated and described in subsequent chapters.

A fully connected Bayes network will describe any joint distribution of its variables, but does not exploit any independencies and so is as inefficient and memory consuming as representing the joint distribution directly. There exist [3.1][3.2][3.3][3.4] algorithms for learning suitable network topologies from example data. However, they require large quantities of data (exponential on the number of nodes). Furthermore, the resulting networks may be sub-optimal since they neglect domain specific knowledge, and nor do they use “hidden” variables (more later).

Choosing an appropriate network topology is a way to constrain the joint distribution of random variables (classification, observations and “explanatory” variables) with application domain specific knowledge.

The following sections will consider various network topologies suitable for classification and describe the network used for rock and meteorite classification. More references to Bayes network topologies are [3.6][3.7][3.8][3.9][3.11][3.12][3.13][3.14][3.15][3.16][3.17]

### 3.1. Naïve Bayes architecture



**Figure 3.1** Naïve Bayes classifier for  $X$  given conditionally independent data  $d_1, d_2, \dots, d_N$ .

A naïve Bayes classifier is the simplest Bayes network classifier, having the structure shown in Figure 3.1. All observations  $d_1, d_2, \dots, d_N$  are conditionally independent of each other given  $X$

$$(Eqn 3-1) \quad \bar{\mathbf{D}} = [d_1 \dots d_N]$$

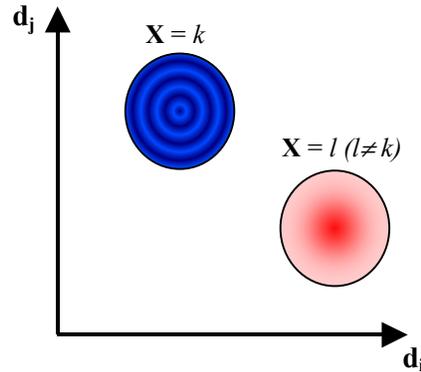
$$(Eqn 3-2) \quad p(\bar{\mathbf{D}} | \mathbf{X}) = \prod_{i=1}^N p(d_i | \mathbf{X})$$

Therefore, using Bayes rule, the posterior probability of  $X$  given new data is

$$(Eqn 3-3) \quad P(\mathbf{X} | d_{i+1}, d_1 \dots d_i) = \frac{p(d_{i+1} | \mathbf{X})P(\mathbf{X} | d_1 \dots d_i)}{\sum_k p(d_{i+1} | \mathbf{X} = k)P(\mathbf{X} = k | d_1 \dots d_i)}$$

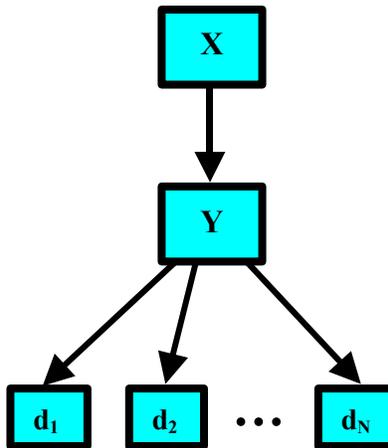
This is equivalent to computing the posterior probability on  $\mathbf{X}$  given  $\mathbf{d}_{i+1}$  with  $P(\mathbf{X} | \mathbf{d}_i \dots \mathbf{d}_1)$  substituted for the (marginal) prior  $P(\mathbf{X})$ .

The naïve Bayes network is very efficient, permitting trivial compounding of evidence. For this reason it is widely (sometimes unknowingly) used for classification. Unfortunately, there are serious limitations in the probability densities (or distributions)  $p(\bar{\mathbf{D}} | \mathbf{X})$  that it can represent (Figure 3.2), making it unsuitable for most “real world” applications.



**Figure 3.2** Conditional probability densities  $p(\bar{\mathbf{D}} | X=i)$  ( $i$  are possible states of  $X$ ) representable by the naïve Bayes network of Figure 3.1. Each density is subject to the requirement that the elements  $d_1 \dots d_N$  of  $\bar{\mathbf{D}}$  be independent of each other (represented here by drawing them as spherical density functions).

### 3.2. Naïve Bayes network with intermediate nodes

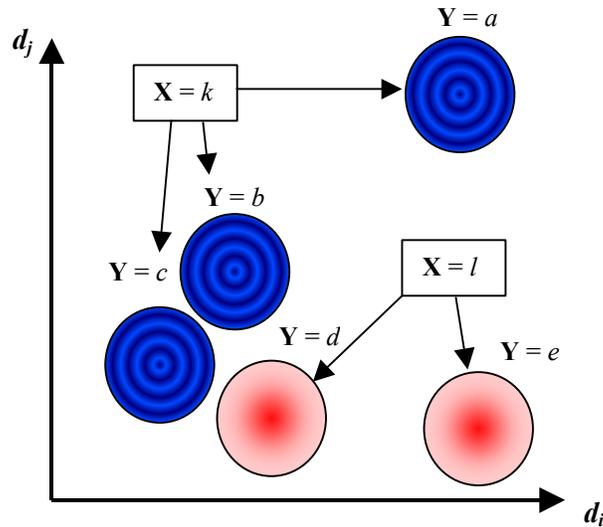


**Figure 3.3** Previous naïve Bayes network modified by addition of intermediate node  $\mathbf{Y}$ . This allows certain dependencies amongst  $\mathbf{d}_1 \dots \mathbf{d}_N$  given  $\mathbf{X}$ , however, they must now be independent given  $\mathbf{Y}$ .

Introducing an intermediate (or *hidden*) node  $\mathbf{Y}$  in between the node  $\mathbf{X}$  and the data nodes  $\mathbf{d}_1 \dots \mathbf{d}_N$  (Figure 3.3) allows the representation of more complex data distributions, no longer requiring that they be conditionally independent given  $\mathbf{X}$ . The densities  $p(\bar{\mathbf{D}} | \mathbf{X})$  (Figure 3.4) are mixture models, with mixing coefficients  $P(\mathbf{Y} | \mathbf{X})$ :

(Eqn 3-4) 
$$p(\bar{\mathbf{D}} | \mathbf{X}) = \sum_{k=1}^M P(\mathbf{Y} = k | \mathbf{X}) p(\bar{\mathbf{D}} | \mathbf{Y} = k)$$

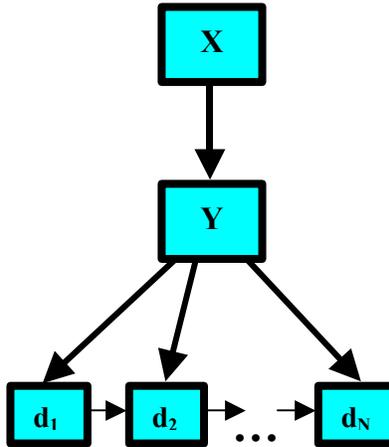
The only restriction is that  $\mathbf{d}_1 \dots \mathbf{d}_N$  are independent given  $\mathbf{Y}$ .



**Figure 3.4** Densities represented by network with intermediate unit. For each  $i$ ,  $p(\bar{\mathbf{D}} | \mathbf{X}=i)$  is a mixture model, with the restriction that in each model component the data are independent.

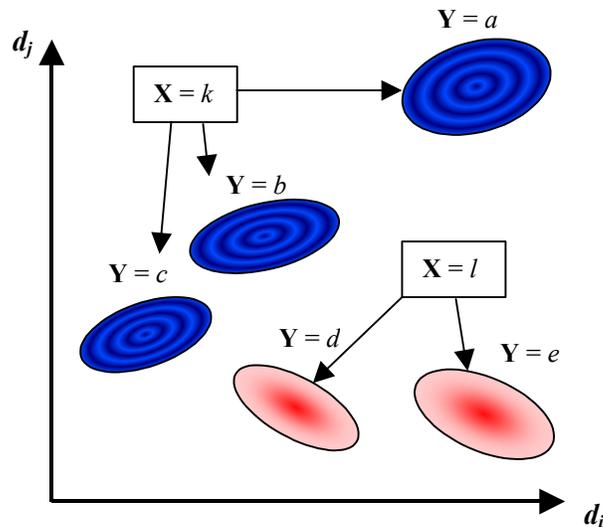
Crucial to the success of this network architecture for classifying  $\mathbf{X}$  given  $\bar{\mathbf{D}}$  (that is, computing  $p(\bar{\mathbf{D}} | \mathbf{X})$ ) is the identification of suitable intermediate (or hidden) variable(s)  $\mathbf{Y}$ , representing groupings into which the data naturally divides. It is possible to use a clustering algorithm [3.10] to identify such a partition of variable  $\mathbf{Y}$  of  $\bar{\mathbf{D}}$ . Unfortunately, this presents problems with network training, when the conditional probabilities are learned from a set of examples (Chapter 4). Stable learning methods require that example data, with *all* nodes assigned definite states. This cannot be done if  $\mathbf{Y}$  represents a set of abstract states, and each example data point is only assigned a membership probability to each state. Otherwise,  $\mathbf{Y}$  can be deduced using expert knowledge from the problem domain, in which case determining the value of  $\mathbf{Y}$  for each example point is possible.

### 3.3. Bayes network with weak data dependencies and intermediate node



**Figure 3.5** “Semi-naïve” Bayes network classifier, with an intermediate node, incorporating limited dependencies amongst  $d_1 \dots d_N$ . In this example, only direct dependence between adjacent members of  $\bar{\mathbf{D}}$  is allowed.

Sometimes the presence of the intermediate node  $Y$  is insufficient to remove any direct dependencies between the observed data. This happens particularly when two observations are derived from similar sensing modalities, such as measuring the albedo of a rock sample with two different instruments (a color camera and spectrometer in this application).

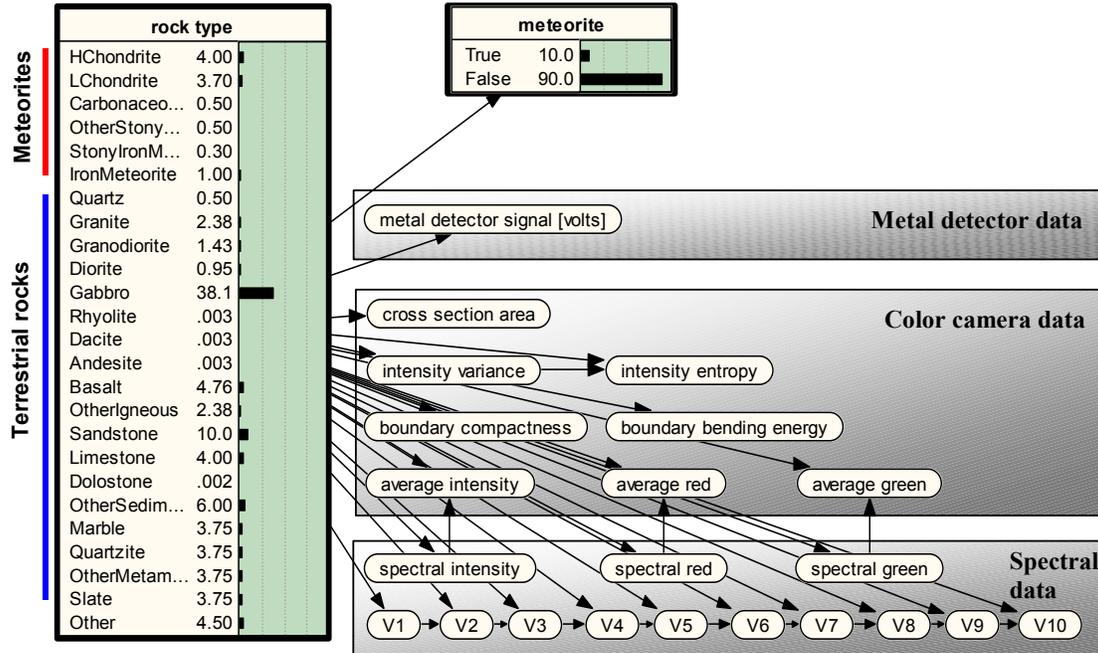


**Figure 3.6** Conditional densities represented by Bayes network in Figure 3.5. Note that the conditional densities  $P(\bar{\mathbf{D}} | \mathbf{X})$  is still a mixture model, but the components may now each have limited dependencies amongst the data. The exact nature of these dependencies depends on the problem domain. It is incumbent on the designer to choose feature vectors  $\bar{\mathbf{D}}$  that minimize these.

Modeling this requires links connecting those data nodes that cannot be made conditionally independent given  $Y$ . The precise topology of these links depends on the application problem domain. Naturally, it is

incumbent upon classifier designers to choose the observable data (or *features*) so as to minimize cross dependencies. This requires close attention to the physics of the sensors, and the processing applied to their raw signals to extract features to put into the Bayes network (see the subsequent chapters on sensors and sensor feature extraction to see how this is done for the rock and meteorite classification problem).

### 3.4. The rock and meteorite Bayes network classifier



**Figure 3.7** Bayes network for terrestrial rock and meteorite discrimination, and general rock classification, using features calculated from robotic sensor data (images, spectra and metal detector data.)

The Bayes network topology chosen to model the relationship between whether a rock is terrestrial or a meteorite, and the observed sensor data is shown in Figure 3.7. It is based on the network of Figure 3.5. The intermediate node **Y** is in this case the rock type of the sample. Its possible states consist of the most common classes of meteorites and terrestrial rocks according to various references on meteorites [3.18][3.19][3.20][3.21] and geology [3.27][3.28], making sure to include rock classes found in the Elephant Moraine region of Antarctica where this classifier is to operate [3.23][3.24][3.25][3.26].

The node **X** in this case corresponds to a binary random variable, **meteorite**, indicating whether the rock sample is terrestrial or a meteorite. A change from the architecture Figure 3.5 is that **meteorite** (**X**) is conditioned upon **rock type** (**Y**), not the other way around. It is more natural to consider **rock type** as the explanatory variable from which all effects follow, rather than **meteorite**. This does not affect the distribution represented by the network [3.15], but requires specifying the marginal densities on rock types instead of the conditional density  $P(\text{rock type} \mid \text{meteorite})$  and the marginal density  $P(\text{meteorite})$ .

In this case, the prior probabilities on rock types for Elephant Moraine are roughly specified in [3.23][3.24][3.25][3.26], and explicitly set in this network. The state of **meteorite** given **rock type** is deterministic (the list of possible rock types is partitioned into terrestrial rocks and meteorites) and also set

explicitly. What remains to be specified are the conditional densities of the sensor feature vectors, which have to be learned from training data (next chapter).

The indicated dependencies between the feature vectors depend on the specifics of each sensor, and are explained in the later chapters about the choice of sensors and feature vectors.

### 3.5. References

- [3.1] Spirtes, P., “Directed Cyclic Graphical Representation of Feedback Models”, Proceedings of the Eleventh *Conference on Uncertainty in Artificial Intelligence*, ed. by Philippe Besnard and Steve Hanks, Morgan Kaufmann Publishers, Inc., San Mateo, 1995.
- [3.2] Spirtes, P., and Meek, C., “Learning Bayesian Networks with Discrete Variables from Data”, in *Proceedings of The First International Conference on Knowledge Discovery and Data Mining*, ed. by Usama M. Fayyad and Ramasamy Uthurusamy, AAI Press, pp. 294-299. 1995.
- [3.3] Spirtes, P., Glymour, C., Meek, C., “Learning Bayesian Networks from Data”, *AISB Quarterly*, 1996
- [3.4] Spirtes, P., Meek, C., and Richardson, T., “Causal Inference in the Presence of Latent Variables and Selection Bias” , in *Proceedings of the Eleventh Conference on Uncertainty in Artificial Intelligence* , ed P. Besnard and S. Hanks, Morgan Kaufmann Publishers, Inc., San Mateo, pp 499-506. 1995.
- [3.5] Castillo, Enrique, Jose Manuel Gutierrez and Ali S. Hadi, *Expert Systems and Probabilistic Network Models*, Springer, New York, 1997.
- [3.6] Charniak, Eugene “Bayesian networks without tears” in *AI Magazine*, 12(4), 50-63, Winter 1991.
- [3.7] Henrion, Max, John S. Breese and Eric J. Horvitz, “Decision Analysis and Expert Systems” in *AI Magazine*, 12(4), 64-91. Winter 1991.
- [3.8] Jensen, Finn V., *An Introduction to Bayesian Networks*, Springer-Verlag, New York. 1991.
- [3.9] Jensen, Frank, Finn V. Jensen and Soren L. Dittmer, “From influence diagrams to junction trees” in Proceedings of the 10<sup>th</sup> conference on *Uncertainty in Artificial Intelligence*, Seattle, 1994.
- [3.10] Lauritzen, Steffen L. “The EM algorithm for graphical association models with missing data” in *Computational Statistics and Data Analysis*, 19(2), 191-201. 1995.
- [3.11] Lauritzen, Steffen L. and David J. Spiegelhalter, “Local computations with probabilities on graphical structures and their application to expert systems” in *Journal of the Royal Statistics Society B*, 50(2), 157-194. 1988.
- [3.12] Neapolitan, Richard E., *Probabilistic Reasoning in Expert Systems: Theory and Algorithms* , John Wiley & Sons, New York. 1990.
- [3.13] Pearl, Judea (1988) *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*, Morgan Kaufmann, San Mateo, CA. 2nd edition 1991
- [3.14] Shachter, Ross D. “Probabilistic inference and influence diagrams” in *Operations Research*, 36(4), 589-604. 1988.

- [3.15] Shachter, Ross D., “Evidence absorption and propagation through evidence reversals” in Proceedings of the 5<sup>th</sup> Workshop on *Uncertainty in Artificial Intelligence*, 303-308. 1989.
- [3.16] Smith, James E., Samuel Holtzman and James E. Matheson “Structuring conditional relationships in influence diagrams” in *Operations Research*, 41(2), 280-297. 1993.
- [3.17] Spiegelhalter, David J., A. Philip Dawid, Steffen L. Lauritzen and Robert G. Cowell, “Bayesian analysis in expert systems” in *Statistical Science*, 8(3), 219-283. 1993.

#### Meteorites

- [3.18] McSween, H.Y., *Meteorites and Their Parent Planets*, Cambridge University Press, pp 6-7, 1987.
- [3.19] Dodd, R.T., *Meteorites: a petralogic-chemical synthesis*, Cambridge University Press.
- [3.20] Heide, Fritz, *Meteorites*, University of Chicago Press, 1964.
- [3.21] Wasson, John T., *Meteorites: classification and properties*, Springer-Verlag, 1974.
- [3.22] W. Cassidy, W., Harvey, R., Schutt, J., Delisle, G., and Yanai, K., “The Meteorite Collection Sites of Antarctica”, *Meteoritics*, 27(5), pp 490 – 525, 1992.

#### Elephant Moraine

- [3.23] Taylor, K.S., *Lithologies and Distribution of Clasts in the Elephant Moraine, Allan Hills South Victoria Land, Antarctica*, M.S. Thesis submitted to Kent State University Graduate College.
- [3.24] Faure, G, Kallstrom, M.L. and Mensing, T.M., “Classification and Age of Terrestrial Boulders in the Elephant and Reckling Moraines”, *Antarctic Journal of the U.S.*, 19(5), pp28-29, 1984.
- [3.25] Faure, G., and Taylor, K.S., “The Geology and Origin of the Elephant Moraine on the East Antarctic Ice Sheet”, *Antarctic Journal of the U.S.*, 20(5), pp 11-12, 1985.
- [3.26] Faure, G., Taylor, K.S., and Jones, L.M., “Hydrothermal Calcite in the Elephant Moraine”, *Antarctic Journal of the U.S.*, 21(5), pp 21, 1986.

#### Geology

- [3.27] Dietrich, R.V. and Skinner, B.J., *Rocks and Rock Minerals*, John Wiley & Sons, Inc, 1979.
- [3.28] Blatt, H., Jones, R.L., “Proportions of Exposed Igneous, Metamorphic, and Sedimentary Rocks”, *Geological Society of America Bulletin*, v. 86, pp 1085-1088, August 1975.

## Bayes Network Training

The previous chapter described the Bayes network classifier topology but left the conditional densities of the data vectors unspecified. These need to be learned from example data. This process is often referred to as network training. Network re-training is when estimates of the conditional densities are further refined using additional training data.

This chapter describes the Bayesian learning algorithm used to estimate the conditional densities, which was taken from [4.1]. It only works for estimate discrete probability distributions; therefore all the random variables in the network, corresponding to continuous sensor feature vectors, need to be suitably discretized. This introduces additional problems of populating the conditional probability matrices given the limited training data. This is solved by generating more training examples by perturbing the original ones, similar to the approach of [4.2].

Classification results of the trained network on laboratory data are presented.

### 4.1. Bayesian parameter learning

Consider the parameter  $\bar{\theta} = (\theta_1, \dots, \theta_M)$  representing the proportions of the different states ( $1 \dots k$ ) of a population drawn from some probability distribution, and let the (discrete) random variable  $\mathbf{y}$  be a single observation from that distribution. Thus  $\bar{\theta}$  is also a random variable with the properties

$$(Eqn 4-1) \quad \theta_1, \dots, \theta_M \in [0,1], \quad \sum_{i=1}^M \theta_i = 1$$

and

$$(Eqn 4-2) \quad P(\mathbf{y} = k | \bar{\theta}) = P(\mathbf{y} = k | \theta_k) = \theta_k$$

It follows that

$$(Eqn 4-3) \quad \begin{aligned} P(\mathbf{y} = k) &= \int P(\mathbf{y} = k | \theta_k) P(\theta_k) d\theta_k \\ &= \int \theta_k P(\theta_k) d\theta \\ &= E\{\theta_k\} \end{aligned}$$

Learning the density  $p(\bar{\theta})$  from examples therefore solves the problem of learning the probability distribution  $P(\mathbf{y})$ . Using Bayes rule we get

$$(Eqn 4-4) \quad p(\bar{\theta} | \mathbf{y} = k) \propto \theta_k p(\bar{\theta})$$

Thus, given a prior distribution for  $\bar{\theta}$ , a posterior density can be calculated given a new independent observation. Consider the case where  $\mathbf{Y}$  consists of multiple (independent identically distributed) observations, in which each state  $k$  is observed  $(\alpha_k - 1)$  times. The posterior density becomes

$$(Eqn 4-5) \quad p(\bar{\theta} | \mathbf{Y}) \propto \theta_1^{\alpha_1 - 1} \theta_2^{\alpha_2 - 1} \dots \theta_M^{\alpha_M - 1} p(\bar{\theta})$$

If the prior distribution  $p(\bar{\theta})$  is assumed to be uniform, then the posterior density is immediately recognized as the *Dirichlet* distribution [4.3]

$$(Eqn 4-6) \quad p(\bar{\theta} | \mathbf{Y}) = \frac{\Gamma(\alpha_1 + \dots + \alpha_k)}{\Gamma(\alpha_1) \dots \Gamma(\alpha_k)} \theta_1^{\alpha_1 - 1} \dots \theta_k^{\alpha_k - 1}$$

$$(Eqn\ 4-7) \quad E\{\theta_j | \mathbf{Y}\} = \frac{\alpha_j}{\alpha_0} \quad \text{var}\{\theta_j | \mathbf{Y}\} = \frac{\alpha_j(\alpha_0 - \alpha_j)}{\alpha_0^2(\alpha_0 + 1)}$$

where

$$(Eqn\ 4-8) \quad \alpha_1 \dots \alpha_k > 0$$

$$\alpha_0 \equiv \sum_{i=1}^k \alpha_i$$

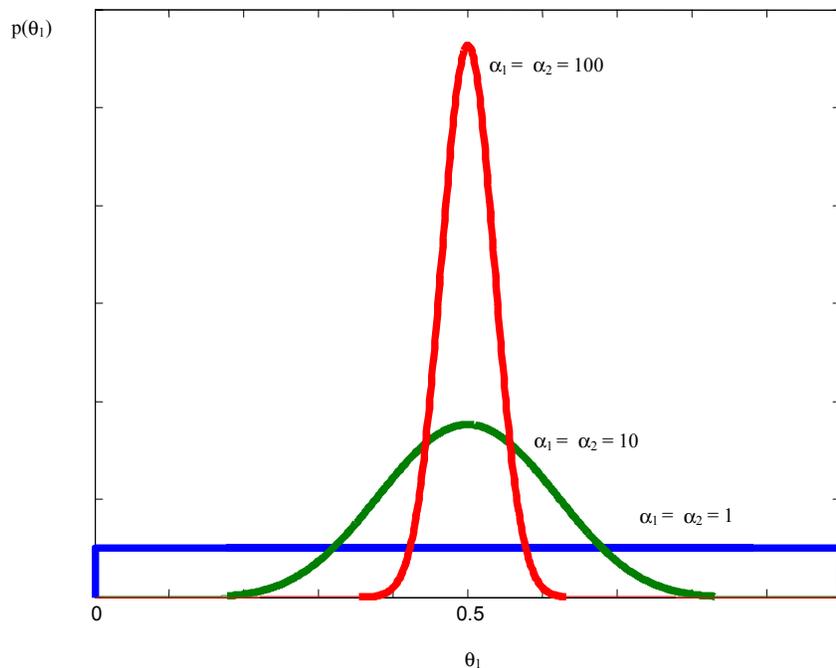
The Dirichlet distribution is fully characterized by the parameters  $\alpha_1 \dots \alpha_M$ , sometimes referred to as the “sample sizes”. Furthermore, Dirichlet distributions are a *conjugate* class of distributions in the sense that if the prior distribution is Dirichlet ( $\alpha_1 = \alpha_2 \dots = \alpha_M = 1$  corresponds to a uniform distribution), then so is the posterior density:

$$(Eqn\ 4-9) \quad p(\bar{\theta}) = \text{Dirichlet}(\bar{\theta}; \alpha_1, \dots, \alpha_k, \dots, \alpha_M) \Rightarrow p(\bar{\theta} | \mathbf{y} = k) = \text{Dirichlet}(\bar{\theta}; \alpha_1, \dots, \alpha_k + 1, \dots, \alpha_M)$$

Therefore

$$(Eqn\ 4-10) \quad E\{\theta_k | \mathbf{y} = k\} = \frac{\alpha_k + 1}{\alpha_0 + 1}, \quad E\{\theta_{j \neq k} | \mathbf{y} = k\} = \frac{\alpha_j}{\alpha_0 + 1}.$$

By interpreting  $\bar{\theta}$  as a “prior on the prior”  $P(\mathbf{y})$ , and assuming it has a Dirichlet distribution, it is simple to learn the likely values of  $P(\mathbf{y})$  from examples using equations (Eqn 4-3), (Eqn 4-9) and (Eqn 4-10). It is standard practice to initially assume  $\alpha_1 = \alpha_2 \dots = \alpha_M = 1$  (i.e. a uniform initial prior on  $P(\mathbf{y})$ ), but other values may be used to reflect earlier training, or other constraints. Note that if the sample sizes are large, then subsequent training will have little effect on the distribution (Figure 4.1).



**Figure 4.1** Dirichlet distribution’s corresponding to different sample sizes.

In order to learn the conditional probability matrices we assume independence between the conditional probabilities given by different conditional antecedents, and subsequently learn each one by example using

the assumption that it is a Dirichlet distribution. This algorithm is described in [4.1] and implemented by the Netica [4.4] Bayes network program.

Note that example data can be assigned a weight  $w$ , in which case (Eqn 4-9) becomes

$$(Eqn 4-11) \quad \text{If } p(\bar{\theta}) = \text{Dirichlet}(\bar{\theta} ; \alpha_1, \dots, \alpha_k, \dots, \alpha_M)$$

Then

$$p(\bar{\theta} | \mathbf{y} = k; w) = \text{Dirichlet}(\bar{\theta} ; \alpha_1, \dots, \alpha_k + w, \dots, \alpha_M)$$

A prior with  $\alpha_1 = \alpha_2 \dots = \alpha_M = 0$  is the uninformative prior, using it with example data set  $\mathbf{Y}$  gives

$$(Eqn 4-12) \quad E\{\theta_k | \mathbf{Y}\} = \frac{\alpha_k}{\alpha_0}$$

This is the simple ratio of specific cases  $k$  to the total number of cases. However, this estimate is unstable, especially if there is a limited number of example cases with which to train, and some of true probabilities are small. Using it is equivalent to assuming that the probability of heads is 1 given a single coin toss and subsequently observing George Washington's (or Queens) head.

### 4.1.1. Learning probabilities from uncertain data

Consider the problem of learning  $\bar{\theta}$  given an uncertain observation of  $\mathbf{y}$ , of the form  $\mathbf{U} = [P_1, \dots, P_M]$  representing the probabilities that  $\mathbf{y}$  has state  $k \in \{1, \dots, M\}$ .

$$(Eqn 4-13) \quad p(\bar{\theta} | \mathbf{U}) = \sum_{k=1}^M p(\bar{\theta} | \mathbf{y} = k) P(\mathbf{y} = k | \mathbf{U}) \\ = \sum_{k=1}^M p(\bar{\theta} | \mathbf{y} = k) P_k$$

If  $p(\bar{\theta})$  is a Dirichlet distribution

$$(Eqn 4-14) \quad p(\bar{\theta}) = \text{Dirichlet}(\bar{\theta}; \alpha_1, \dots, \alpha_k, \dots, \alpha_M) \\ \Rightarrow \\ p(\bar{\theta} | \mathbf{U}) = \sum_{k=1}^M P_k \text{Dirichlet}(\bar{\theta}; \alpha_1, \dots, \alpha_k + 1, \dots, \alpha_M)$$

This implies that  $p(\bar{\theta} | \mathbf{U})$  is no longer a Dirichlet distribution. It is a mixture model of distinct Dirichlet distributions. In fact,  $p(\bar{\theta} | \mathbf{U})$  and  $p(\bar{\theta})$  are in general no longer conjugate densities, regardless of the prior density  $p(\bar{\theta})$ .

More tractable is an observation  $\mathbf{O}$  characterized by  $[Q_1, \dots, Q_M]$  where  $Q_k = P(\mathbf{O} | \mathbf{y} = k)$ . Therefore

$$(Eqn 4-15) \quad P_k = p(\mathbf{y} = k | \mathbf{O}) = \frac{P(\mathbf{O} | \mathbf{y} = k) P(\mathbf{y} = k)}{\sum_j P(\mathbf{O} | \mathbf{y} = j) P(\mathbf{y} = j)} \quad (\text{Bayes rule}) \\ = \frac{Q_k E\{\theta_k\}}{\sum_j Q_j E\{\theta_j\}}.$$

(Eqn 4-14) becomes

$$(Eqn 4-16) \quad p(\bar{\theta} | \mathbf{O}) = \sum_{k=1}^M \frac{Q_k E\{\theta_k\}}{\sum_j Q_j E\{\theta_j\}} P(\bar{\theta} | \mathbf{y} = k)$$

Once again, if  $p(\bar{\theta})$  is a Dirichlet distribution (c.f. (Eqn 4-14)), then  $p(\bar{\theta} | \mathbf{O})$  is not a Dirichlet distribution, but a mixture model of Dirichlet distributions. It violates the conjugacy requirement and is intractable as more observations are made. Subsequent observations would produce mixture models with  $M^2$ ,  $M^3$ , and so on terms. This shows that it is in general not feasible to rigorously learn a distribution from uncertain observations.

A common, non-rigorous hack is to approximate

$$(Eqn 4-17) \quad p(\bar{\theta} | \mathbf{U}) \approx \text{Dirichlet}(\bar{\theta}; \alpha_1 + P_1, \dots, \alpha_M + P_M)$$

This is equivalent to computing the posterior distribution of  $\bar{\theta}$  given the definite observations  $\{\mathbf{y}=k | k=1 \dots M\}$ , each weighted according to their probability  $P_k = P(\mathbf{y}=k)$ .

The implication is that the conditional densities for a hidden node in a Bayes network (Figure 4.2) cannot be rigorously learned from example data unless that data also specifies the values of the hidden variable. Therefore, to learn all the densities in the rock classifier Bayes network of the previous chapter, it is

necessary that the training set of sensor features specify the exact rock type of each rock sample, not just whether it is a meteorite or not.



Figure 4.2 Bayes network with hidden node Y.

## 4.2. Quantization of continuous parameters and the “Curse of dimensionality”

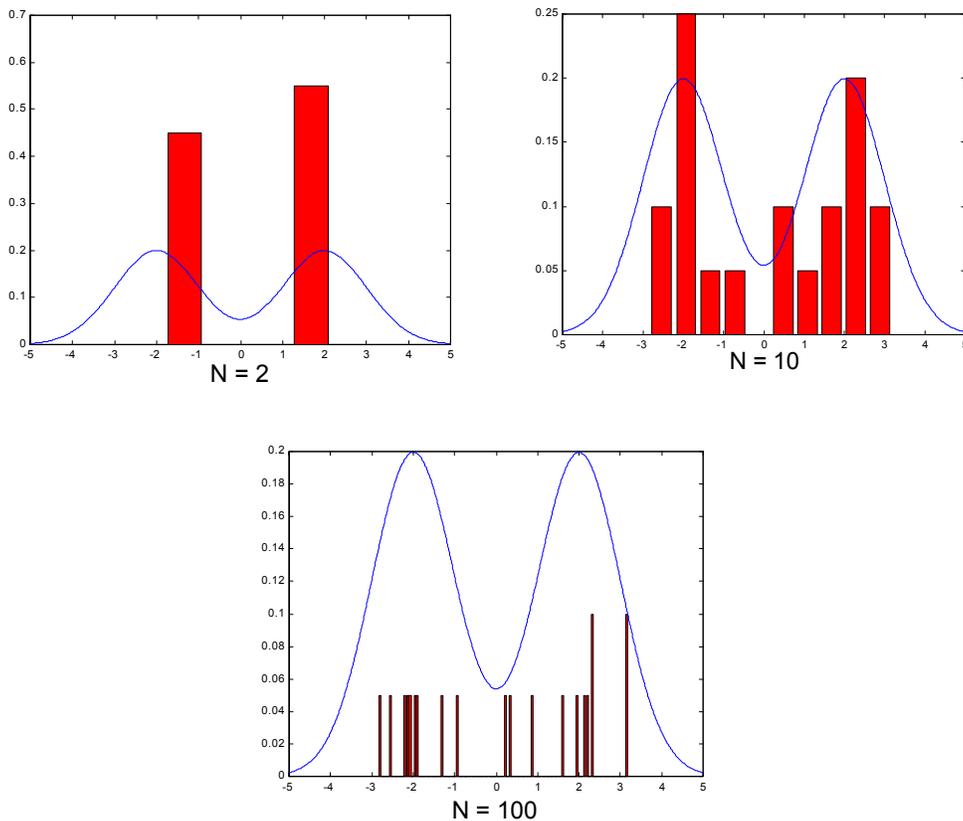
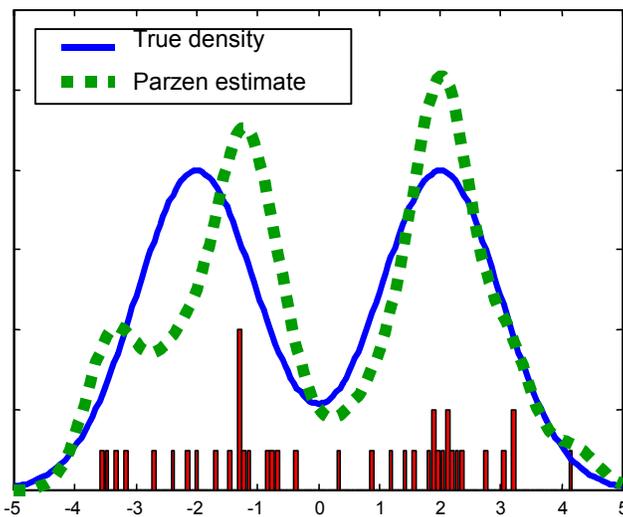


Figure 4.3 Approximating continuous random variable densities by discrete distributions. With limited training cases, learning these distributions becomes difficult as the number of states  $N$  of the discrete distribution is increased, leading most of the discrete states to be unpopulated. In this example, the continuous density is a mixture of two Gaussians, from which 20 samples are used to generate histograms with different numbers of bins ( $N$ ).

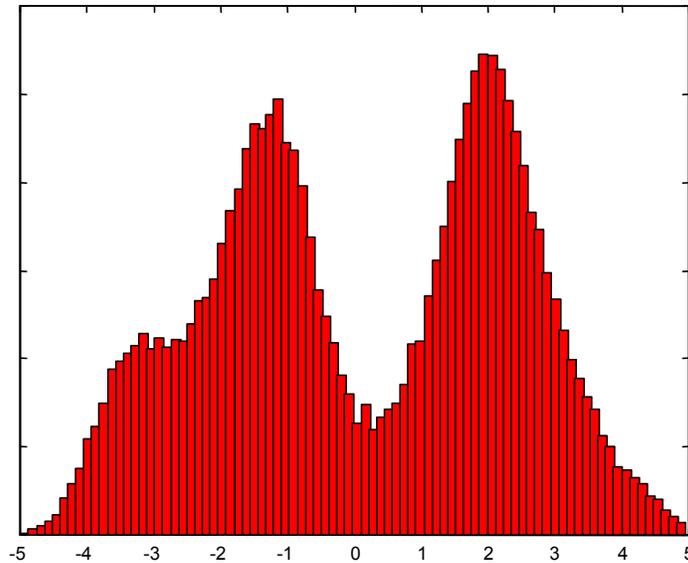
Most Bayes network implementations, including the one used here, are unable to deal with continuous variables. In addition, the method outlined above for learning probabilities applies only to discrete data. For these reasons it is necessary to quantize any continuous random variables. But doing this leads to a trade-off. The finer the quantization, the more representative the Bayes network model is, but at the expense of greater complexity and memory usage. Most serious however, is the limitations on network training. In order to learn the (conditional) probabilities for each state, sufficient examples are needed so as to populate all possible states. As the number of states increases, the learned density from a fixed number of samples deteriorates (c.f. Figure 4.3). In general, if a random variable has  $P$  parent, and  $N$  is the typical number of states, then  $O(N^{P+1})$  example data points are needed. This is essentially a restatement of the so-called “curse of dimensionality”, relating the number of samples needed to populate a  $P+1$  dimensional space. If it is known that the continuous density to be approximated varies smoothly, then a technique analogous to Parzen estimation [4.5] can be used to “fill in” the gaps. In Parzen estimation, a continuous density is inferred from a finite set of samples by summing a set of *kernel* densities, each one centered on a corresponding sample data point. The kernel densities are chosen to reflect local variations in the density, and are typically concentrated in the center and tail off with increasing distance. Gaussians are commonly used.



**Figure 4.4** Parzen density estimate, using 20 samples shown by histogram, shown by 100 bin histogram, from true density. Kernel functions are Gaussians with 0.5 standard deviation. The histogram shows how the sample points would populate a 100 state, uniform interval, discrete approximation of the continuous distribution.

The Parzen density estimate based on a few samples can be used to resample the distribution, thus generating sufficient samples to populate all states, and reliably learn the probabilities of each using Dirichlet update. Indeed, this does not even require the computation of the full Parzen estimate. It is sufficient to generate more samples by perturbing each original sample slightly according to some noise model (analogous to a Parzen kernel function).

It is important that the new simulated samples be weighted so that their combined weight is equal to the weighting of the original sample. Otherwise, undue confidence will be given to the learned distribution, as it would seem that many more samples were available than was the case.



**Figure 4.5** Histogram (100 bins) of points re-sampled from Parzen estimate in Figure 4.4. 50 new points are generated from each original sample point by adding random noise from the Parzen kernel function (a zero mean, 0.5 standard deviation Gaussian). This process interpolates between the original samples and produces a histogram more accurately reflecting the original distribution.

### 4.2.1. Sensor noise models

To generate multiple examples of the (continuous) sensor feature vectors given a single sensor measurement a perturbation or noise model is needed. Assume a single sensor reading  $S$ , from which a feature vector  $\bar{\mathbf{D}} = [\mathbf{d}_1 \dots \mathbf{d}_N]$  is computed. The goal is to perturb, or otherwise add noise to,  $\bar{\mathbf{D}}$  in a manner consistent with the local density of  $\bar{\mathbf{D}}$ . There are two options:

- Perturb  $\bar{\mathbf{D}}$  directly. At it's simplest this is adding independent Gaussian noise to each component of  $\bar{\mathbf{D}}$ . More sophisticated is adding noise from complex noise models. The most complicated is some non-linear perturbation of  $\bar{\mathbf{D}}$  using noise generated by a process accurately representing experimental and other variations in the sensor signal  $S$ .
- Perturb sensor signal  $S$  and re-compute  $\bar{\mathbf{D}}$ . Given a sensor noise model (additive or otherwise), the process by which signal variations are reflected in the feature vector  $\bar{\mathbf{D}}$  is taken care of. Assuming such a sensor model is available this is preferable to the former option. Such a sensor noise model was developed for the spectrometer, and is described in detail in Chapter 5.

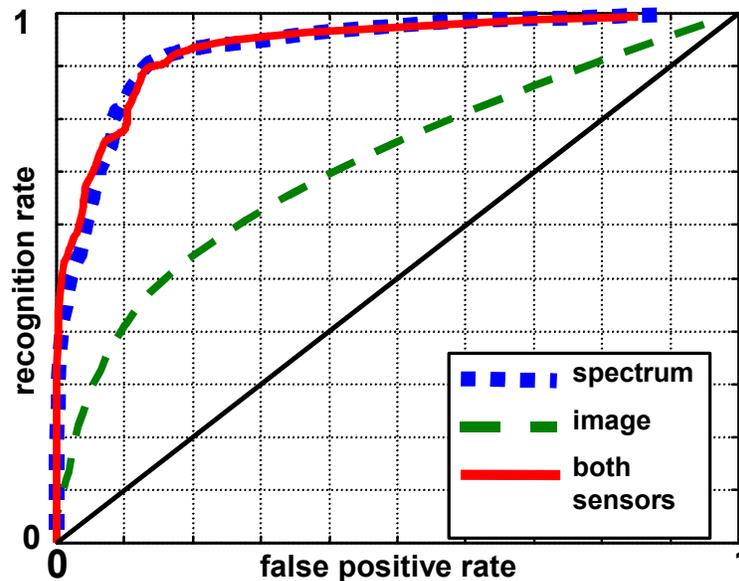
## 4.3. Laboratory training results

Training spectra and color images were acquired from Antarctic terrestrial rock and meteorite specimen collections at Ohio State University and the Johnson Space Center (where all U.S. Antarctic meteorites are kept). This data was acquired in the laboratory using a handheld digital camera and the Nomad spectrometer. Although every effort was made, field conditions could *not* be replicated. Prior to going to the field with the full robot, it was not known how the field conditions and the autonomous placement of the sensors would affect the data.

The rock and meteorite samples were chosen to include all rock and meteorite types likely to be encountered at Elephant Moraine in Antarctica. However, it was not necessary for the sample sets to be statistically unbiased, since the class conditionals are learned independently for each rock or meteorite type. However, for each rock type, the data set used to learn the conditionals for that rock type must be unbiased. As varying numbers of spectra were acquired from each rock sample, it was necessary to weigh each spectral sample to enforce this.

This data was used to train the Bayes network of the previous chapter. To gauge classifier effectiveness the dataset was partitioned into separate training (80% of samples) and test (20% of samples) set, trained on the former and tested on the latter. Repeating this several times it is possible to test the classifier on the whole data set, without ever training and testing on the same data, using a technique known as cross validation.

Recall that the classifier computes a probability of a rock sample being a meteorite given sensor measurements. This probability relies on the possibly arbitrary values for the rock type priors, so it is not necessarily the “true” probability that a rock sample is a meteorite. It might better be interpreted as a classifier confidence. A sample is labeled as a meteorite if the classifier confidence exceeds some decision threshold. Classification errors occur when a meteorite is not recognized, or a terrestrial rock is falsely identified as a meteorite. Choosing a decision threshold requires a trade-off between the meteorite recognition rate (fraction of times a meteorite is recognized as such) and the false positive rate (fraction of terrestrial rocks labeled as meteorites). A high decision threshold means that only samples with a high classifier confidence will be labeled as meteorites. Therefore, a low recognition and false positive rate can be expected. Conversely, a low decision threshold means that most meteorites will be recognized, but at the expense of picking up many terrestrial rocks.



**Figure 4.6** ROC curve of classifier performance recognizing meteorites in laboratory training data, using 20% cross validation.

An ROC curve is a plot of a classifier recognition rate versus false positive rate. It indicates the recognition rate versus false positive rate trade-off that must be made for any binary classifier. It represents

classification performance independently of the decision threshold or the assumed class priors. A random classifier (“head’s, it’s a meteorite, tails, it’s a rock”) has equal recognition and false positive rates regardless of the priors or decision threshold. It is the least useful of all classifiers. The best possible classifier allows a 100% recognition rate for a 0% false positive rate. Its ROC curve is a step function.

The ROC curve of Figure 4.6 shows the performance of the Bayes network rock and meteorite classifier on the laboratory data set after training for different combinations of sensor data. Assuming a 20% maximum acceptable false positive rate, this classifier recognizes 55% of the meteorites using color images only, and 95% if the spectrometer is used.

Note the superior classification with spectra versus classification from color images only. In fact, given a spectrum, a color image does not provide extra information. Nevertheless, because an image is easier to acquire than a spectrum a camera has value for the preliminary assessment of samples.

#### 4.4. References

- [4.1] Spiegelhalter, D.J., Philip, D.A., Lauritzen, S.L., and Cowell, R.G., “Bayesian analysis in expert systems” in *Statistical Science*, 8(3), p219-283, 1993.
- [4.2] Pomerleau, D., “ALVINN: An Autonomous Land Vehicle In a Neural Network”, in *Advances in Neural Information Processing Systems 1*, Morgan Kaufmann, 1989.
- [4.3] Gelman, A., Carlin, J.B., Stern, H.S. and Rubin, D.B., *Bayesian Data Analysis*, Chapman & Hall, 1995.
- [4.4] Netica, <http://www.norsys.com>
- [4.5] Fukunaga, K., *Statistical Pattern Recognition*, 2<sup>nd</sup> ed, Academic Press, 1990.
- [4.6] Jordan, M.I., “Why the Logistic Function? A Tutorial Discussion on Probabilities and Neural Networks”, *Computational Cognitive Science Technical Report 9503*, August 1999.

## **Part II**

# **Sensors and Sensor Deployment**

## Spectroscopy

Reflection spectroscopy is the primary sensor on Nomad for identifying rocks, and determining if they are meteorites. It has been used extensively by the geological community to identify minerals in rocks, either remotely or in the laboratory [5.1][5.2][5.3][5.4][5.5], making it a logical choice for robotic rock classification. The JPL Rocky-7 prototype planetary rover and the Athena instrument payload package for the next Mars rover both feature reflection spectrometers.

A reflection spectrometer works by illuminating a rock sample with white light and measuring the fraction reflected at a range of wavelengths. The shape and intensity of the reflection spectrum is indicative of the minerals present in the rock, and therefore what type of rock it is and whether it could be a meteorite. The spectral signatures of most minerals are beyond the near IR portion of the EM spectrum [5.1]. However, previous experiments with a visible to IR spectrometer established that the visible to near IR (400 nm to 1000 nm) region is sufficient to distinguish many rocks from meteorites, even though discrimination between terrestrial rock types is poor. Doing spectroscopy beyond the near IR is considerably more expensive and difficult than in this region.

Reflection spectroscopy is one category of many different types of spectroscopy. Other forms of spectroscopy used to analyze minerals and rocks are thermal emission spectroscopy (TES), Raman spectroscopy, and laser induced breakdown spectroscopy (LIBS).

Thermal emission spectroscopy involves measuring the thermal infrared emission of an object, usually over a wavelength range of 6 – 80 microns. Characteristic sharp emission peaks at various wavelengths in the spectrum identify mineral species. Furthermore, the spectra from an aggregate of minerals (such as a rock) are relatively easy to deconvolve, being a linear mix of the component mineral spectra. Emission spectra can be acquired from a distance with telescopic optics, and do not rely on the external illumination of a sample. Because of this, and the relative ease with which they can be interpreted, TES is often the instrument of choice for robotic geology. Unfortunately, they are very expensive and delicate instruments, costing up to several million dollars for a space-qualified instrument. They also require cryogenic cooling and long signal integration times.

Raman spectroscopy works by illuminating a sample with intense monochromatic (laser) light, stimulating the sample to re-emit light at wavelengths shifted by some characteristic amount from the wavelength of the exciting light source. The spectra obtained are very effective for determining the chemical composition of a sample, provided the correct excitation source is used. Unfortunately, the returned signal can be extremely faint, making it difficult for a robot to acquire it autonomously in the field. Only very recently, have robust portable Raman spectrometers become widely available.

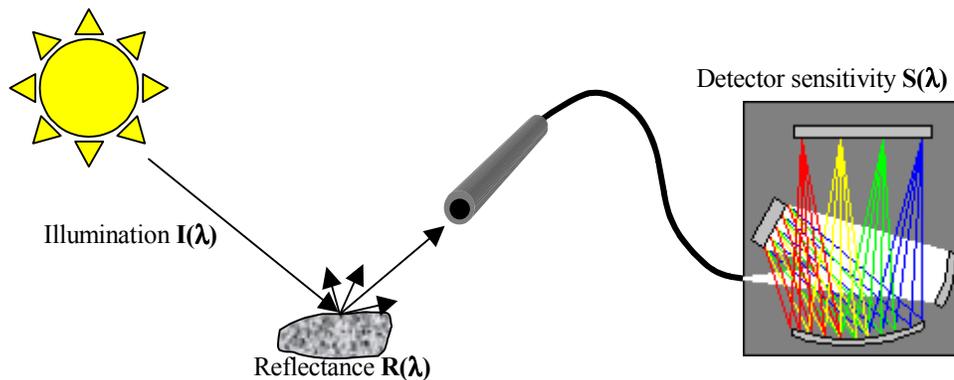
Laser induced breakdown spectroscopy is the most promising sensor modality for robotic rock classification. It works by vaporizing a minute piece of a sample with a high-powered laser pulse. The vaporized sample is subsequently ionized and made to emit a characteristic spectrum. Very high signal to noise ratio's are possible, so these measurements can be done at a distance (approximately 10 m) and in direct sunlight. The resulting spectra are characterized by sharp emission peaks at discrete wavelengths, and are a linear combination of the spectra from the constituent components of the sample. It is therefore

relatively easy to interpret. Although apparently an ideal sensor, a unit suitable for robotic deployment has only recently been developed, and it was not available for experimentation.

Reflection spectra are the hardest to interpret of all the spectra. They have the least sharp peaks, and are affected by the physical as well as chemical properties of a rock. Significantly, the reflection spectrum of a rock is not necessarily a linear combination of the spectra of its constituent minerals. Depending on grain size and other factors the resultant spectra can be a non-linear function of the constituent mineral spectra. To make matters worse, certain minerals such as iron oxide, will dominate the spectrum of a rock, even if only present as trace contaminants. A further complication is that, through weathering, many rocks acquire a coating of “desert varnish” that obscures the true nature of the rock.

However, in spite of their drawbacks, reflection spectrometers are still effective at detecting many different chemical species in a sample, and can characterize many common rock types. They are widely used instruments for industrial processes, such as sorting plastics, and are robust and affordable instruments available “off the shelf”. This is extremely important for the robotic Antarctic meteorite search project, which operated on an aggressive timeline. Furthermore, conditions in Antarctica are extreme and logistical support is minimal, so instruments need to be reliable and require a minimum of management.

## 5.1. Principles of reflectance spectroscopy



**Figure 5.1** Typical reflectance spectrometer setup. A sample is illuminated by some light source  $I(\lambda)$ , of which  $I(\lambda)R(\lambda)$  is reflected. The reflected light is collected by an optical apparatus and directed to a diffraction grating that splits it into the component wavelengths, directing each one to a detector.

Reflection spectroscopy works by measuring the fraction of light,  $R(\lambda)$ , reflected by a rock sample at different wavelengths  $\lambda$ . Figure 5.1 shows the typical arrangement for carrying out this measurement. A sample is illuminated by a light source with intensity  $I(\lambda)$  and some of the reflected light is directed to a diffraction grating (or prism) which separates it into its component wavelengths (colors). A bank of detectors measures the intensity of the light at each wavelength. If  $S(\lambda)$  is the combined detector and spectrometer optics sensitivity, the measured signal  $x(\lambda)$  is given by

$$(Eqn 5-1) \quad x(\lambda) = I(\lambda)R(\lambda)S(\lambda) + D(\lambda)$$

$D(\lambda)$  is the ambient DC offset signal, due thermal noise in the detector and ambient illumination not accounted for by the light source  $I(\lambda)$ . Switching off the light source (or covering the spectrometer light

collector) allows this to be measured. However, both the illumination  $I(\lambda)$  and detector sensitivity  $S(\lambda)$  are hard to measure independently, necessary if (Eqn 5-1) is to be solved for a direct measurement of  $R(\lambda)$ .

The solution is to calibrate the spectrometer and light source with a standardized reflectance target with a known reflectance spectrum  $R_{ref}(\lambda)$ . If  $x_{ref}(\lambda)$  is the measured spectrometer signal from the reference target, then by subtracting out the dark current from both signals and taking their ratio's, the illumination and detector sensitivity terms cancel out:

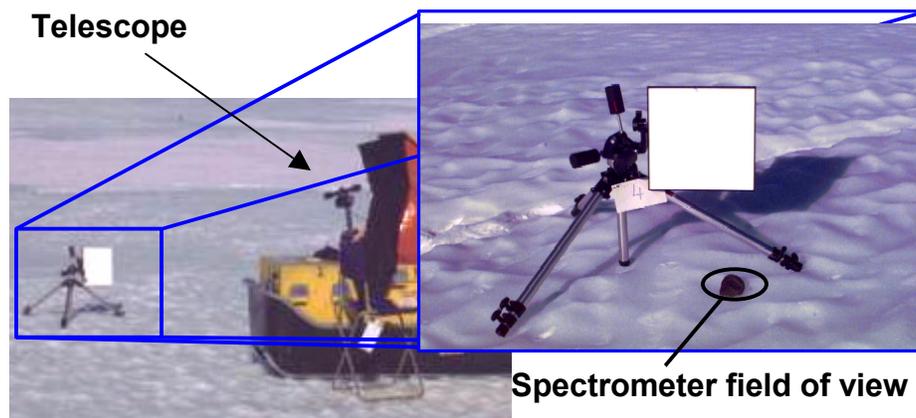
$$(Eqn\ 5-2) \quad \frac{x(\lambda) - D(\lambda)}{x_{ref}(\lambda) - D(\lambda)} = \frac{R(\lambda)}{R_{ref}(\lambda)}$$

The most common reference or calibration target is a substance with the trade name Spectralon™. It is almost uniformly white from the visible wavelengths to well into the infrared spectrum, and is a lambertian reflector.

To accurately measure reflectance spectra this way requires that:

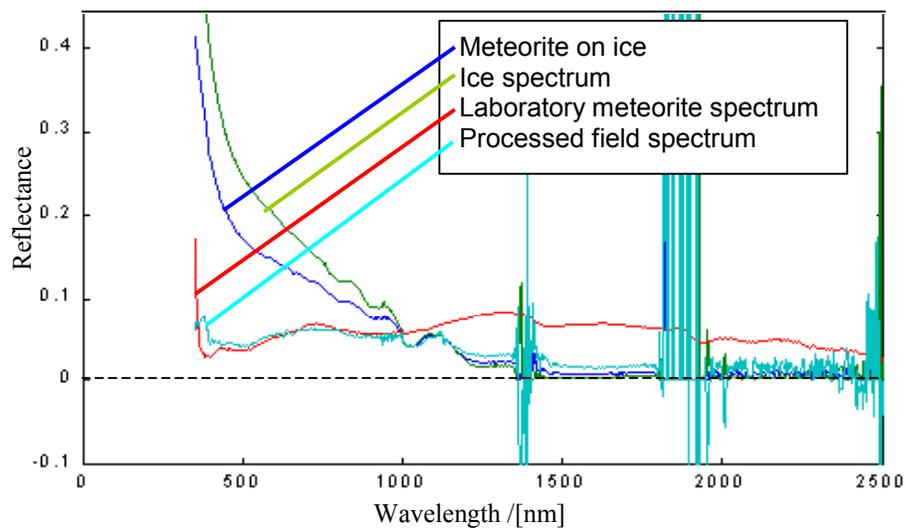
- Both rock sample and calibration target be exposed to the same light source. Furthermore, the relative geometry between the spectrometer light collector, the viewed section of the rock sample or calibration target and light source must be kept constant. This complicates using the sun as a light source if a calibration target cannot be placed next to the rock sample. Altering this geometry leads to a scaling error in (Eqn 5-2).
- The calibration target must be kept clean. Fortunately, this is not a problem in Antarctica, which is an exceptionally clean environment. It could be a problem in desert environments, such as Mars.
- The DC offset term  $D(\lambda)$  must be properly accounted for. If artificial illumination is used, the offset due to uncontrolled ambient illumination can be significant. The uncontrolled lighting must either be excluded or properly measured.

## 5.2. Autonomous spectroscopy in the field



**Figure 5.2** Experimental arrangement for acquiring spectra remotely using a telescopic fore optic attached to the spectrometer fiber optic feed. Note the large white spectral calibration target placed near the rock sample, and the large field of view of the telescope.

Accurately measuring the reflectance spectra of a rock in the field with an autonomous robot is challenging. It is preferable to acquire spectra from rocks remotely as this saves time and reduces complexity. This can be done using sunlight to illuminate the samples and a telescopic fore optic to collect the reflected light. A 100% white reflectance standard is needed to calibrate the system and determine the incident sunlight. The calibration target should be close to the sample and oriented similarly with respect to the sun. Such an instrument was tested on rocks and meteorites on the Antarctic ice sheets in 1997 and the results were not encouraging for small rock samples the size of typical meteorites (Figure 5.3). It is difficult to obtain a field of view small enough to exclude light from the surrounding ice. The uncontrolled shape of the rocks and their position with respect to the sun and the spectrometer fore-optic cause uneven illumination and shadows on the areas visible to the spectrometer. Accurately aiming the telescope is difficult, and frequent recalibrations are necessary because of the transient nature of solar and other ambient illumination.



**Figure 5.3** *The problems with remotely acquired spectra of rocks on the ice sheet. A telescope is used to collect light for spectral analysis. The signal from the ice overwhelms that from the rock or meteorite sample. If the background ice spectrum is known, it can be subtracted from the combined signal to deduce the rock spectrum with only limited success.*

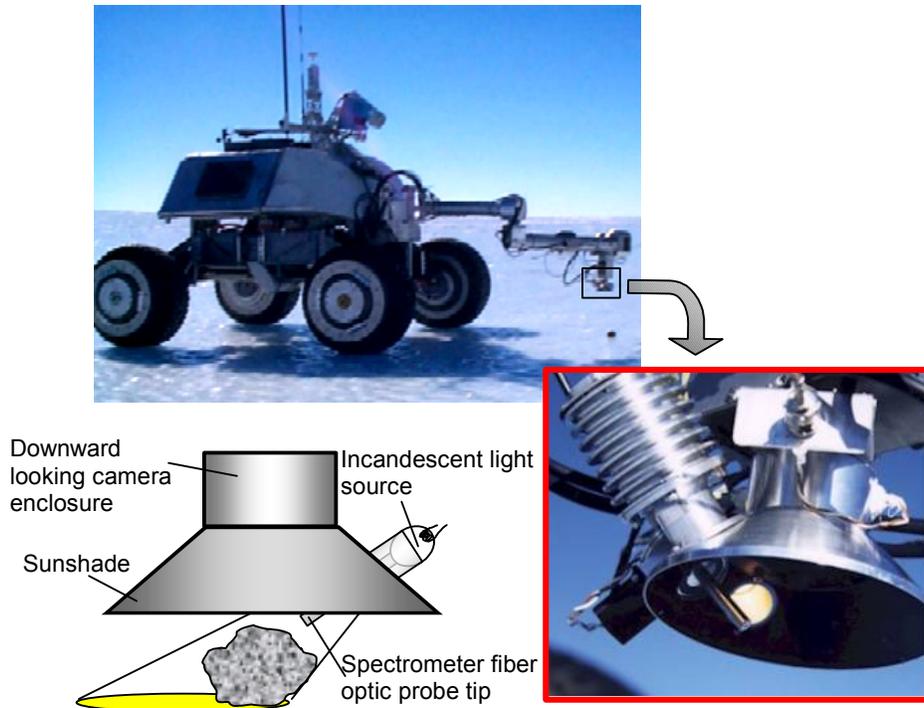
Because of these difficulties, it is necessary to acquire spectra using a fiber optic probe placed directly over rock samples to collect the reflected light. This requires a robot manipulator arm [SAS/Wagner ref] with suitable sensors and software to autonomously place the probe, not an easy task to do with sub-centimeter precision.

A special sensor head (Figure 5.4) was designed to cope with the unpredictable and intense ambient illumination. Two incandescent projector bulbs intensely illuminate the area of the rock sample visible to the probe. A sunshade, painted black on the inside, reduces ambient lighting of the sample. The lights are arranged to provide uniform lighting from the probe to 2 cm below and horizontally.

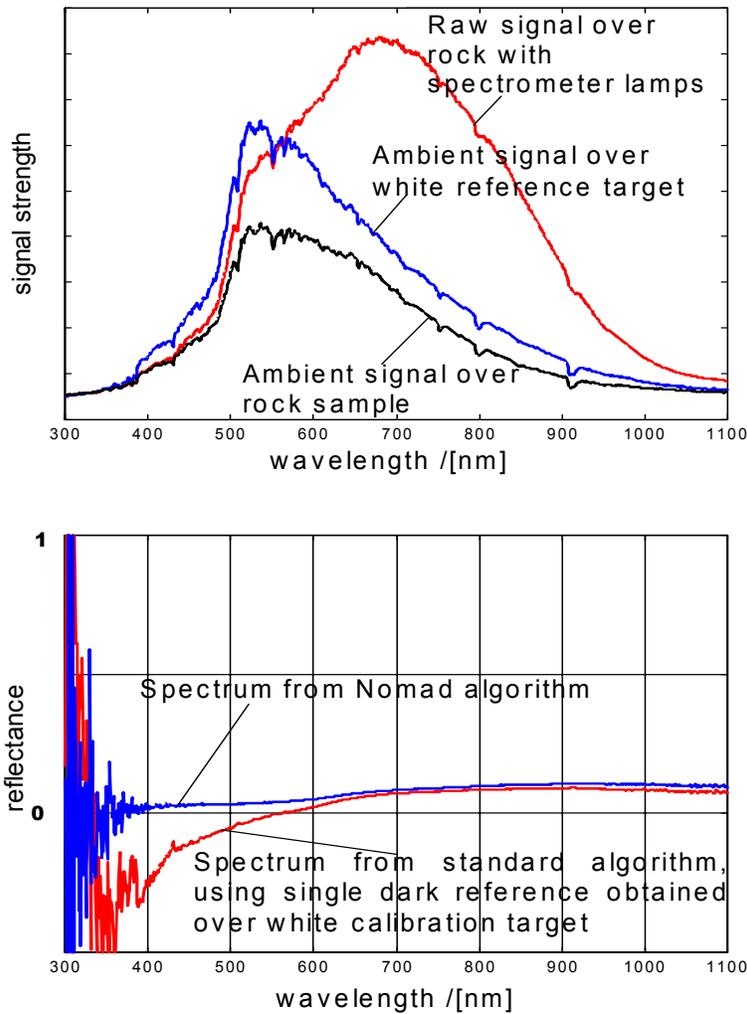
The ambient lighting, along with any instrument signal offset, is measured by taking spectral measurements with the lights off. The effect of ambient lighting is thus accounted for and eliminated without any complicated light tight seal over the probe and sample area.

The uniform lamp lighting reduces the required probe placement precision to approximately 1 to 2 cm. A downward looking camera above the probe allows horizontal positioning, while a contact switch indicates when the assembly is in contact with a rock sample.

Calibration is simple with this manipulator-based actively illuminated system. A white reference target is placed near the home position of the robot manipulator arm. Prior to measuring a new sample, the fiber-optic probe is placed over the reference target and measurements taken with the lamps on and off. The difference is computed and a similar measurement performed on the rock sample. The reflectance spectrum is computed by dividing the raw sample difference by the white reference difference.



**Figure 5.4** *Spectrometer sensor head, deployed by manipulator arm onto rock samples on ice.*



**Figure 5.5** Top: Unprocessed spectral signal from a sample rock with lamps illuminated, and spectral signals with lamps off when probe is over the white reflectance standard and over the rock. Note the strong signal due to ambient sunlight. Bottom: Computed reflectance spectra using Nomad's spectrometer algorithm and a traditional algorithm. Negative values for spectrum are nonsensical. The high noise below 400 nm is due to the low UV illumination from the spectrometer lamps.

### 5.3. Spectrometer noise model

The need for a spectrometer noise model for network training purposes was indicated earlier (Chapter 4). Given a spectrum  $R(\lambda)$  of a rock sample, the goal is to perturb it slightly to generate a new spectrum  $R^*(\lambda)$  that is consistent with the sensor and the rock sample from which the original spectrum was obtained.

The simplest perturbation is to add random Gaussian noise to the spectrum acquired by the Nomad. More realistic is adding the noise to the raw spectral data used to compute the reflectance spectrum (Eqn 5-2):

$$(Eqn 5-3) \quad R^*(\lambda) = \frac{x(\lambda) - D(\lambda) + \varepsilon_1(\lambda)}{x_{ref}(\lambda) - D(\lambda) + \varepsilon_2(\lambda)}$$

Where  $\varepsilon_1(\lambda), \varepsilon_2(\lambda) \sim N(\mu = 0, \sigma)$  are additive white noise processes.

For the purposes of this model the white reference measurement  $x_{ref}(\lambda)$  and dark measurement  $D(\lambda)$  can be considered constants for all spectral measurements. Using this, and the fact that the noise processes are uncorrelated, (Eqn 5-3) can be rewritten as

$$(Eqn\ 5-4) \quad R^*(\lambda) = R(\lambda) + \varepsilon(\lambda)$$

Where

$$(Eqn\ 5-5) \quad \varepsilon(\lambda) \sim N(\mu = 0, \sigma_R(\lambda))$$

is a Gaussian noise process, and

$$(Eqn\ 5-6) \quad \sigma_R^2(\lambda) \approx \left[ \frac{\sigma}{x_{ref}(\lambda)} \right]^2 \left[ 1 + \frac{1}{x_{ref}^2(\lambda)} \right]$$

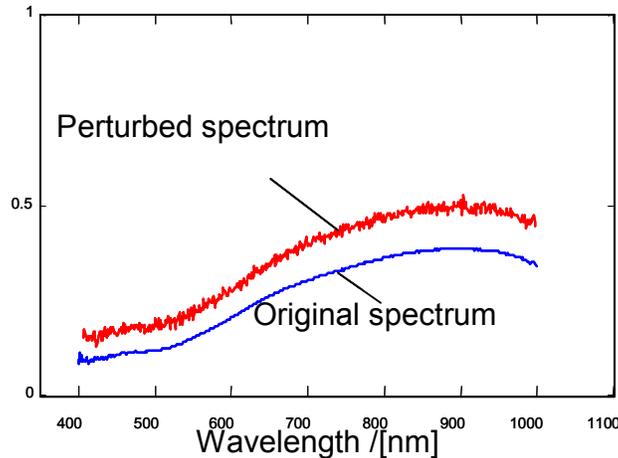
using a first order Taylor series approximation.

Note that now additive noise is no longer constant across  $\lambda$ . It results from the two noise processes in (Eqn 5-3) and depends on the degree to which the sample is illuminated at each wavelength. If the sample is not illuminated, then a poor signal is returned and there is a low signal to noise ratio, which is consistent with this noise model. Note that the Nomad spectrometer lamps emit most energy in the visible wavelengths, and the spectrometer CCD sensor is most sensitive in the visible wavelength. Therefore,  $x_{ref}(\lambda)$  will be small in the UV and IR wavelengths, giving a very noisy spectra in these regions (Figure 5.5), and reflected in the noise model output (Figure 5.6).

Recall, furthermore, that spectral measurements are subject to significant scaling errors because of the uncontrolled geometry of rock surfaces. Therefore, it is consistent to add a random scaling factor to the perturbation equation (Eqn 5-4):

$$(Eqn\ 5-7) \quad R(\lambda)^* = \Lambda [R(\lambda) + \varepsilon(\lambda)]$$

where  $\Lambda \sim N(\mu = 1, \sigma_\Lambda)$ , truncated to fit in interval [0.5, 1.5], is the random scaling factor.



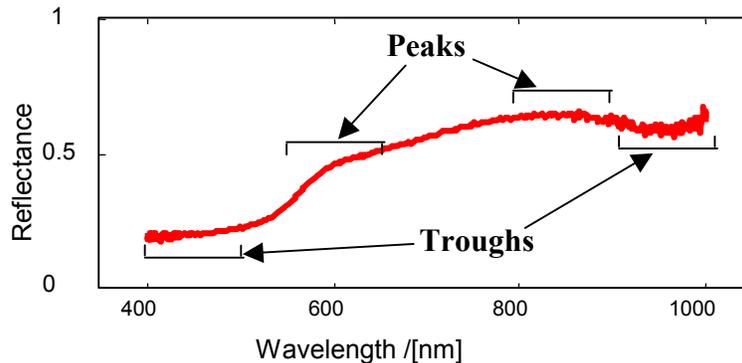
**Figure 5.6** Rock spectrum (hydrothermally altered red Basalt) acquired by robot, and perturbed spectrum used for network training purposes.

## 5.4. Spectral feature extraction

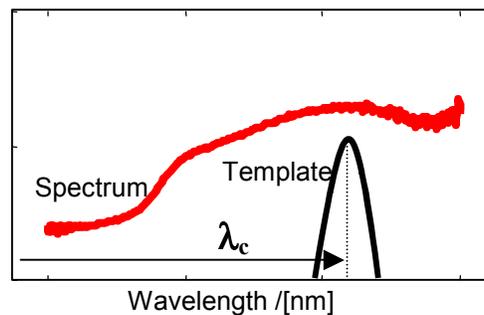
Attempting to statistically model a complete spectrum for all rock types is both intractable and unnecessary. A typical spectrum may consist of a 1000+ reflectance measurements at different wavelengths. In this section we describe features that characterize the spectra of different rocks, yet are relatively robust to sensor artifacts and have limited statistical interdependencies as required by the Bayes network topology described in Chapter 3.

### 5.4.1. Spectral peak and trough detection

Peaks and troughs (Figure 5.7) are what characterize the spectra of different minerals, and therefore different rocks. Automatically detecting them in spectra acquired in the field presents difficulties. Clearly simple thresholding is not suitable. [5.4] attempt to remove the continuum parts of the spectrum, leaving only the troughs (corresponding to absorption peaks), which they locate, by either thresholding or examination of the first and second derivatives. Once troughs (absorption peaks) are detected, they measure the center, width and depth. Unfortunately, while their methods work on laboratory data extending from the visible to mid IR (8  $\mu\text{m}$ ), they are less effective on field-measured spectra extending into the near IR only. The signal noise, the troughs are generally not well defined and measured spectra tend to be improperly scaled. Furthermore, because there are fewer absorption bands in these wavelengths, the continuum peaks (Figure 5.7) assume added importance for classification. They are not detected by the above methods.



**Figure 5.7** Reflectance spectral peaks and troughs corresponding to a measurement from an Iron Oxide stained Basalt rock in Elephant Moraine, Antarctica



**Figure 5.8** Template function  $T(\lambda_c)$  for detecting peaks and troughs in a spectrum using normalized correlation.

The normalized cross correlation between a spectrum and a suitable template (Figure 5.8) is a better way of detecting peaks and troughs centered at some wavelength  $\lambda_c$ :

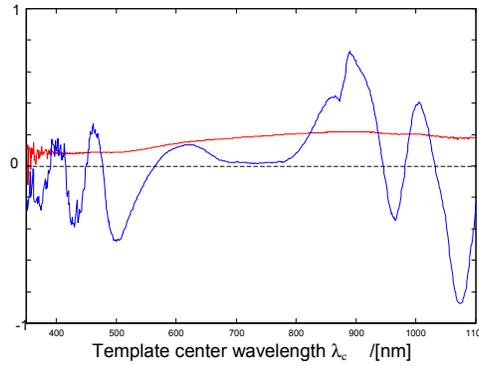
$$(Eqn 5-8) \quad c = \frac{C_{ST}}{\sigma_S \sigma_T}$$

$$C_{ST} = \int_{-W/2}^{W/2} [R(\lambda - \lambda_c) - \mu_S][T(\lambda - \lambda_c) - \mu_T] d\lambda$$

$$\sigma_S^2 = \int_{-W/2}^{W/2} [R(\lambda - \lambda_c) - \mu_S]^2 d\lambda, \quad \mu_S = \int_{-W/2}^{W/2} R(\lambda - \lambda_c) d\lambda$$

$$\sigma_T^2 = \int_{-W/2}^{W/2} [T(\lambda - \lambda_c) - \mu_T]^2 d\lambda, \quad \mu_T = \int_{-W/2}^{W/2} T(\lambda - \lambda_c) d\lambda \quad (\text{Both constant for a given template.})$$

$R(\lambda)$  is the reflectance spectrum,  $T(\lambda)$  is the template function centered at 0 and with width  $W$ . The normalization coefficients cancel out, so they are ignored. Figure 5.9 shows the normalized correlation of a truncated Gaussian shaped template with a typical spectrum. Note the positive correlation values near spectral peaks and negative correlation values near troughs.

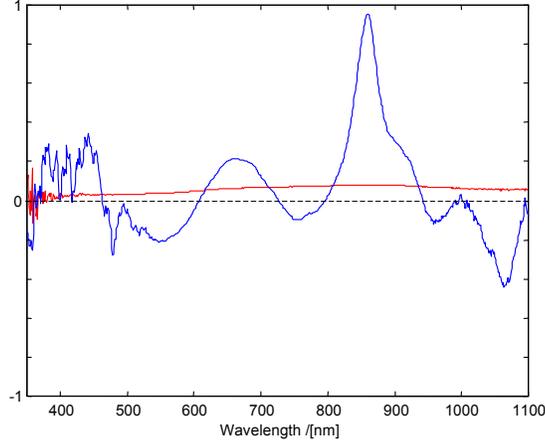


**Figure 5.9** Normalized correlation of template function in Figure 5.8 with a rock spectrum as the template center wavelength  $\lambda_c$  is varied.

Detecting peaks and troughs with this measure has several advantages:

- The correlation coefficient  $c$  is insensitive to spectrum scale and offset. This is important for classification as field measured spectra often differ from true reflection spectra by a random scale factor.
- Both spectral peaks and troughs are detected.
- Varying the width of the templates optimizes the algorithm for the detection of different width peaks and troughs.
- The output is bounded, with 1 representing a definite peak, -1 a definite trough, and numbers in between lesser degrees of correlation.

A problem with this peak and trough metric occurs with very flat spectra. In such cases, the value of  $\sigma_S$  is very small, so the normalized correlation coefficient  $c$  (Eqn 5-8) is overly sensitive to even the smallest peak or trough (Figure 5.10). This follows directly from scale independence, since even the smallest deviations in the spectrum can be magnified sufficiently to appear like a large peak or trough.



**Figure 5.10** The problem with detecting peaks and troughs using normalized correlation, which is overly sensitive whenever the spectrum is flat. This is a direct consequence of the scale independence.

The ideal detector would be insensitive to peaks and troughs below a certain level, but otherwise scale invariant. Clearly here has to be some trade-off. Consider adding uncorrelated Gaussian noise  $\varepsilon(\lambda) \sim N(\mu = 0, \sigma)$  to the spectrum before computing the correlation coefficients. Substituting into (Eqn 5-8) get

$$(Eqn 5-9) \quad R'(\lambda) = R(\lambda) + \varepsilon(\lambda),$$

$$c' = \frac{C_{ST}}{\sigma_T \sqrt{\sigma_S^2 + \sigma^2}},$$

Then only peaks or troughs that are significant compared to the noise level are detected by (Eqn 5-9) yet scale invariance is still maintained for those features that are much bigger than the noise. However, adding the ‘virtual’ noise decreases the range of the correlation coefficients:

$$(Eqn 5-10) \quad |c'| \leq 1 / \sqrt{1 + \frac{\sigma^2}{\sigma_S^2}}$$

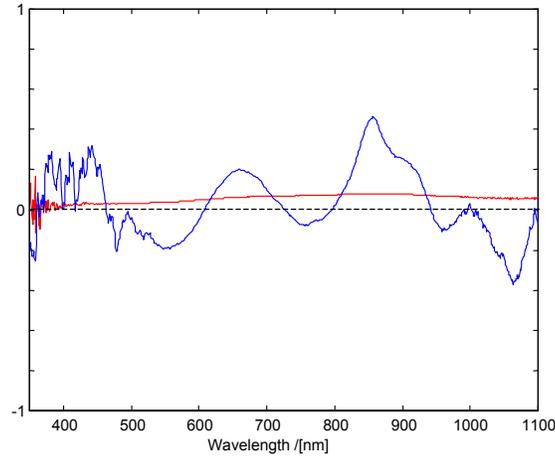
therefore

$$c' \rightarrow 0 \text{ as } \sigma \rightarrow \infty$$

It is desirable to rescale (Eqn 5-9) to keep the range approximately constant regardless of the virtual noise level. Multiplying by  $(\sigma+1)$  does the trick. Therefore, define the function V for gauging the strength of a peak or trough at a specific wavelength, given a suitable template function centered on that wavelength:

$$(Eqn 5-11) \quad V := \frac{C_{ST}(\sigma+1)}{\sigma_T \sqrt{\sigma_S^2 + \sigma^2}}$$

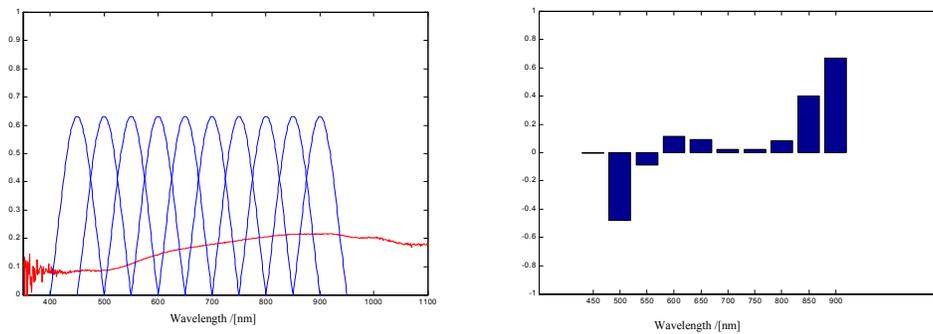
$$\lim_{\sigma \rightarrow \infty} V \in [-1, 1]$$



**Figure 5.11** Response of the function  $V(\text{spec}, \text{template})$  (Eqn 5-11) to the spectrum from Figure 5.10, using a virtual noise level of 0.02. Note the decreased response at wavelengths where the spectrum is flat and noise free, whereas response is unchanged in areas where the spectrum varies significantly more than the virtual noise.

Choosing a finite set (Figure 5.12) of templates  $\{T_1, \dots, T_N\}$ , define the feature vector  $(V_1, \dots, V_N)$  on a spectrum  $R$  by  $V_k = V(R; T_k)$ . With the right templates, this feature vector quantifies the peaks and troughs in a spectrum and therefore helps classify the rock that produced the spectrum in the first place.

Choosing the right templates directly impacts classifier performance. They must be sensitive to the peaks and troughs that discriminate the spectra of different rocks. These occur at different wavelengths and with different widths. However, to constrain the classification complexity it is necessary to minimize the number of templates. Furthermore, overlapping templates give values that depend on each other; even knowing which rock type produced the spectrum.



**Figure 5.12** Spectral peak and trough templates (left) set for computing a feature vector (right) quantifying the peaks and troughs in the example spectrum plotted with the templates. The feature vector,  $[V_1, \dots, V_{10}]$ , is shown as a bar plot on the right, with each element plotted against the central wavelength of its corresponding feature template. The example spectrum comes from an Iron Oxide stained Basalt, examined by Nomad in Antarctica. Note the strong negative correlation corresponding to the trough at 450 nm and the positive correlation around the peak at 900 nm.

Figure 5.12(left) shows the set of templates (along with an example spectrum), plotted against wavelength, used for the classification of rocks [5.6]. The feature vector  $[V_1, \dots, V_{10}]^T$ , computed by correlating these

templates with the example spectrum using the formula in (Eqn 5-11) is shown as a bar plot in Figure 5.12(right), with each component of the feature vector plotted against the corresponding template center wavelength for ease of interpretation. *Note: this format will be henceforth adopted for displaying spectra along with their template correlations. The same set of templates is used throughout.*

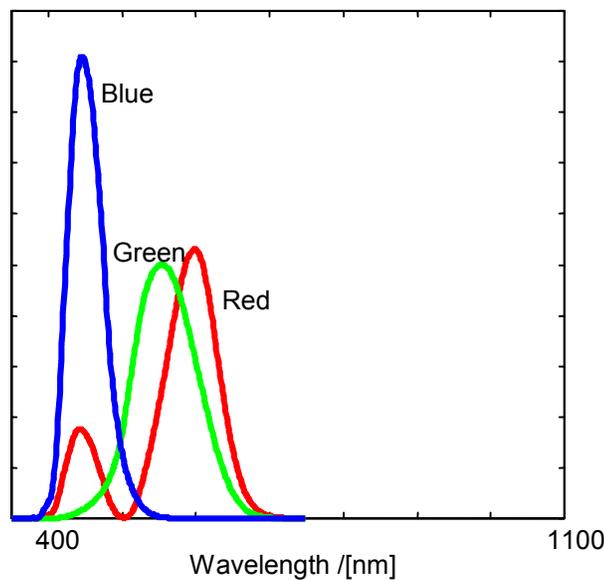
Since only adjacent templates overlap, it is safe to assume that non-adjacent template correlations are not directly statistically dependent, justifying the Bayes network classifier architecture of Figure 3.7. They cover the wavelengths from 400 – 900 nm. Outside this range, spectral measurements with the Nomad spectrometer have proven to be unreliable and noisy.

### 5.4.2. Spectral red, green and blue indices

By itself, the spectral feature vector  $(V_1, \dots, V_N)$  does not fully describe a spectrum for rock classification purposes. They are insensitive to the average reflectance over different wavebands, and do not indicate trends. In fact, they are insufficient to differentiate between a perfectly black and a perfectly white object. They are deliberately designed this way because of the scaling errors in measured spectra. Nonetheless, the errors are not so great that, if correctly accounted for by the Bayes network classifier, the average reflectance cannot be used. This is especially so for the recognition of meteorites which tend to be very dark objects with relatively flat spectra, while there are many rocks also with flat spectra but greater average reflectance (albedo).

Given a spectral response function  $F(\lambda)$ , such as the red, green and blue spectral response functions of the human eye or a typical CCD camera (Figure 5.13), the albedo  $A_F$  of a spectrum  $R(\lambda)$  in the waveband covered by  $F(\lambda)$  is given by

$$(Eqn 5-12) \quad A_F = \int_0^{\infty} F(\lambda)R(\lambda)d\lambda$$



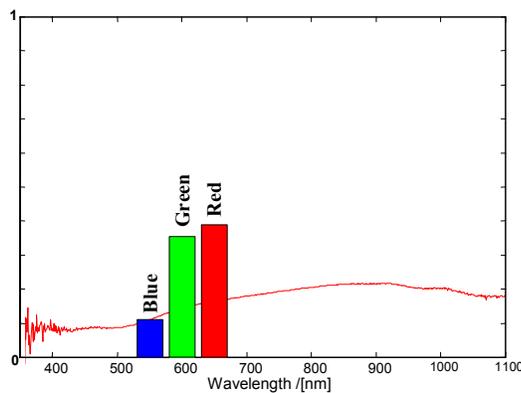
**Figure 5.13** Red, Green and Blue spectral response functions.

The albedos  $\{A_1, \dots, A_M\}$ , corresponding to response functions  $\{F_1, \dots, F_M\}$  indicate the spectrum signal level at different wavelength bands, and together indicate trends. However, they are not optimal for classification because they depend on each other. The indices

(Eqn 5-13)  $\{I, A_1/I, \dots, A_{M-1}/I\}$ , where  $I = A_1 + A_2 + \dots + A_M$

are preferable, since only  $I$ , representing overall spectral intensity, depends on the spectrum scale. The remaining terms are scale invariant, but indicate the relative spectral magnitude at different wavebands.

For the Bayes network rock classifier, the indices (Eqn 5-13) are computed using the CCD spectral response functions of Figure 5.13, giving *intensity*, normalized *red* and normalized *green*. These functions are chosen because they relate directly to the pixel colors in color image of a rock sample, as seen in the classifier Bayes network diagram (Figure 3.7). Note furthermore that all the uncertainty in the spectrum scale is fully accounted for by the *intensity* variable, further simplifying the Bayes network.



**Figure 5.14** Bar plot of normalized Red, Green and Blue indices for a typical rock spectrum, plotted against the center wavelength's of their respective spectral response functions Figure 5.13. Note: this display format henceforth be used to indicate the RGB indices of a spectrum without further explanation.

## 5.5. Conclusions

Given a spectrum from a rock sample, the feature vector  $[V_1, \dots, V_{10}, \text{Intensity}, \text{Red}, \text{Green}]^T$  is computed. This characterizes the peaks and troughs in the spectrum, the overall intensity and any trends. While not necessarily optimal for distinguishing rock types these features are designed to control statistical interdependencies as required by the Bayes network classifier topology.

Worthwhile work, perhaps using genetic algorithms, would be to determine optimal templates and response functions while constraining statistical dependencies.

## 5.6. References

[5.1] Hunt, G., "Spectral Signatures of Particulate Minerals in the Visible and Near Infrared", *Geophysics*, vol. 42, no.3, p. 501-513, April 1977.

[5.2] Clark, R.N. "Spectroscopy and Principles of Spectroscopy", *Manual of Remote Sensing*, A. Rencz, Editor, John Wiley and Sons, Inc. 1999.

- [5.3] Clark, R.N., T.V.V. King, M. Klejwa, G. Swayze, and N. Vergo, "High Spectral Resolution Reflectance Spectroscopy of Minerals", *Journal Geophysical Research*, 95, pp12653-12680, 1990.
- [5.4] Clark, R.N. and Swayze, G.A., "Mapping Minerals, Amorphous Materials, Environmental Materials, Vegetation, Water, Ice and Snow, and Other Materials: The USGS Tricorder Algorithm", *Summaries of the Fifth Annual JPL Airborne Earth Science Workshop*, January 23- 26, R.O. Green, Ed., JPL Publication 95-1, p. 39-40, 1995.
- [5.5] Clark, R.N., Swayze, G.A., Koch, C., Gallagher, A., and Ager, C., "Mapping Vegetation Types with the Multiple Spectral Feature Mapping Algorithm in both Emission and Absorption", *Summaries of the Third Annual JPL Airborne Geosciences Workshop, Volume 1: AVIRIS Workshop*. JPL Publication 92-14, 60-62, 1992.
- [5.6] Pedersen, L., Apostolopoulos, D., Whittaker, W., Roush, T. and G. Benedix "Sensing and data classification for robotic meteorite search" Proceedings of *SPIE Photonics East* Conference. Boston, USA, 1998.
- [5.7] Clark, R.N., Swayze, G.A., Gallagher, A.J., Gorelick, N., and Kruse, F., "Mapping with Imaging Spectrometer Data Using the Complete Band Shape Least-Squares Algorithm Simultaneously Fit to Multiple Spectral Features from Multiple Materials", *Proceedings of the Third Airborne Visible/Infrared Imaging Spectrometer (AVIRIS) Workshop*, JPL Publication 91-28, 2-3, 1991.
- [5.8] Clark, R.N., Gallagher, A.J. and Swayze, G.A., "Material Absorption Band Depth Mapping of Imaging Spectrometer Data Using a Complete Band Shape Least-Squares Fit with Library Reference Spectra", *Proceedings of the Second Airborne Visible/Infrared Imaging Spectrometer (AVIRIS) Workshop*. JPL Publication 90-54, 176-186, 1990.

## Color Image Processing

Autonomously distinguishing terrestrial rocks from meteorites in color images is difficult (Figure 6.1). Not only can they look alike, but also the bright polar environment reduces contrast and color balance. Nonetheless, a color camera is an important sensor because color images can be acquired from a distance. Unlike when using a spectrometer, the robot does not have to be maneuvered as extensively to bring a rock sample into view. A camera is a good screening sensor to initially evaluate objects from a distance so that those candidates obviously not interesting can be identified as such and bypassed in favor of samples that are more promising.



**Figure 6.1** Color image of meteorite (arrow) and terrestrial rocks on the ice, showing the possible difficulty of visually distinguishing between the two. Features that aid in doing so are the rock sizes, color, texture and shape. Meteorites tend to be several centimeters in diameter, dark colored with a uniform texture, and a rounded and smooth exterior. [Photo: ANSMET]

Meteorites are usually distinguished by their dark and uniform coloration. This is because their surfaces melt and burn upon atmospheric entry. Because of this, they also tend to have a rounded shape. Finally, most Antarctic meteorites are rarely larger than several centimeters in diameter, whereas terrestrial rocks are often quite large. These features do not distinguish meteorites from all terrestrial rocks. Basalts are also uniformly dark and weathered sedimentary rocks are well rounded. However, their purpose is to identify rocks obviously not meteorites, reserving more costly sensors such as the spectrometer to investigate the more difficult rocks.

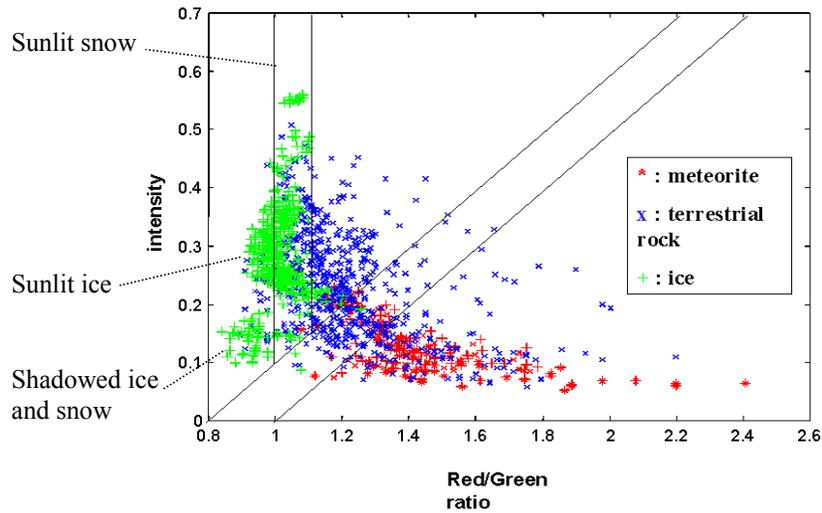
Given an accurate statistical model, even imperfect and inexact features can contribute evidence for discriminating between rocks and meteorites. This chapter shows simple metrics for quantifying the color, texture, shape and size of rock samples in color images that are used for classification.

### 6.1. Object/background segmentation

In order to extract features describing a rock in an image, the rock must first be segmented from the ice background. This is less simple than it would appear (Figure 6.2). Variable lighting conditions and shadows make simple thresholding on intensity impractical. Not all rocks are dark colored; some granites or calcites can be almost white, whilst the ice in the rock shadow can be very dark. Individually classifying

each pixel in an image as rock or ice, based on its intensity and color, is insufficient to determine the rock/ice boundary.

In color images, rocks are compact and connected objects, surrounded by the ice background. Furthermore, it is only necessary to determine the likely rock/ice boundary curve, not classify every single pixel. Segmentation does not need to be optimal, but simply good enough to allow the extraction of useful features that, on average, improve classification performance.



**Figure 6.2** Scatter plot of pixel observations from rocks, meteorites, snow and ice. The observed features are pixel intensity (scaled to  $[0,1]$ ) and Red/Green ratio. The intensity is sensitive to lighting conditions and is clearly not sufficient by itself to distinguish rocks from the ice background using a simple thresholding algorithm. The Red/Green ratio is more lighting independent, but still cannot distinguish all rocks from all types of ice or snow. Using both of these features, it is still necessary to use some shape based or continuity constraint to robustly segment rocks from ice in all circumstances.

### 6.1.2. Markov chain segmentation model

The camera system on Nomad is designed to return images similar to Figure 6.16, along with scaling information and the centroid of the rocklike pixels. It can be assumed that there are no rocks close together; so all rocks are fully surrounded by an ice background. If pixels are examined along a radial projecting out from the centroid, at some unique boundary distance from it they will transition from all being object (rock) pixels to all being background (ice) pixels. This sequence of pixels can be modeled by a Partially Observable Markov Chain (Figure 6.3). The sequence of states  $\{X_1 \dots X_M\}$  represent whether the sequence of pixels along a ray emanating from the object centroid are due to the object or the background. The boundary point  $k$  is that such that  $\{X_1 \dots X_k\}$  are all object, and  $\{X_{k+1} \dots X_M\}$  are all background. Given observations on each pixel in this sequence (intensity and Red/Green ratio), the probability of  $k$  being the boundary point is

$$(Eqn 6-1) \quad P(k | \bar{O}) = \frac{P(\bar{O} | k)P(k)}{P(\bar{O})}$$

Expanding the terms

$$(Eqn 6-2) \quad P(\bar{O} | k) = \prod_{i=1}^k P(O_i | X_i = object) \prod_{i=k+1}^M P(O_i | X_i = background)$$

$$(Eqn 6-3) \quad P(k) = \prod_{i=2}^k P(X_i = object | X_{i-1} = background) \times \\ P(X_{k+1} = background | X_k = object) \times \\ \prod_{i=k+2}^M P(X_i = background | X_{i-1} = background)$$

It follows that

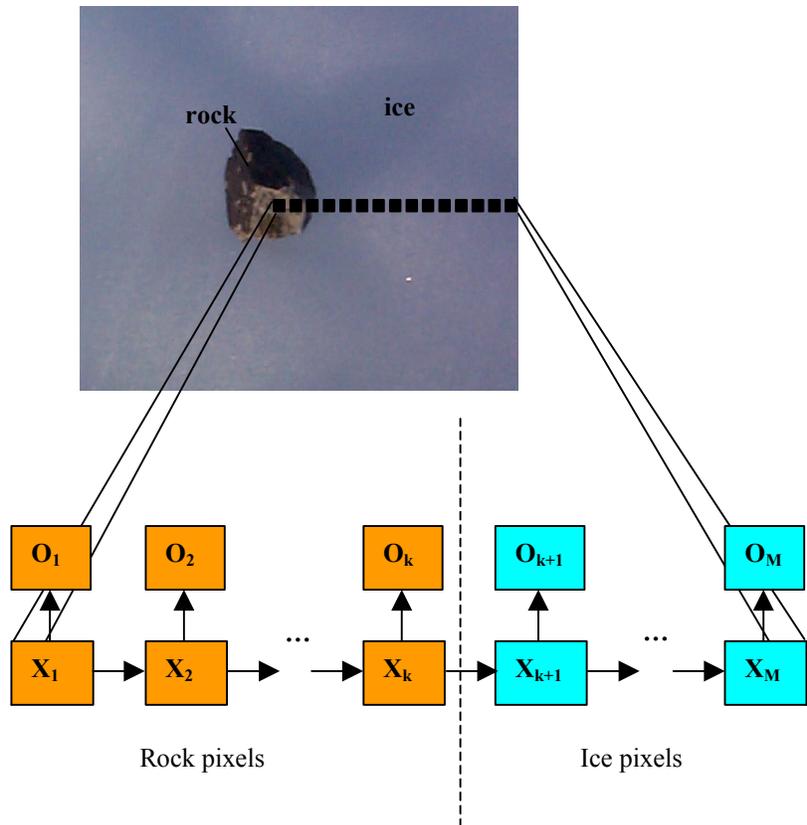
$$(Eqn 6-4) \quad \log[P(k | \bar{O})] = \\ \sum_{i=1}^k \log[P(O_i | X_i = object)] + \sum_{i=1}^k \log[P(O_i | X_i = object)] + \\ \sum_{i=2}^k \log[P(X_i = object | X_{i-1} = object)] + \log[P(X_{k+1} = background | X_k = object)] + \\ \sum_{i=k+2}^M \log[P(X_i = background | X_{i-1} = background)] - \log[P(\bar{O})]$$

The most probable boundary point is that which maximizes (Eqn 6-4). Note that  $\log[P(\bar{O})]$  does not depend on  $k$  and therefore does not affect the maximization.

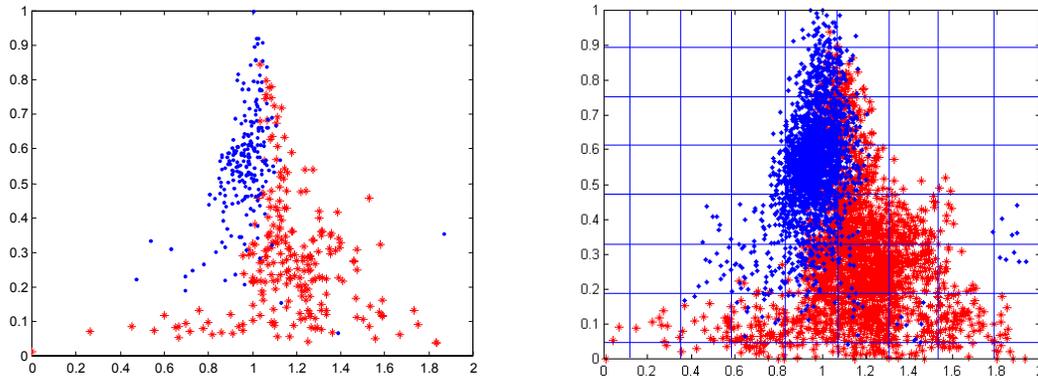
The conditional statistical relationship of pixel observations to pixel states  $P(O_i | X_i)$  learned from examples (Figure 6.4), using the Bayesian learning algorithm detailed in [6.1] and **Chapter 4**. Note that these probabilities depend significantly on the robot location and camera system, and must be learned anew for each location.

The conditional probabilities  $P(X_{i+1} | X_i)$  can also be learned this way or assigned a priori, since they are not expected to change much. It is reasonable to assume that they are the same for every  $i$ . Furthermore, if  $P(X_{i+1} = rock | X_i = rock) = P(X_{i+1} = ice | X_i = ice)$  for every  $i$ , then these terms are no longer a factor in maximizing (Eqn 6-4).

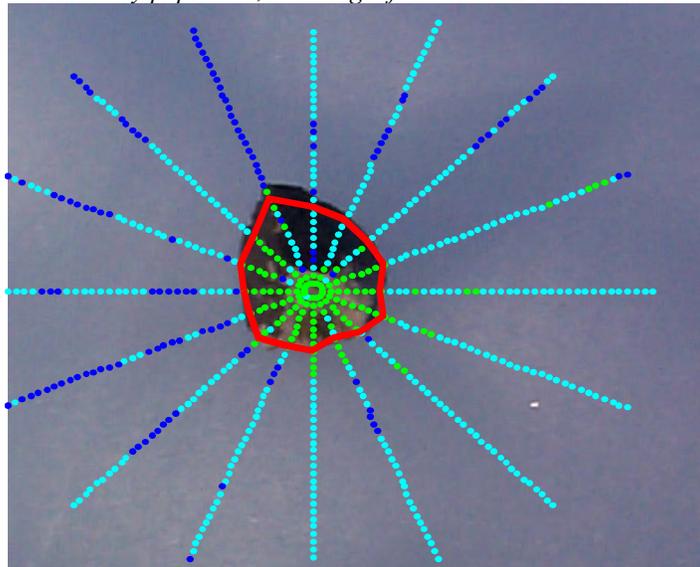
To compute the rock ice boundary many radials from the rock centroid are examined, and the optimal boundary point for each determined independently (Figure 6.5, Figure 6.6). Linking these boundary points produces a boundary polygon, with resolution determined by the number of radials and the size of the rock. For efficiency only points every 5 mm along the radials are examined, and each radial is limited to 30 cm (this is possible since the image scaling is known).



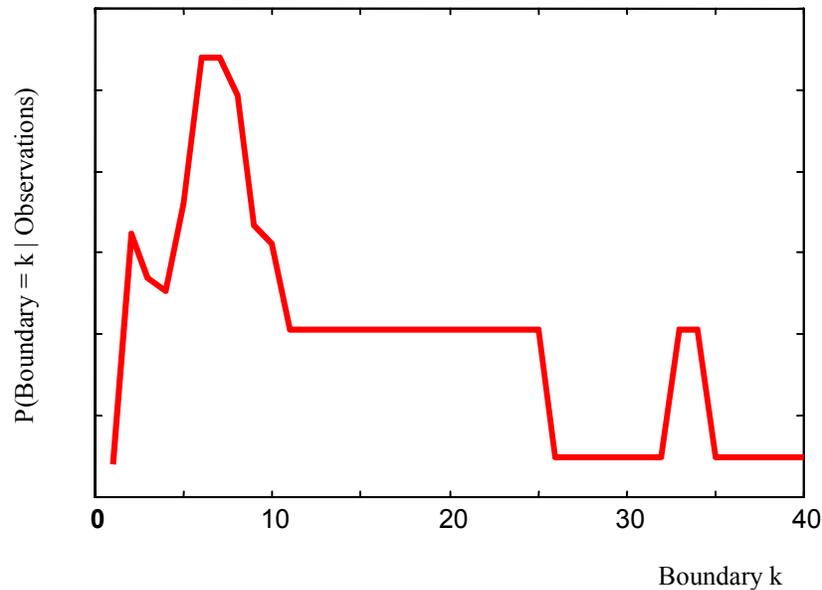
**Figure 6.3** Partially Observable Markov Chain model for segmenting a rock or meteorite from an ice or snow background in a color image.  $\{X_1 \dots X_M\}$  represent the state (rock or ice) of a line of pixels in the image, emanating from a point on the rock.  $\{O_1 \dots O_M\}$  represents observations on those pixels, in this case their red/green ratio and intensity. It is assumed that rocks are convex objects and sufficiently far apart from each other that there exists a single boundary value  $k$  such that  $\{X_1 \dots X_k\}$  are all rock, and  $\{X_{k+1} \dots X_M\}$  are all ice. Under this formulation finding the rock/ice boundary given observations on the pixels reduces to determining the most likely value of  $k$  (Eqn 6-4).



**Figure 6.4** Scatter-diagram of pixel intensity versus red/green ratio for rocks (red) and ice (blue) pixels in Antarctic images (left). By discretizing the observation space (right), with each grid representing a possible observation state, the conditional probability matrix  $P(O_i|X_i)$  is determined using the robust Bayesian learning algorithm of [6.1]. Adding noise to the samples to generate many more ensures that all observation states are reasonably populated, allowing a fine discretization.



**Figure 6.5** Computed rock/ice boundary from Figure 6.16, using the Markov chain model. Pixels, spaced every 5 mm along rays emanating from the rock centroid to a distance of 30 cm, are examined and the most likely boundary point determined for each radial. Only looking at pixels on radials is more efficient than looking at every single pixel in the image, and makes use of the rock compactness in determining the boundary curve.



**Figure 6.6** Likelihood of boundary along one of the radials in Figure 6.5. Note the clear peak around the actual boundary, and the tapering off at the extremes.

## 6.2. Image features extraction

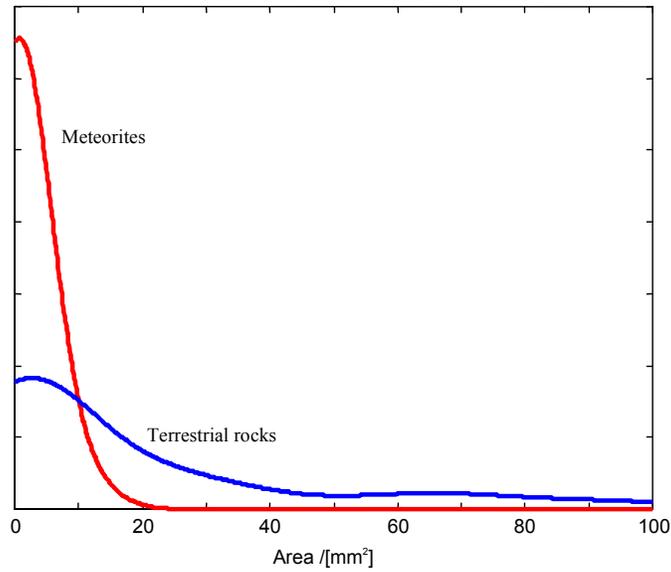
Features distinguishing meteorites from terrestrial rocks are, as previously mentioned, size, dark color, uniform texture and rounded shape. Given the rock/ice boundary curve, area and color are easy to determine. Texture and shape are harder to quantify.

Following is the set of features found to be both reasonable for classification and relatively easy to extract from segmented images. They are demonstrated on two data sets: Patriot Hills and JSC/OSU. The Patriot Hills data set consists of field images of rocks and meteorites taken in the Patriot Hills moraine, Antarctica, with a digital camera. Most of these rocks are sedimentary or metamorphous specimens, carried to the area by the ice flow. The meteorites in this set are not native to the area. They were planted there as no meteorites have been found in Patriot Hills. Nonetheless, they are representative of meteorites found in Antarctica. The OSU/JSC data set consists of laboratory images of Antarctic rock and meteorites collections at Ohio State University and Johnson Space Center respectively. The rocks are from the Elephant Moraine region of Antarctica and consist mostly of basalts and dolerites, dark colored igneous rocks. This set does not represent the distribution of rocks at Elephant Moraine, since only small rocks were collected shipped back (for obvious reasons).

- *Cross-sectional area*  $[mm^2]$

The cross sectional area of a rock indicates size of rock. It is computed using the image scaling information computed from the camera zoom values and the inferred distance of the camera to the sample.

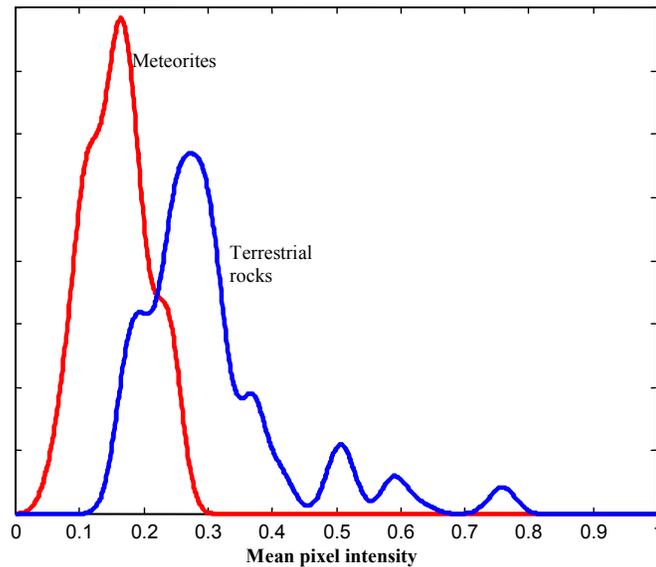
Most meteorites range in size from about 1 cm to 10 cm in diameter, while terrestrial rocks can be a lot larger, as seen with the Patriot Hills data set (Figure 6.7).



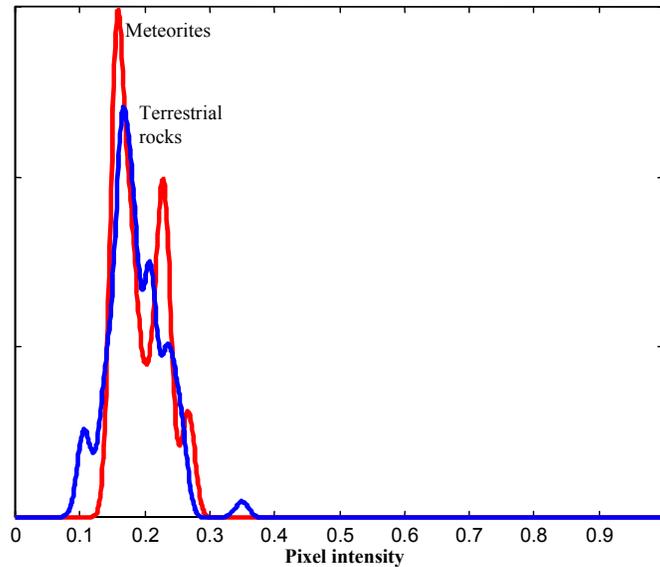
**Figure 6.7** Meteorite and terrestrial rock cross sectional area distribution, for rocks in the Patriot Hills field data set. Note the greater variation in rock size, and how meteorites tend to be smaller than the terrestrial rocks examined.

- Mean pixel color

Given the rock/ice boundary, the mean red, green and blue values for all the rock pixels (interior to the boundary) can be calculated. However, most information is encoded by the intensity  $I = R + G + B$ . The normalized color indices  $R/I$  and  $G/I$  measure the color, and have the advantage of not depending on the intensity of the ambient illumination.



**Figure 6.8** Meteorite and terrestrial rock pixel intensity distributions for Patriot Hills, Antarctica, field data set. Most of the terrestrial rocks are light colored marble, limestone and quartzite



**Figure 6.9** Meteorite and terrestrial rock pixel intensity distributions for JSC/OSU laboratory data set. Terrestrial rocks in this set contain many dark colored basalts and dolerites. Furthermore, different sources of illumination were used for the meteorites and the rocks.

Figure 6.8 shows a clear distinction in brightness between meteorites and terrestrial rocks in the Patriot Hills data set. However, the distinction is mediocre with the JSC/OSU laboratory data set (Figure 6.9), largely because this dataset consisted mostly of basalts and dolerites, igneous rocks themselves quite dark.

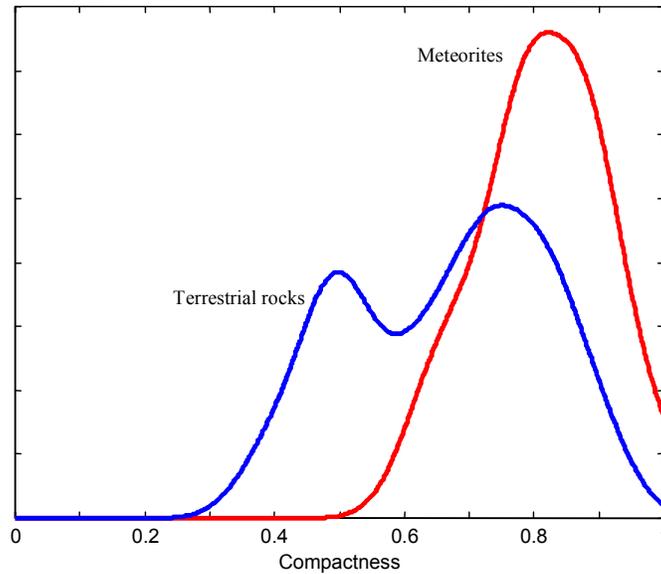
Different lighting conditions may also explain some of the discrepancies between Figure 6.8 and Figure 6.9, although the meteorite intensities are apparently consistent.

- *Compactness*

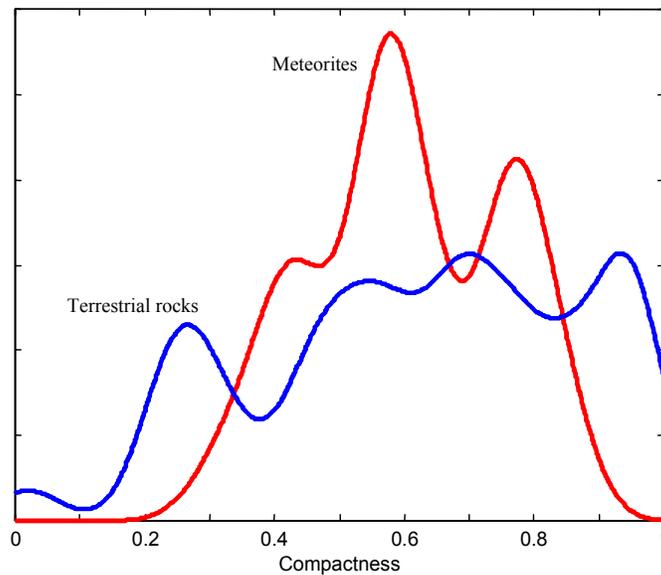
Compactness is a measure of the similarity of the rock boundary to a circle. It is the ratio of the rocks cross sectional area to the area of a circle with the same perimeter.

(Eqn 6-5) 
$$\text{Compactness} = \frac{4 \text{ Area}}{\text{Perimeter}^2}$$

A well rounded rock, such as an old and worn sedimentary sample, will have a compactness near one, whereas a jagged rock will be less compact. Figure 6.10 and Figure 6.11 indicates the measured compactness distributions for the laboratory and field data sets respectively. The laboratory data indicates a reasonable distinction between the measured compactness of rocks and meteorites, the field data less so. The difference in the rock distributions is unsurprising because of the differences in the rocks. However, compactness is sensitive to the accuracy with which the image of the rock is segmented from the background. In the field data the images of meteorites were sometimes cluttered, causing occasional errors in the segmentation, hence the discrepancy between the measured meteorite compactness in the two sample sets (which also consisted of different meteorites).



**Figure 6.10** Compactness distribution for rocks and meteorites in the JSC/OSU laboratory data set. Meteorites are more rounded while the terrestrial rocks tended to be angular, as reflected by the greater average compactness of meteorites.



**Figure 6.11** Compactness distribution for rocks and meteorites in the Patriot Hills, Antarctica, field data set. Many of the terrestrial rocks were highly eroded sedimentary rocks which are quite rounded, while others were recently broken, resulting in angular contours.

- *Bending Energy*

Bending energy is another measure of the angularity of the boundary [6.4], being analogous to the potential energy of a steel spring wrapped around the rock. Angular rocks will have a higher bending energy than smooth rounded ones.

(Eqn 6-6) 
$$BE = \frac{1}{Perimeter} \oint |\kappa(s)|^2 ds$$

where  $\kappa$  is the curvature of the boundary curve  $\bar{x}(s)$ , parameterized by arc length  $s$ :

$$(Eqn 6-7) \quad \kappa(s) = \left| \frac{de_1(s)}{ds} \right|$$

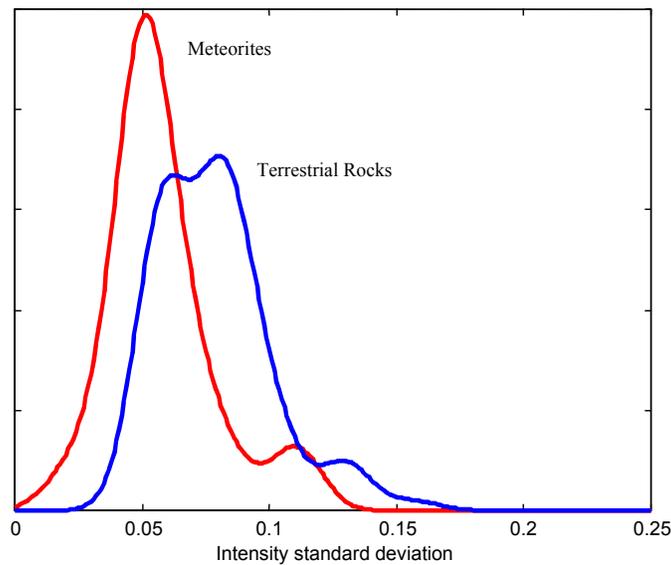
$$e_1(s) = \frac{d\bar{x}(s)}{ds}$$

In practice, bending energy has turned out to be too sensitive to the quality of segmentation to be an effective discriminator.

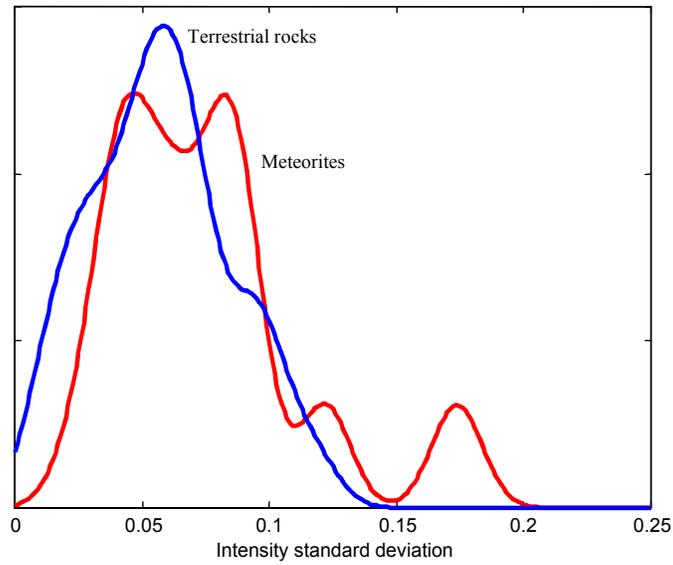
- *Standard deviation and entropy of pixel intensities*

Meteorites tend to have smooth surfaces without much texture or variation in color. Conversely, some terrestrial rocks, such as granite, have a speckled appearance. The goal here is a simple measure that may help exclude some rocks obviously not meteorites, using the minimal amount of computation.

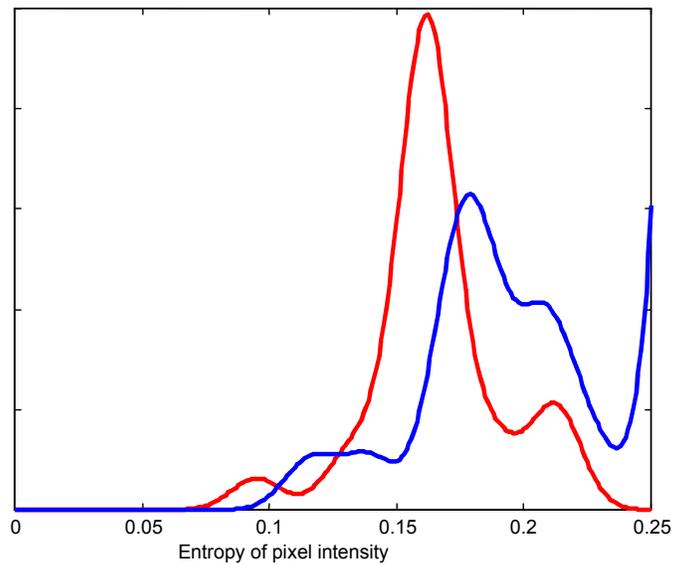
A smooth and uniform surface implies a narrow distribution of pixel intensities over the visible part of an object, whereas a non-homogenous surface is characterized by a wider distribution. These properties of the distributions can be quantified by the standard deviation and the entropy [6.8][6.4] of the pixel intensities.



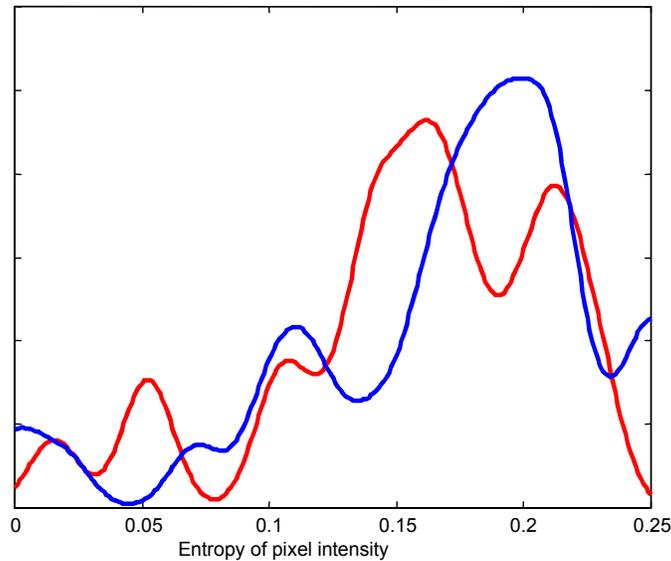
**Figure 6.12** Distribution of standard deviation of pixel intensities, JSC/OSU laboratory data set. This set included several granite samples, characterized by a speckled texture.



**Figure 6.13** Distribution of standard deviation of pixel intensities, Patriot Hills, Antarctica, field data set. Rocks in this set are mostly marble, limestone or quartzite, none of which has much texture. Many of the meteorites were broken, with a lighter colored interior visible, resulting in the bimodal distribution in the meteorite texture.



**Figure 6.14** Distribution of pixel intensity entropy for meteorites and terrestrial rocks, JSC/OSU laboratory data set.

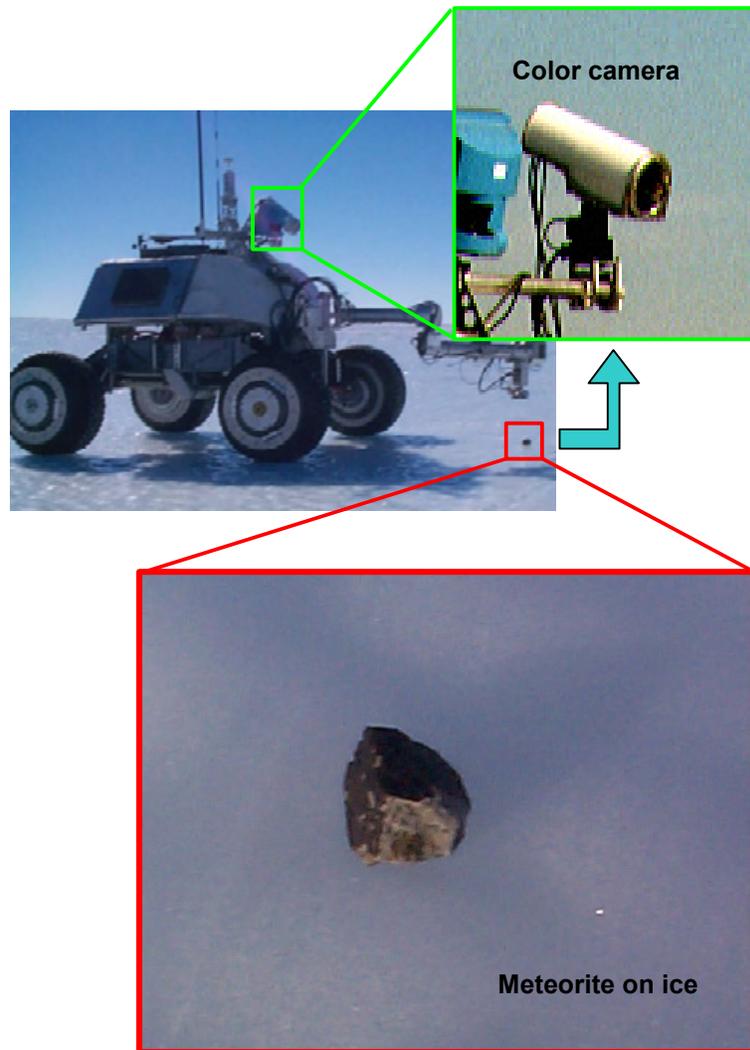


**Figure 6.15** *Distribution of entropies of pixel intensity for meteorites and terrestrial rocks, Patriot Hills, Antarctica, field data set.*

Both the entropy and standard deviation of the intensity distributions for rock samples show a slight, but not decisive, difference between terrestrial rocks and meteorites (Figure 6.12, Figure 6.13, Figure 6.14, and Figure 6.15). The laboratory data set shows a bigger difference, largely because it contained many granite samples amongst the terrestrial rocks and because the pictures were of higher resolution. This dependence on resolution could be a problem except that Nomad's camera is equipped with a zoom lens and so can limit excessive disparities in resolution between images.

### 6.3. Nomad high-resolution camera

Robotically acquiring images of rocks suitable for classification requires some care. Images must have sufficient resolution, contrast and color balance. Varying lighting conditions must be accounted for, especially if trying to measure the intrinsic color and albedo of rocks. Determining the size of objects requires that the scaling of images be known, requiring therefore that the camera lens be calibrated and the distance to rock samples be known. All this must be accomplished autonomously and in the harsh polar conditions.



**Figure 6.16** Color camera on Nomad for acquiring close-up images of rocks and meteorites on the Antarctic ice. A pan and tilt mechanism points the camera at rock samples and a zoom lens enables the acquisition of high-resolution close-up images. This shows an ideal image for classification. Most field images have less resolution and contrast.

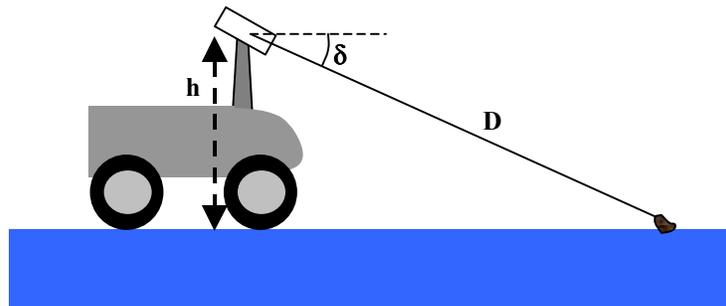
It was estimated that the visual classification of rocks requires images with a spatial resolution of at better than 0.5 mm/pixel. Nomad's camera (Figure 6.16) is equipped with a zoom lens to provide this resolution at 5 m, and to zoom out so the camera can be used to search the ice for rocks. Originally, a separate 360° panoramic camera was planned for the latter task but its resolution proved insufficient for detecting meteorite-sized objects several centimeters in diameter.

The camera system (Figure 6.16) consists of a 640x480 3-chip CCD camera with a zoom lens on a pan / tilt unit. The field of view can be varied from 28° to 1.9°. The pan and tilt unit is attached to Nomad's sensor mast 1.8 m above the ground, enabling a look-ahead distance of 5m. The pan and tilt unit allows close-up images of rocks within a large workspace in front of the robot.

The 3-chip color CCD camera was chosen to ensure correct color pixel registration, although this has proved to be unnecessary for this classifier. While getting good images, the 3-chip camera's lens requirements restricted the range of zoom lens's that could be used.

To acquire a close-up image once a sample has been located in a wide-angle search image, the lens is progressively zoomed until the sample fills about 75% of the image. At each stage, small movements of the pan and tilt are used to center the sample in the image.

The spatial resolution of the zoomed in image has to be computed in so that the rock size can be deduced. This is determined by multiplying the camera angular resolution by the distance from the camera CCD to the rock sample, deduced by assuming the world is flat and projecting a ray from the camera to the sample (Figure 6.17).



**Figure 6.17** Flat world assumption used to compute the distance **D** from the camera CCD sensor to a rock visible in an image. The height of the CCD above the ground, and the horizontal declination of the ray from the camera CCD to the rock are computed from the pant and tilt values of the camera, the position of the rock in the image and the lens focal length.

With the flat world assumption the distance from the camera to the rock sample is

$$(Eqn 6-8) \quad D = \frac{h}{\sin \delta}$$

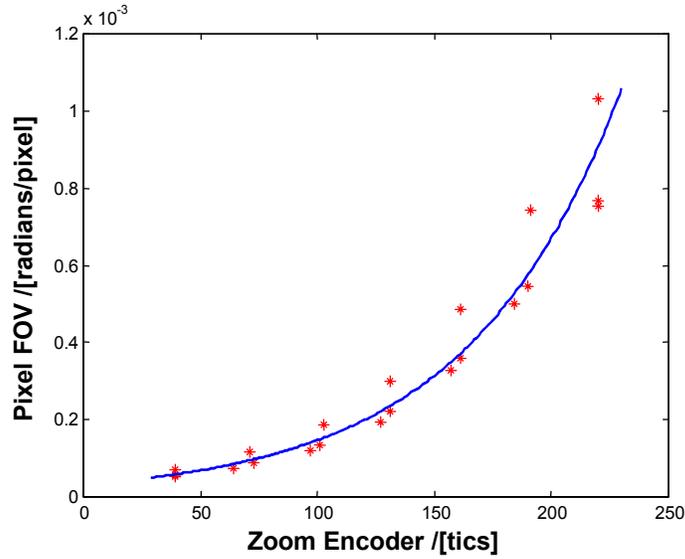
The angular resolution is a function of the lens's focal length (zoom setting). However, the zoom encoder on the lens was uncalibrated, and the published specifications gave no indication of the functional relationship between the zoom encoder values  $Z$  (0...255) and the actual focal length  $f$  (mm). Assuming a logarithmic variable resistor as the lens zoom transducer, it follows that

$$(Eqn 6-9) \quad f = \alpha e^{-\beta Z} \text{ [mm]}$$

and therefore the angular field of view of each pixel  $Fov$  is given by

$$(Eqn 6-10) \quad Fov = A e^{\beta Z} \text{ [radians/pixel]}$$

The constants  $A$  and  $\beta$  were determined by measuring  $Fov$  on a set of calibration images taken with different zoom settings and the computing the best fit (Figure 6.18).



**Figure 6.18** Non-linear relationship between the cameras zoom encoder and the angular field of view of each pixel. This had to be experimentally determined from a sequence of calibration images and an appropriate (ad hoc) model fitted (Eqn 6-10). With distance, the pixel field of view is needed to determine the size of objects in an image.

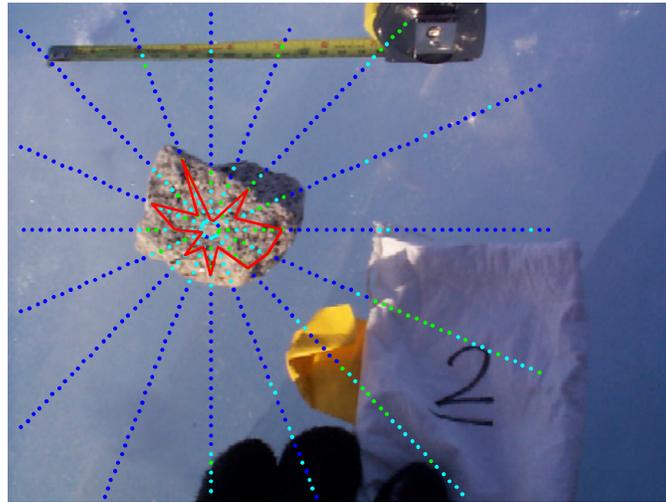
Auto-focus software was implemented using [6.9], based on the algorithm of as the lens does not provide this capability. The focus quality is calculated as the sum of pixel intensity differences over a centered image window bounded by columns  $C_0$  to  $C_n$  and rows  $R_0$  to  $R_m$ , as shown in the equation

$$(Eqn 6-11) \quad FocusQuality = \sum_{c=C_0...C_n} \sum_{r=R_0...R_m} |Intensity(r,c) - Intensity(r,c+1)|$$

The optimum focus setting is that which maximizes *FocusQuality*.

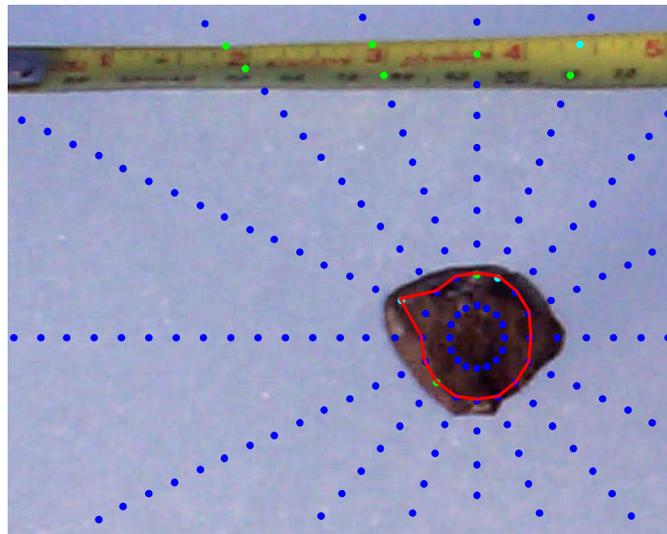
The camera and lens also support automatic iris control. Image light levels are dominated by the snow and ice in the background of all images so the iris setting depends mainly on the illumination of the background and the fraction it takes up of the image, not the color and albedo of the rock being examined. The constancy of the background makes it a crude reflectance standard for setting the aperture. Rock pixel colors are thus with respect to the “ice standard” and in theory independent of the illumination.

## 6.4. Conclusions



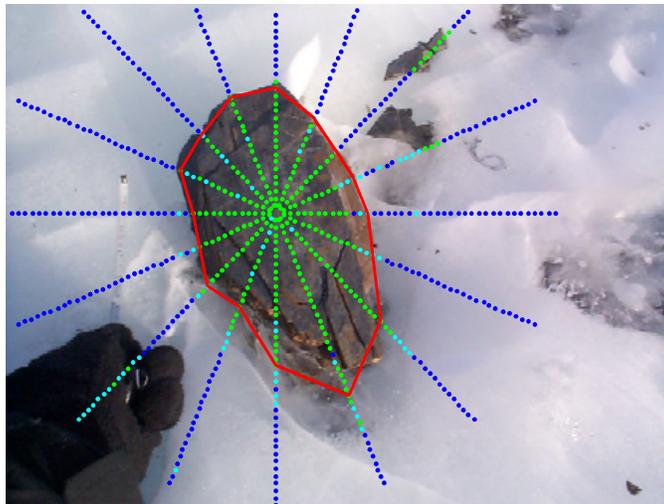
<i>Intensity</i>	<i>Area</i>	<i>Standard dev.</i>	<i>Entropy</i>	<i>Compactness</i>	<i>BE</i>
0.52	1.38	0.13	0.26	0.19	7.8

**Figure 6.19** Segmentation and extracted features of a light colored granite rock on ice. This demonstrates the inferior segmentation for some light colored rocks on a dark ice background. Nonetheless, the computed boundary is conservative and does not enclose any ice pixels, therefore the computed color and texture features are unaffected. Furthermore, better segmentation can be expected for dark colored meteorites and dark colored rocks, which are most likely to be confused with them.



<i>Intensity</i>	<i>Area</i>	<i>Standard dev.</i>	<i>Entropy</i>	<i>Compactness</i>	<i>BE</i>
0.17	0.33	0.061	0.18	0.87	14.4

**Figure 6.20** Segmentation and extracted features for a meteorite. Note the superior segmentation to Figure 6.19, and the lower intensity, lower area, lower standard deviation, greater entropy and greater compactness than seen for many terrestrial rocks.



<i>Intensity</i>	<i>Area</i>	<i>Standard dev.</i>	<i>Entropy</i>	<i>Compactness</i>	<i>BE</i>
0.31	24.5	0.077	0.22	0.78	0.30

**Figure 6.21** Segmentation and feature vector for a large sandstone rock in Antarctica. Note how the rock is distinguished from several different backgrounds, including snow, sunlit and shadowed ice. Furthermore, the segmentation is resistant to the extraneous objects in the image. The area and intensity distinguish this rock from the meteorite in Figure 6.20, although the standard deviation, entropy and compactness do not.

The previous plots demonstrate that simple features extracted from images of rocks and meteorites are feasible for sometimes, but not always, distinguishing the two. This is acceptable given that the purpose of the camera on Nomad is to exclude obvious terrestrial rocks from consideration and not the primary means of distinguishing rocks from meteorites.

It is probable that better metrics exist for visually distinguishing rocks from meteorites, but visual classification is not the topic of this thesis. Apart from being a difficult topic, it does not directly address the robotic issues or rock classification, such as sensor fusion, or modeling the environment. In this thesis, the main interest in visual classification is to demonstrate multi-sensor fusion and to improve spectral classification.

## 6.5. References

- [6.1] Spiegelhalter, D.J., Philip, D.A., Lauritzen, S.L. and Cowell, R.G., “Bayesian analysis in expert systems” in *Statistical Science*, 8(3), p219-283., 1993.
- [6.2] Pedersen, L., Wagner, M., Apostolopoulos, D., Whittaker, W., “Autonomous Robotic Meteorite Identification in Antarctica”, in *IEEE International Conference on Robotics and Automation*, Seoul, Korea, 2001.
- [6.3] Wagner, M., Apostolopoulos, D., Shillcutt, K., Shamah, B., Whittaker, W., “The Science Autonomy System of the Nomad Robot”, in *IEEE International Conference on Robotics and Automation*, Seoul, Korea, 2001.

- [6.4] Jain, A.K., *Fundamentals of Digital Image Processing*, Prentice-Hall, 1989.
- [6.5] Fukunaga, K., *Introduction to Statistical Pattern Recognition*, Academic Press, 2nd edition, 1990.
- [6.6] Bishop, C.M., *Neural Networks for Pattern Recognition*, Oxford University Press, 1995.
- [6.7] Gelman, A., Carlin, J.B., Stern, H.S. and Rubin, D.B., *Bayesian Data Analysis*, Chapman & Hall, 1995.
- [6.8] Papoulis, A., *Probability, Random Variables, and Stochastic Processes*, 3rd edition, McGraw-Hill, Inc., 1991.
- [6.9] Krotkov, E., “ Focusing”, *International Journal of Computer Vision*, vol. 1, no. 3, October 1987, pp. 223-237. (Reprinted in *Physics-Based Vision*, Jones and Bartlett, 1992.)

## Electromagnetic sensors

Electromagnetic sensors, generally taken to mean metal detectors and magnetometers, are sensitive to conducting and magnetic media, especially metallic iron. They have been successfully used for landmine clearance, meteorite search, and treasure hunting to name a few of their applications.

This chapter describes how they may be deployed from a robot and a series of meteorite search trials done with Nomad in the Atacama Desert in 1997. Although effective for meteorite search, logistical reasons prevented their subsequent use with Nomad in Antarctica.



**Figure 7.1** Robotic deployment of electromagnetic sensors (metal detector and magnetometers) in the Atacama Desert. Sensors are mounted on a non-metallic (polypropylene) platform, a distance away from the robot to minimize electromagnetic interference from the robot vehicle.

Approximately 90% of all meteorites, including the so-called stony meteorites, contain significant amounts of metallic iron and nickel [3.20][3.21]. In sufficient quantities, both of these are detectable with metal detectors from a few centimeters, depending on the amount. Furthermore, because both Iron and Nickel are ferromagnetic materials, and so affect the ambient magnetic field. In contrast, most, but not all, terrestrial rocks contain conductive or magnetic materials, and so do not register with metal detectors and magnetometers.

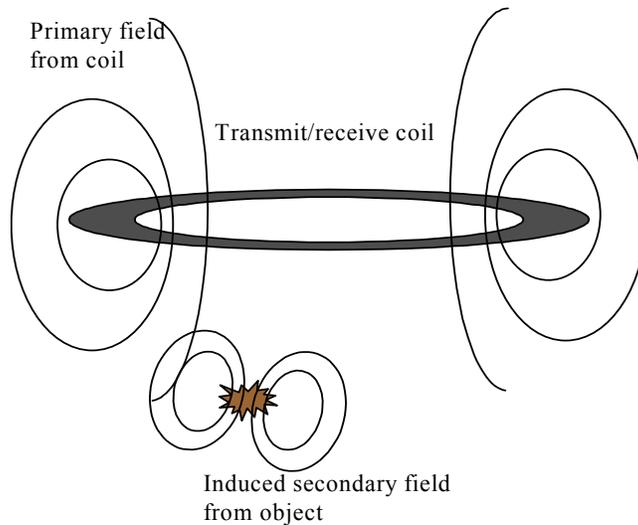
## 7.1. Metal detectors

Metal detectors (Figure 7.2), more strictly known as eddy current sensors because of the way they work, are obvious devices for detecting those meteorites containing metallic iron or nickel. Furthermore, they are effective search instruments and can locate buried objects.

An eddy current sensor works by pulsing current through a coil, creating a magnetic field that induces a secondary field from any conducting objects nearby. A receiving coil senses this secondary field (Figure 7.3) using various proprietary electronics packages unique to each manufacturer [7.8][7.9]. The two main classes of eddy current sensors are pulse induction sensors and continuous wave sensors, the difference being whether the transmit coil is pulsed or energized by a continuous alternating current. The primary advantage of the former is that CW devices cause some landmines to explode. Manufacturers also claim better rejection of certain conducting soils and more immunity to EM interference.



**Figure 7.2** Pulse 8x® metal detector, a pulse induction type of eddy current sensor used to find meteorites in the Atacama Desert. [Photo: J.W.FisherMfg product brochure]



**Figure 7.3** Schematic of the operating principles of a metal detector or “eddy current sensor”. An electric current in the device transmitter coil creates a large magnetic field, which interacts with any conducting object nearby inducing eddy currents to flow in them. These in turn generate a secondary magnetic field that induces a signal in the receiver coil.

Without going into the details of the myriad metal detector models available the following points apply

- Return signal strength decreases dramatically with distance,  $\sim 1/d^6$ . Therefore detection range is inherently limited. Maximum detection range is comparable to the search coil diameter, but large coils are insensitive to small objects.
- The search coil of a metal detector is an excellent radio antenna, so certain models are very sensitive to RF interference. Furthermore, they are also sources of RF and magnetic interference.
- The detection footprint of a metal detector is slightly larger than the search coil size. Hence, to be used for meteorite search, rather than simply identification, a large coil is desirable. Elongated coils are sometimes used to cover a large area and yet still be able to detect small objects.
- Metal detectors have a minimum dwell time if they are to detect an object, typically about 0.5 seconds. This limits the speed with which they can be scanned over an area.
- Sensor response to different metallic objects is difficult to predict [Yoga Das references], depending on composition, size, shape and orientation. Experimentation is the best way to verify sensor appropriateness.

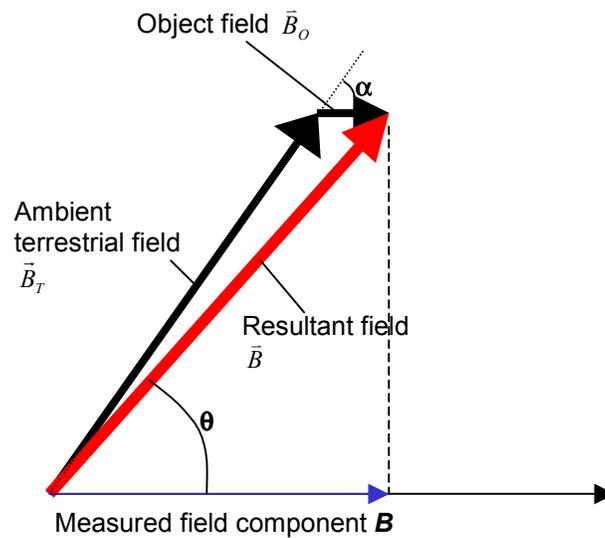
A large number of metal detector models were tested on some meteorite samples. Top of the line military mine detectors have comparable performance to the better civilian models for the treasure hunting industry, but are more expensive and have a camouflage finish. All detectors incorporate circuitry that sounds an audio alarm when the signal strength exceeds some variable threshold. Some designs attempt to suppress soil background signals, usually by looking for signal differences at different places. These require that the detector be continuously scanned. If the detector permits direct access to the receive coil signal for computer processing then all these other features are quite unnecessary.

## 7.2. Magnetic field measurement

There exist two types of magnetometers: vector sensors and total field sensors. Vector sensors measure a single component of the magnetic field vector at the sensor location. They are the most common type of sensor, and include fluxgate sensors, optical fiber interferometers, Hall effect sensors and super-conducting quantum interference devices (SQUID's). Total field sensors are sensitive to the magnitude, but not direction, of the magnetic field. However, most varieties of these sensors, such as magnetic resonance magnetometers need to be approximately aligned with the magnetic field in order to function.

The powerful and pervasive geomagnetic field that is significantly bigger than the field due to a meteorite complicates measuring the magnetic field due to an object. If  $\vec{B}_E$  is the (primary) ambient geomagnetic field and  $\vec{B}_O$  is the (secondary) field associated with the object, then the resultant magnetic field vector  $\vec{B}$  is given by

$$(Eqn 7-1) \quad \vec{B} = \vec{B}_O + \vec{B}_E$$



**Figure 7.4** Relationships between the terrestrial geomagnetic field vector, a meteorite magnetic field vector, the resultant magnetic field vector and quantities measurable with magnetometers, which are sensitive to either a single component or the magnitude of the resultant magnetic field vector.

Since the primary field is larger than the secondary, then, to first order, the magnitude of the resultant field is equal to the magnitude of the primary field plus the magnitude of the projection of the secondary field onto the primary:

$$(Eqn 7-2) \quad |\vec{B}| \approx |\vec{B}_E| + |\vec{B}_O| \cos \alpha$$

Therefore, a total field sensor, measuring only to this magnitude, is insensitive to any fields perpendicular to the primary.

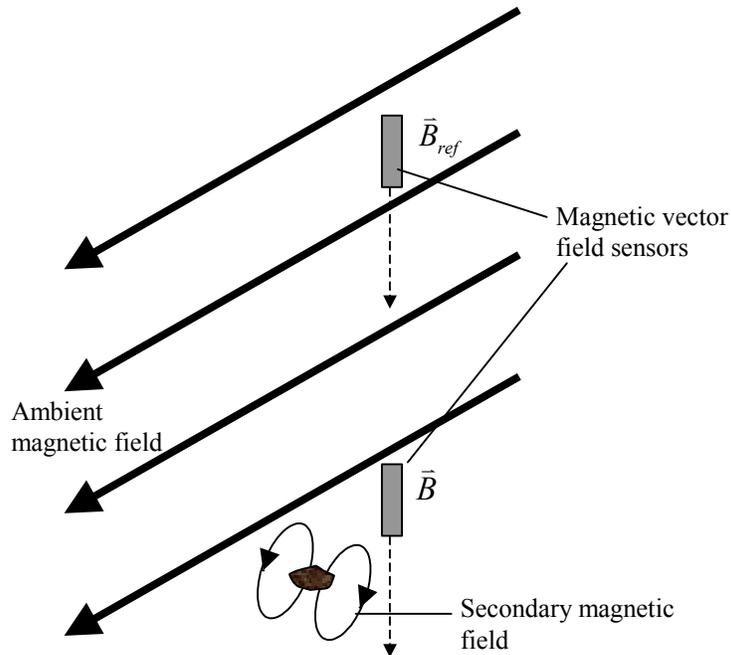
The output of a vector magnetometer, unlike a total field magnetometer, varies with sensor orientation and in direct proportion to the total field strength.

$$(Eqn 7-3) \quad B = |\vec{B}| \cos \theta$$

$$\approx \left( |\vec{B}_E| + |\vec{B}_O| \cos \alpha \right) \cos \theta$$

A sensor operating in a large primary field (such as the Earth's) is therefore particularly sensitive to changes in orientation. In fact, variations in the measured field component due to changes in sensor attitude can easily exceed those due to the small secondary field of a meteorite. If the sensor is mounted on a mobile robotic vehicle it will be subject to motion and keeping it at a fixed orientation with respect to the ambient geomagnetic field is futile, since the vehicle itself will be in motion. This is where total field systems hold the advantage, as they are insensitive to changes in sensor orientation. However, in practice they still need to be approximately aligned with the ambient field in order to function.

The secondary field of small objects like meteorites is short range, decaying rapidly with distance from the object. The ambient geomagnetic field is relatively constant in space, so sensors measuring the field some distance away from any secondary magnetic field sources can be used to infer the primary field near a rock or meteorite sample. Subtracting this from the measured field near a sample allows the primary field to be deduced. This is the principle behind the magnetic gradiometer, typically consisting of two precisely aligned vector sensors: a sensor near the sample and a reference some distance away.



**Figure 7.5** *Magnetic gradiometer configuration. A reference measurement of the magnetic field is made far from the source of the secondary field to infer the ambient (primary) field  $\vec{B}_E \approx \vec{B}_{ref}$  at a sensor placed by the source, which measures a component of the resultant field  $\vec{B} = \vec{B}_O + \vec{B}_E$ . The two vector magnetometers are precisely co aligned so that they both measure the same vector component.*

With this configuration, the primary field near an object is approximately given by a reference:

$$(Eqn 7-4) \quad \vec{B}_E \approx \vec{B}_{ref}$$

Because the sensors are co-aligned, they both measure the same field components. A component of the objects secondary field is therefore given by

$$(Eqn 7-5) \quad B'_O = \vec{B}_O \cos \theta$$

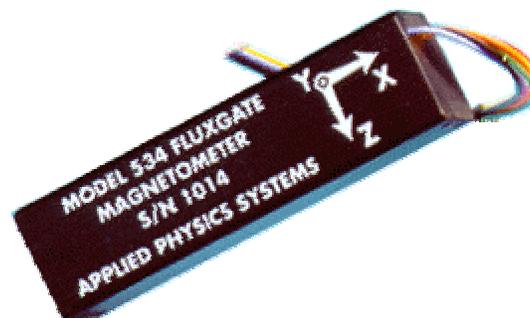
$$\approx B' - B'_{ref}$$

Where  $\theta$  is the angle between the angle between the secondary field and the magnetometer preferred axis. The effect of the ambient field is therefore accounted for, regardless of the sensor orientations or any common mode variations in the ambient field. There is still variation in signal strength with orientation, but *only if* a secondary field is present. No secondary field perpendicular to the magnetometer axis will be detected. A graver problem from the robotic standpoint is keeping the two sensors precisely co-aligned, especially if they are to be in the harsh environment. Nonetheless, this configuration has proven suitable for application such as landmine and unexploded ordinance detection (Figure 7.6)



**Figure 7.6** Ferex® magnetic gradiometer system for landmine and unexploded ordinance detection. Note the vertical probe, containing two co-aligned magnetic vector sensors. This unit was tested on various meteorite specimens in 1997, verifying that they had detectable magnetic fields. [Photo: Foerster Systems product brochure].

Complete sensing of the full magnetic field vector can be achieved with three orthogonal vector sensors (Figure 7.7). The field magnitude can be calculated from this, and will be orientation invariant provided the sensors have zero offset and identical scaling. Arranged in a gradiometer configuration 3-axis sensors can be used to determine the full secondary field around an object. This is clearly more sensitive than the single axis gradiometer, but still requires careful alignment or calibration of the sensors.



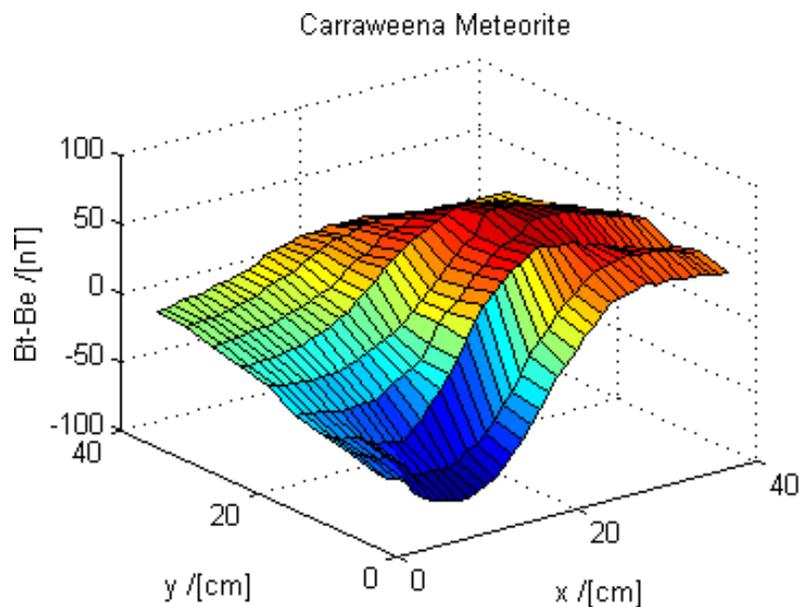
**Figure 7.7** 3-Axis fluxgate magnetometer for measuring the total magnetic field vector, tested during robotic meteorite detection trials in the Atacama Desert, Chile, in 1997. [Photo: Applied Physics Systems].

### 7.2.1. Magnetic field of meteorites

For ferrous objects, such as Iron and Nickel containing meteorites, the secondary field itself consists of two components: an induced field produced by the interaction of the object with the primary (terrestrial) field, and that due to the permanent, or remnant, magnetization of the object:

$$(Eqn 7-6) \quad \vec{B}_O = \Lambda \vec{B}_E + \vec{B}_p$$

Therefore, any ferrous object will produce a weak secondary field in the presence of the primary field. A permanently magnetized object produces a stronger, more detectable field. To investigate the magnetization of meteorites the magnetic field of several meteorite samples was measured with a Caesium vapor total field magnetometer at Canada's Defense Research Establishment Suffield (DRES) "foam dome" facility in Medicine Hat, Canada. Measurements were taken at regular intervals on a plane above the samples (Figure 7.8).



**Figure 7.8** Total magnetic field strength measured on a plane 3 cm above a meteorite sample. This is a common stony Chondrite meteorite, approximately 20 x 15 x 15 cm in size and containing metallic iron filings.

Some of the stony meteorites had fields up to 9 times stronger than that predicted for the same mass of unmagnetized iron by a computer model developed at DRES [7.4]. Furthermore, when turned upside down, the field of many meteorites is a direct inversion of the former field. These facts suggest a strong remnant magnetization in the three stony meteorites (Chondrites) tested, and a moderate remnant field for a small (1 cm diameter) iron meteorite.

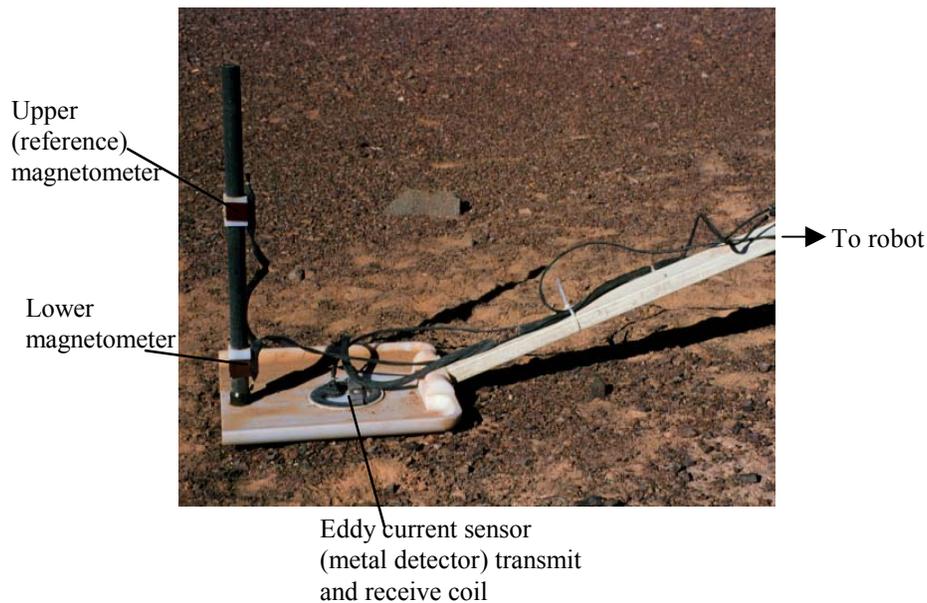
Laboratory detection distances varied from 13 cm (small iron meteorite) to approximately 25 cm (Carraweena stony meteorite). Deviations of total field strength from the ambient geomagnetic field ranged from 50nT (iron meteorite at 11cm distance) to 90nT (Carraweena at 19cm). However, it should be noted that the meteorites available for testing were large specimens found in Australia, and not necessarily representative of all those typically found in Antarctica, with the exception of the small iron fragment.

Antarctic rocks were also tested. They exhibited negligible magnetization (about 10nT peak for one sample, 20 times less than the predicted field of the same amount of iron).

By comparison, the Earth's field varies from 30,000nT at the magnetic equator to 100,000nT at the magnetic poles. It is relatively constant over distances of several hundred meters (the magnetic field gradient is  $\sim 10$  nT/km) and, barring magnetic storms, does not change significantly over short time scales [3.21].

### 7.3. Atacama robotic sensor deployment

In order to evaluate magnetometers and metal detectors as sensors for detecting meteorites from a robot in actual field conditions a series of trials were conducted with Nomad in the Atacama Desert, Chile in 1997 (Figure 7.1). A pulse induction type metal detector (J.W.Fisher Pulse 8x) with an 8" coil and two 3-axis magnetometers (Figure 7.7) arranged in a gradiometer configuration were mounted on a non-metallic platform dragged behind the robot (Figure 7.9). The platform was arranged 8' away from the robot to minimize the effect of EM interference generated by it.



**Figure 7.9** *Electromagnetic sensor platform, dragged behind the Nomad robot vehicle during meteorite detection trials in the Atacama Desert. The platform is completely non-metallic, made mostly of polypropylene and fiberglass. All cabling is shielded copper, which did not affect the sensors.*

The sensors were connected by shielded copper cables to the differential inputs of an A/D, using a standard interference reducing cable arrangement. The magnetometers were sampled directly, while the metal detector had a custom built output. Because of the magnetometer 400 Hz pass band the sensors outputs were sampled 800Hz, and with 16 bit precision.

#### 7.3.1. Atacama signal processing

The eddy current sensor generates significant 250Hz magnetic fields that are picked up by the magnetometers. In addition, in spite of the distance from Nomad, there was still significant vehicular noise

interference. Fortunately, since the vehicle speed was approximately 30cm/s and the sensor footprint is approximately 20cm, signals of interest are less than 10Hz, while most of the noise is at higher frequencies. Low pass filtering with a 10Hz cut-off frequency and down-sampling to 20Hz improved the data and reduced its volume.

The top magnetometer (low pass filtered) readings are used to infer the primary field (geomagnetic field and common mode low frequency vehicular interference field emissions) measured by the bottom magnetometer. Thence the secondary magnetic field due to a ferrous object can be deduced, and orientation independent quantities (such as magnitude) computed from it. Furthermore, low frequency common mode interference (that cannot be removed with low pass filters) is rejected.

Assuming a uniform magnetic field (such as the Earth's), the fields *measured* (as opposed to the actual fields; the magnetometers response to a field may change with time) by the top and bottom 3-axis vector sensors are related by a rotation, scaling and displacement:

$$(Eqn 7-7) \quad \vec{B}_{bottom} = \tilde{S} \cdot \tilde{R} \cdot \vec{B}_{top} + \vec{D}$$

$$\tilde{S} = \begin{bmatrix} s_1 & 0 & 0 \\ 0 & s_2 & 0 \\ 0 & 0 & s_3 \end{bmatrix} \text{ is a scaling matrix.}$$

$$\tilde{R} \text{ is a rotation matrix.}$$

The rotation reflects the arbitrary orientation of the sensors with respect to each other. The scaling and offset arise from differences in the sensor scale factors and zero offsets.

Thus, the secondary field secondary field *as measured* by the bottom magnetometer *less* any zero offsets due to the instruments is given by

$$(Eqn 7-8) \quad \vec{B}_O = \vec{B}_{bottom} - (\tilde{S} \cdot \tilde{R} \cdot \vec{B}_{top} + \vec{D})$$

From this the secondary field magnitude is computed:

$$(Eqn 7-9) \quad B = \|\vec{B}_O\|$$

Another quantity to gauge the presence of a small secondary magnetic field is the angle between the bottom sensor field and the ambient field predicted to be there from the top sensor measurements:

$$(Eqn 7-10) \quad A = \angle(\vec{B}_{bottom}, \tilde{S} \cdot \tilde{R} \cdot \vec{B}_{top} + \vec{D})$$

These quantities vanish if there is no secondary field and are invariant under variations in the orientation of the sensor platform or common mode interference generated by the robot. However, for this to be so it is essential that the zero-offsets be accounted for.

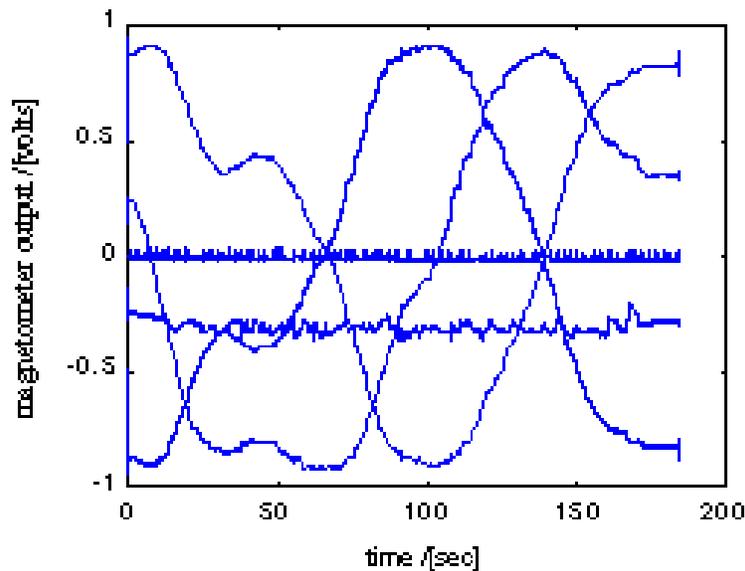
Given calibration data obtained the unknown matrices and offset vector in (Eqn 7-7) are solved for using least squared error estimation and the algorithm in [7.5] to solve for the rotation.

### 7.3.2. Atacama results

Noise and calibration proved to significant issues. The electric motors in the wheels and vehicular iron contributed significantly to the ambient magnetic field. Computers, switching mode power supplies and radio communications equipment generated substantial RF interference. Experience proved the metal detector to be particularly sensitive to the latter. For example, it would saturate whenever personnel in the

robot vicinity used VHF radios. However, interference was hard to predict. The gigahertz communications system on the robot did not appear to affect the metal detector, but its signal would vary greatly in the evenings, even when held away from the ground or any local EM sources. This latter effect was particularly perplexing. It could be speculated that it was caused by evening changes in the ionosphere that affected the signal propagation from distant radio stations, causing sufficient local interference.

The harsh conditions wreaked havoc with the magnetometers, which were exposed to vibrations, shock and temperature extremes. While they remained rigidly aligned with respect to each other, they developed unexplained zero offsets in some channels, and possible scale changes with temperatures. One unit failed completely and had to be replaced. The zero offsets would have, if uncorrected for, greatly degraded the results.

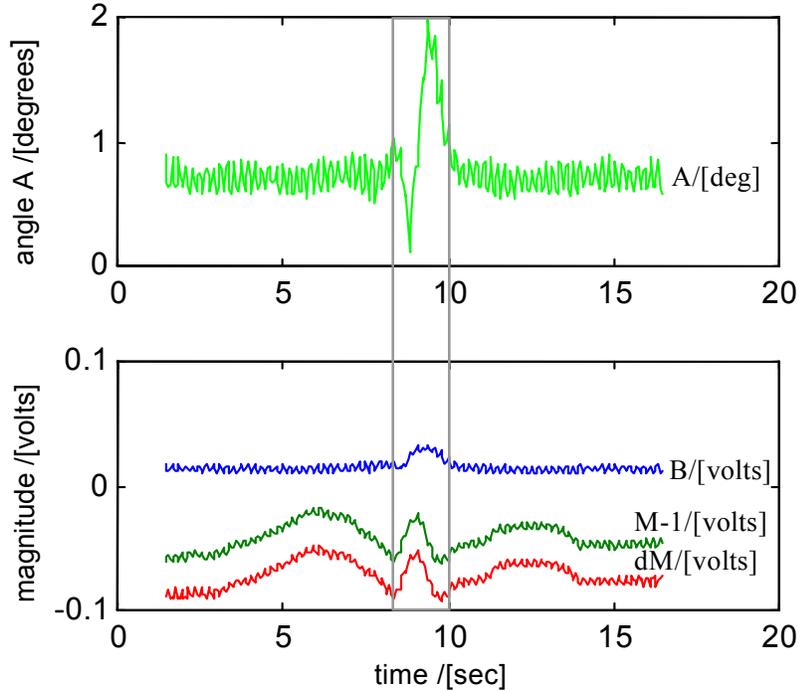


**Figure 7.10** Output from upper and lower magnetometers as the sensor platform is dragged in a circle in an area of the Atacama Desert with no secondary field sources. Note the failure of one channel, giving a zero output.

Magnetometer calibration data was obtained in the Atacama Desert with the rover dragging the sensor platform along a curved path relatively free of ferrous objects (this was not possible in the Pittsburgh test sites because of the amount of iron lying about). Solving for the transformation matrices and offset (Eqn 7-7) between the magnetometers yielded

$$(Eqn\ 7-11) \quad \tilde{R} = \begin{bmatrix} -1.00 & -0.003 & -.001 \\ 0.00 & -1.00 & -0.03 \\ 0.00 & -0.03 & 1.00 \end{bmatrix} \quad \tilde{S} = \begin{bmatrix} 0.98 & 0 & 0 \\ 0 & 1.00 & 0 \\ 0 & 0 & 0.99 \end{bmatrix} \quad \bar{D} = \begin{bmatrix} -0.23 \\ 0.01 \\ -0.01 \end{bmatrix}$$

Note the large zero offset between the X channel of the magnetometers. Figure 7.11 indicates the efficacy of this transformation in computing sensor platform orientation independent quantities A and B that indicate the strength of the secondary field due to a meteorite.



**Figure 7.11** Comparison of the secondary field measures  $B$  (Eqn 7-9) and  $A$  (Eqn 7-10) using the transformation matrices (Eqn 7-11) with the traditional measures  $M$ , the magnitude of the bottom sensor measured field, and  $dM$ , the difference in magnitudes between the fields measured by both magnetometers. This data was generated from readings taken as the magnetometer sled was dragged over a large meteorite, causing it to pitch considerably. Because of a significant offset in the output of the lower sensor, both  $M$  and  $dM$  vary significantly with the pitching of the sensor platform. Conversely,  $A$  and  $B$  do not vary at all except when the sensors pass over the meteorite (indicated by the central rectangle), since they account for the offset. Furthermore, signal level of  $B$  does not depend on the orientation of the secondary field with respect to the primary field, as does  $M$  or  $dM$ , and has better noise rejection.

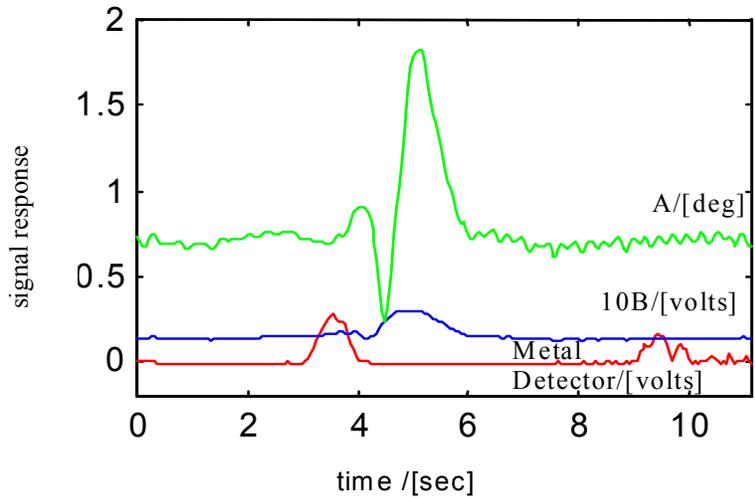
Ideally  $A$  and  $B$  would be both zero in the absence of a secondary field. However, not all vehicular interference is common to both sensors. Furthermore, the three vector sensors in each magnetometer unit are not strictly orthogonal, but may deviate slightly ( $< 0.2^\circ$ ). Therefore, there should also be a slight shear component to the transformation equation. Nevertheless, it is small, and estimating it along with the other parameters in the presence of noise turned out to be unstable.

Further data was collected with real meteorites (olive seed sized iron samples to larger stony samples) both in the Atacama and in Pittsburgh, with the vehicle travelling at approximately 30cm/s. Differing geomagnetic fields and mineralogy affected the results. Additionally, the complete failure of one magnetometer and the partial failure of another in the Atacama necessitated separate calibration for each location, and indicates the probable necessity of multiple recalibration of the system if deployed for extended periods of time.

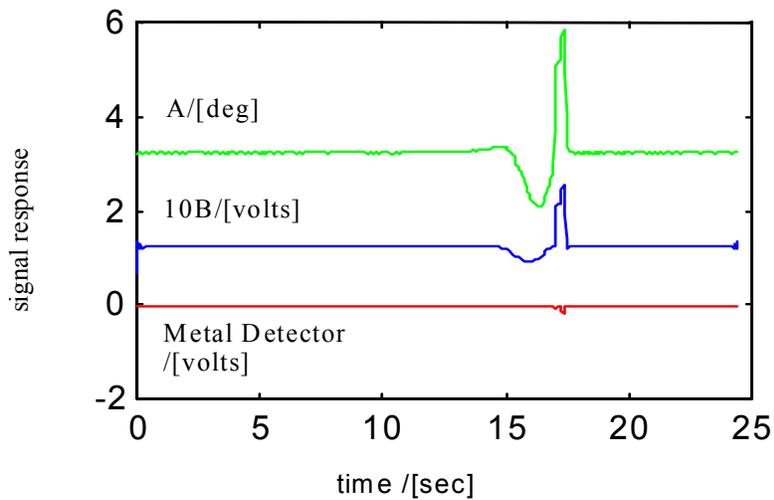
Sensor responses (processed) to a large stony meteorite (Carraweena) are shown in Figure 7.12, Figure 7.13, and Figure 7.14. The different geomagnetic field strengths (60,000nT in Pittsburgh, 30,000nT in the Atacama) and declinations might explain the different magnetomer responses. In addition, there were many

iron pieces in the Pittsburgh slag heap that affected the sensors. In comparison to the field data, the laboratory magnetometer data (obtained by manually moving the meteorite under the sensors) is exceptionally noise free. However, there is a significant DC offset, probably due to the large iron content in the laboratory building.

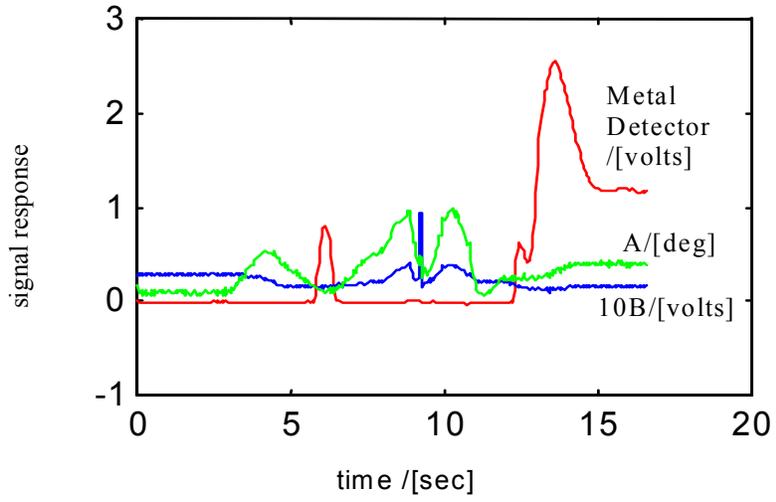
Figure 7.15 and Figure 7.16 show the sensor response to a small iron meteorite (the sample with the smallest magnetic signature, but largest metal detector signature) and the Kulnine stony meteorite (no metal detector response).



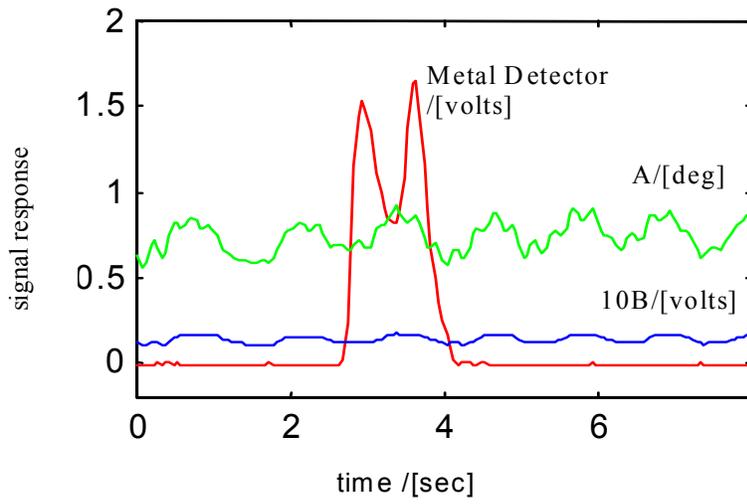
**Figure 7.12** Carraweena meteorite, Atacama desert



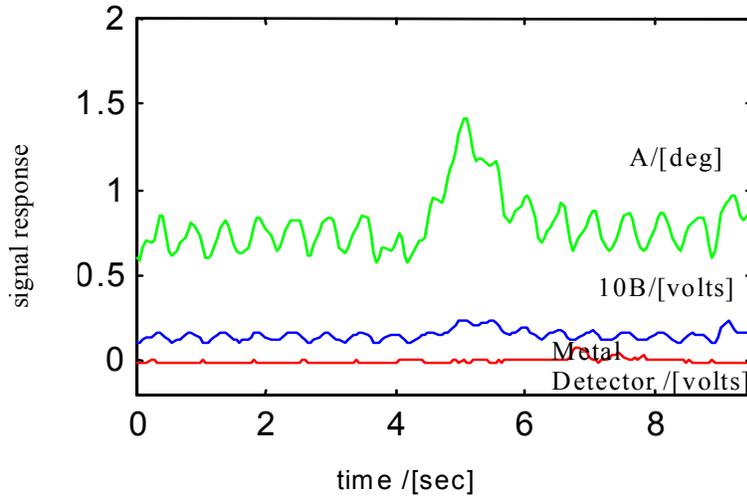
**Figure 7.13** Carraweena meteorite, Pittsburgh laboratory



**Figure 7.14** Carraweena meteorite, Pittsburgh slag heap



**Figure 7.15** Small iron meteorite, Atacama Desert



**Figure 7.16** *Kulnine meteorite, Atacama desert*

### 7.3.3. Atacama lessons

The results show that a simple threshold-crossing detector is less than adequate for autonomous recognition of meteorite electromagnetic signatures. An adaptive thresholding algorithm produced tolerable results but still gave a high number of false positives. Because of the large variation in the signatures, even for the same sample, it is questionable whether a single pattern classifier could be built to successfully resolve all cases.

The sensor configuration developed here is able to accommodate both sensors on a single platform and resolves the vehicular noise and calibration issues mentioned earlier. The magnetometers can be calibrated in the field and the output processed to obtain quantities that are independent of orientation, yet sensitive to localized fields produced by objects near the sensors. The value of this system was proven in the field when the lower magnetometer developed a large zero offset causing naïve computations of the lower field magnitude to depend strongly on sensor orientation.

Signal quality could possibly be further improved by removing periodic interference (c.f. Figure 7.15 and Figure 7.16). However, initial trials were not encouraging as the interference has multiple harmonics and removing them all adversely affected the signal.

From the robotic perspective, the ability to recalibrate sensors in the field is essential. Long duration missions are likely to experience sensor changes. Normally, special facilities would be needed. Without calibration, sensor outputs will vary unacceptably as sensors are moved over rough terrain, with a high number of false positive results.

The combined pulse induction metal detector and magnetometer system is effective for detecting meteorites. The sensors complement each other well, the former responding most strongly to iron meteorites, and the latter to stony meteorites.

The metal detector proved to be a particularly valuable yet troublesome sensor. When deployed manually it was sensitive to all the sample meteorites, and was used to discover two Pallasite (stony iron meteorites) specimens in a known impact site at Imilac in Chile. However, results from a robotic deployment are less

encouraging. The unit was sensitive to vehicular emissions, and would occasionally behave unpredictably for no apparent reason.

If the robot was traveling at excessive speed, the metal detector was unable to respond to metal objects within its footprint. Search speed is limited by the device response time, to less than 30 cm/s for the model used by us.

Integrating the sensors with the robot navigation would greatly improve detection rates. The odometer logs are of particular importance. By associating a position with sensor readings, data taken at different times from different sensors at the same point can be more effectively compared.

## 7.4. References

- [7.1] Heide, Fritz, "Meteorites", University of Chicago Press, 1964
- [7.2] Wasson, John T., "Meteorites: classification and properties", Springer-Verlag, 1974.
- [7.3] McFee, J.; Das, Y.; and Ellingson, R., "Locating and Identifying Compact Ferrous Objects", *IEEE Transactions on Geoscience and Remote Sensing*, Vol. 28, No. 2, March 1990.
- [7.4] McFee, J.; and Das, Y.; "Determination of the Parameters of a Dipole by Measurement of the Magnetic Field", *IEEE Transactions on Antennas and Propagation*, Vol. AP-29, No. 2, March 1981.
- [7.5] Faugeras, Olivier, *Fundamentals in Computer Vision: an advanced course*, Cambridge University Press, 1983
- [7.6] McFee, J. and Das, Y., "Development of an Intelligent System to Locate and Identify Ferrous Objects", *Artificial Intelligence, Expert Systems and Languages in Modeling and Simulation*, C.A. Kulikowski, R.M. Huber and G.A. Ferrate (eds.), 1998
- [7.7] Das, Y., McFee, J., Toews, J., Stuart, G.C., "Analysis of an Electromagnetic Induction Detector for Real-Time Location of Buried Objects", *IEEE Transactions on Geoscience and Remote Sensing*, vol. 28, no. 3, May 1990.
- [7.8] Chesney, R.H., Das, Y., McFee, J.E., and Ito, M.R., "Identification of Metallic Spheroids by Classification of Their Electromagnetic Induction Responses", *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. PAMI-6, no. 6, November 1984.
- [7.9] Das, Y., McFee, J.E., and Chesney, R.H., "Determination of Depth of Shallowly Buried Objects by Electromagnetic Induction", *IEEE Transactions on Geoscience and Remote Sensing*, vol. GE-23, no. 1, January 1985

# **Part III**

## **Experimental Results**

## Antarctica 2000 Nomad Field Trials

In January 2000 A.D. Nomad was deployed to Elephant Moraine, Antarctica, for meteorite search field trials. These were the first real trials of the entire robot system, complete with autonomous search capability, sensor deployment and onboard classification, culminating in the first autonomous robotic field identification of a meteorite.

Elephant Moraine is located near the Trans-Antarctic Mountains. As the Antarctic ice sheets glacially outwards from the interior of the continent they are forced up against the mountains. As the ice rises it is ablated away by wind erosion, causing rocks and meteorites buried in the ice and transported from the continental interior to be stranded on the surface. Because of this concentration mechanism, certain areas such as Elephant Moraine are rich sources of meteorites. To date, many hundreds of samples have been retrieved there, including two Martian meteorites.



**Figure 8.1** *Elephant Moraine, Antarctica, so named because from the air it has the appearance of an elephant. Sub-glacial rocks cause the ice to flow upwards at the moraine. Winds ablate away the surface of the ice causing buried rocks and meteorites, originally transported from the continental interior, to be exposed on the surface.*

Most terrestrial rocks found lying on the ice at Elephant Moraine are Dolerites and Basalts, and some quartz Sandstones<sup>[8.1][8.2][8.3][8.4]</sup>. Dolerite and Basalt are igneous rocks, ranging from black to gray in color, and having minerals in common with many meteorites. Quartz Sandstones are sedimentary rocks, usually light colored, depending on how much contamination is present. Occasionally, samples of subglacial Travertine (“black Limestone”) and black Shale are found. As they look so much like meteorites, these terrestrial rocks are classic “meteowrongs”. Fortunately, these are rare enough that misclassifying them as meteorites and retrieving them is not cause for concern. Many of the Basalts and Dolerites found at Elephant Moraine showed signs of hydrothermal alteration, giving them a slightly red hue due to the presence of Iron Oxide (rust) on their surfaces. This was unexpected, since the rock sample set<sup>1</sup> used to train the Bayes network classifier prior to Antarctic deployment did not contain any examples of these rocks.

---

<sup>1</sup> Elephant Moraine rock sample set collected by and kindly made available by Professor Gunter Faure of the Ohio State University Department of Geology.

To get to Elephant Moraine, Nomad was first shipped to Christchurch, New Zealand and thence to McMurdo base on Ross Island by ski-equipped LC-130 Hercules military transport aircraft, landing at Williams Field, a summer airfield set up on the pack ice near the base. McMurdo, “MacTown” to the locals, is the main US Antarctic base, supporting expeditions to the whole continent, including Amundsen-Scott / South Pole station, the ANSMET meteorite collecting expeditions and Elephant Moraine. It has extensive facilities, including a bowling alley.

After assembly and testing at McMurdo, Nomad was stripped of its manipulator arm and sensor mast and hoisted by helicopter across the Ross Sea and the Trans-Antarctic Mountains to Elephant Moraine. This was the logistically most critical part of the operation, requiring a convoy of two helicopters for safety and a refueling stop at a supply depot. Helicopters were necessary, as Nomad does not fit inside the Twin Otter aircraft, capable of landing on unprepared stretches of ice and snow.

The expedition team subsequently flew to Elephant Moraine from McMurdo, ferrying equipment and personnel there over several Twin Otter flights. Camp was set up on ice sheets at the northern end of the moraine and the robot reassembled. This was hampered by high winds and low temperatures, making it difficult to handle metal connectors and pieces without bulky gloves. Furthermore, the cold almost caused the loss of the fiber optic cable connecting the spectrometer to the probe on the manipulator arm, and several components, including the pan-tilt mounting for the camera, were misaligned during helicopter transport and landing.

Elephant Moraine is too large an area to search exhaustively with Nomad, which searches along a strip approximately 1 m wide while traveling at 30 cm/s, stopping frequently to maneuver or deploy the spectrometer. For this reason it was necessary to first search do human searches of the moraine, either on foot or snowmobile, to identify smaller areas with meteorites where the meteorite search capability of Nomad could be evaluated in the short time available. Furthermore, parts of the moraine are not easily traversable by Nomad, or have such a high density of rocks that Nomad’s visual search algorithm is overwhelmed.

Ten meteorite search demonstrations were performed at Elephant Moraine, during which a total area of approximately 2500 m<sup>2</sup> were searched in around 16 hours spread over 10 days. Three meteorites were found and correctly classified during autonomous searches, out of 42 targets examined. An additional two meteorites were correctly classified during tests performed without patterned searches.

Sensor deployment was performed with mixed results. While Nomad’s control system [SAS reference] autonomously found new rock targets with a 79% success rate and deployed Nomad’s manipulator arm with a 72% success rate, it failed to obtain suitable close-up images of rocks. While zooming in the camera would lose sight of the targets, so it was only possible to obtain images in which the target rock was less than 5% of the image. This is different from the images used to train the classifier. Furthermore, the preponderance of bright snow and ice in the images caused extremely bad contrast, with all rocks looking essentially like meteorites: small, black and featureless. Consequently, every sample had to be examined with the spectrometer to determine if it was a meteorite. This problem was possibly caused by damage to the pan and tilt unit when Nomad was transported into the field slung underneath a helicopter. Furthermore, the cold caused the contact switch for sensing when the spectrometer probe is just above a rock sample to fail. Therefore, human intervention to prevent the spectrometer probe being driven into the ground was required each time a spectrum was obtained

## 8.1. First autonomous robotic meteorite find



**Figure 8.2** *First meteorite found by Nomad, an ordinary Chondrite, on the ice sheets by Elephant Moraine.*

Nomad made the first autonomous identification of a meteorite on January 22, 2000. An area 100 m x 100m where two meteorites were known to be was marked out, and a few terrestrial rocks added to the area. To prevent bias, the meteorite locations were known only to the field guide<sup>2</sup>, who did not otherwise participate in operating the robot. Nomad was commanded to start a search pattern. Each time it located a rock target (dark object on the ice) it would stop and acquire a higher resolution image of the target (Figure 8.3). This was used both to better estimate the target position and for initial classification. Subsequently Nomad maneuvered to bring the target within the arm workspace and the spectrometer probe was visually servoed to the target using the downward looking camera located on the arm wrist, immediately above the fiber optic probe. Upon taking a single spectrum, the arm was retracted and the spectrum sent to the onboard classifier, which returned a confidence value of the sample being a meteorite given the image and spectrum.

Nomad investigated rocks 7 times before stopping at a meteorite (Figure 8.3), for which it returned a confidence of 92%, 3 standard deviations above the confidences for the terrestrial rocks (Figure 8.4) and so distinguishing it. Assuming a decision threshold of 50% means the meteorite was detected and one of the rocks was falsely identified as a meteorite.

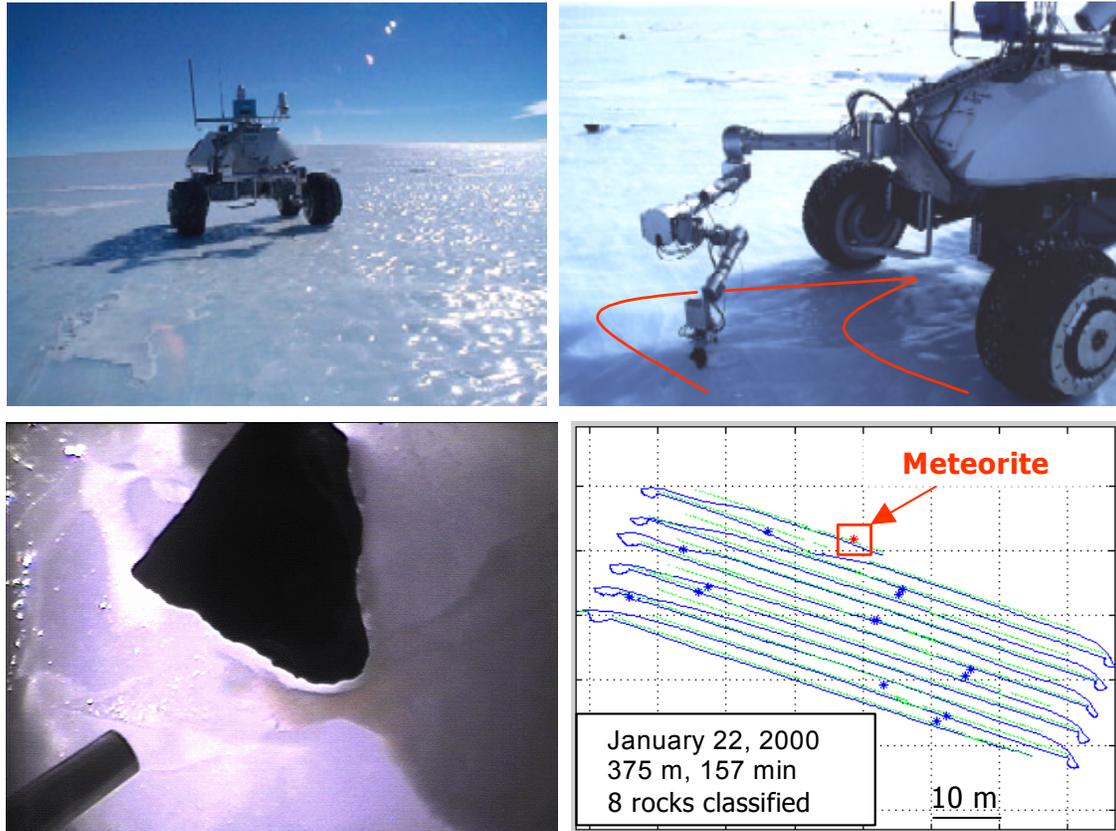
Figure 8.5(i) – (viii) shows the robot acquired images and spectra (along with correlation coefficients showing where peaks or troughs are present, and the classifier computed probabilities that the targets are meteorites given the image, the spectra or both. In the absence of sensor data the classifier assumes a 10% likelihood that a candidate rock is actually a meteorite.

Classification using the images alone is clearly inferior. Nomad was unable to get fully zoomed images of rocks without losing them altogether, probably because the camera pan and tilt mechanism was damaged in

---

<sup>2</sup> John Schutt, ANSMET field guide.

transit from McMurdo. The result is that images could not be centered properly on rocks, causing them to go off the edge as the zoom is increased. There is insufficient resolution to pick out much texture, and because the bright background dominates the scene, there is no contrast or color balance. Hence, all the rocks look like meteorites in the images: small, dark and featureless. This is confirmed by the classifier probabilities in Figure 8.5, all of which show an increased chance of a rock being a meteorite given the images.



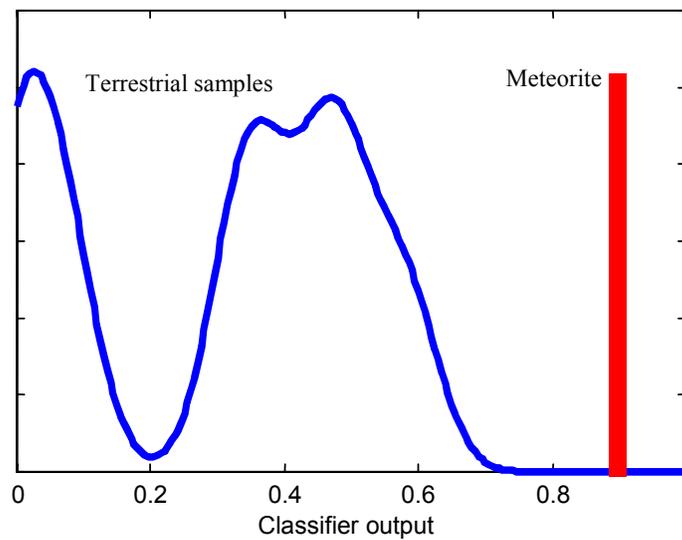
**Figure 8.3** First autonomous robotic meteorite find, in Elephant Moraine, Antarctica, in January 2000 AD. Nomad searched a 100 m x 100 m area where a meteorite was known to be. Upon reaching the meteorite (top left), a high resolution color image was obtained and classified. The spectrometer arm was then deployed (top right), using its own camera to visually servo the spectrometer fiber optic probe onto the meteorite (bottom left). From the image and the spectrum the classifier was 92% confident of the sample being a meteorite, distinguishing it by 3 standard deviations from the other rock samples examined (bottom left).

Examination of the spectra in Figure 8.5(i) – (iv) show bad readings around the blue-green part of the visible spectrum.

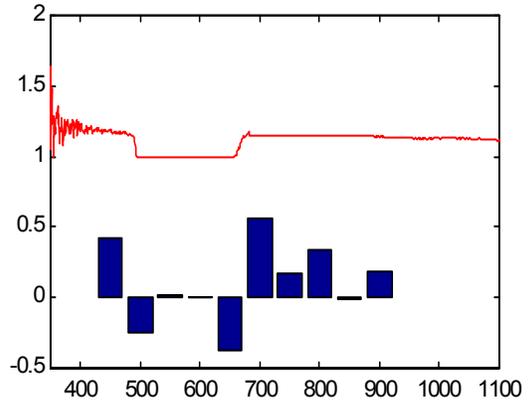
Sunlit blue ice is intense at these wavelengths, fading rapidly at longer wavelengths. These spectra would be consistent with bad positioning of the spectrometer probe so that sunlit ice as well as rock is in the fiber optic field of view. In the blue-green wavelengths the signal from the sunlit ice dominates, saturating the spectrometer detector. Therefore, no change is measured as the spectrometer lamps are switched on and off, giving the almost zero reflectance spectra observed. In the red and infrared parts of the spectrum the

signal is less, so the detector is not saturated and spectrometer algorithm removes the background illumination signal. That which remains is the rock signal, albeit reduced by the fraction of the spectrometer fiber optic probe field of view taken up by rock. Figure 8.5(vi) is clearly an ice reflectance spectrum, showing the spectrometer probe was not even placed near the rock sample. Coincidentally, all these spectra were obtained from small rock samples, while the others came from rocks (and a meteorite) twice as large.

*NOTATION: all spectra in this chapter are plotted along with their associated template correlation coefficients (chapter 5), plotted underneath as bars, centered on the center wavelength of the associated template. A positive value indicate a peak, a negative value a trough, and zero a flat spectrum. Spectrum plots are offset upwards by 1.0 so that they are not obscured by the bar plots. Furthermore, Figure 8.7 onward also plot spectra along with the associated (left to right) Intensity, Green and Red values plotted as bars alongside the spectrum.*



**Figure 8.4** Classifier confidence output distribution from 7 examinations of terrestrial samples and the confidence output on the single meteorite examined. The meteorite find is clearly distinguished by the classifier, which had confidence 3 standard deviations above the confidences on other samples.

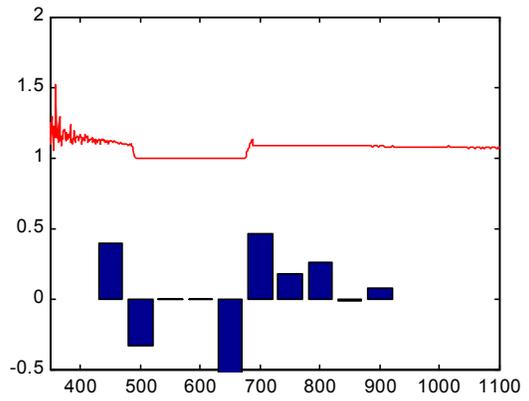
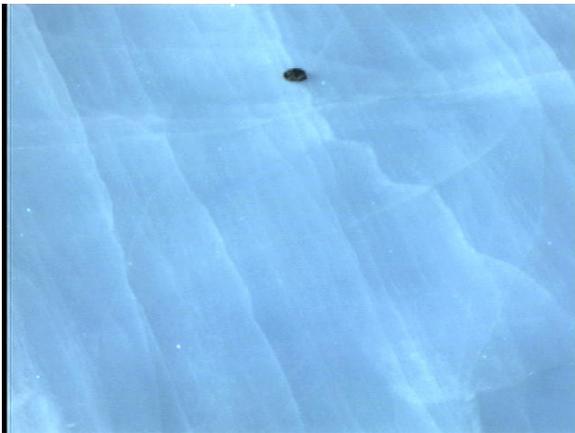


(i) Basalt fragment Bad spectrum, partly of ice.

$$P(\text{meteorite} \mid \text{image}) = 0.31$$

$$P(\text{meteorite} \mid \text{spectrum}) = 0.29$$

$$P(\text{meteorite} \mid \text{image \& spectrum}) = 0.49$$

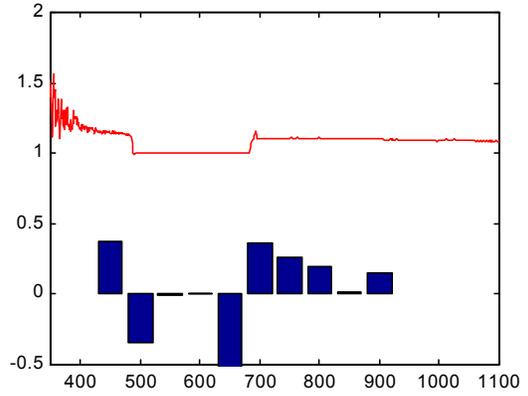
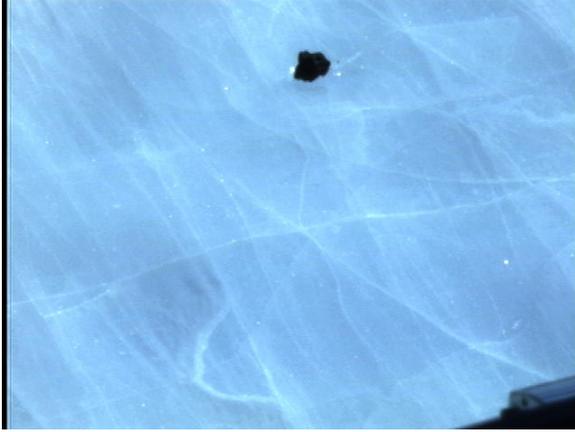


(ii) Same Basalt fragment as 1 above. Bad spectrum, partly of ice.

$$P(\text{meteorite} \mid \text{image}) = 0.37$$

$$P(\text{meteorite} \mid \text{spectrum}) = 0.19$$

$$P(\text{meteorite} \mid \text{image \& spectrum}) = 0.37$$

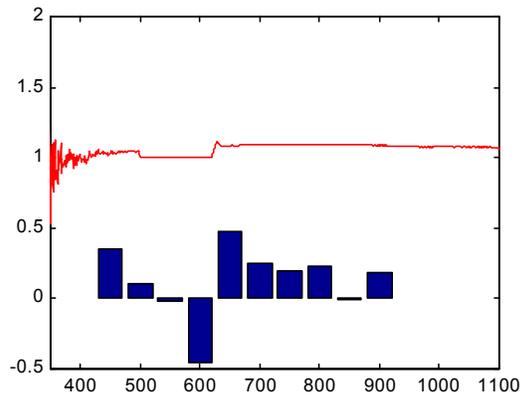
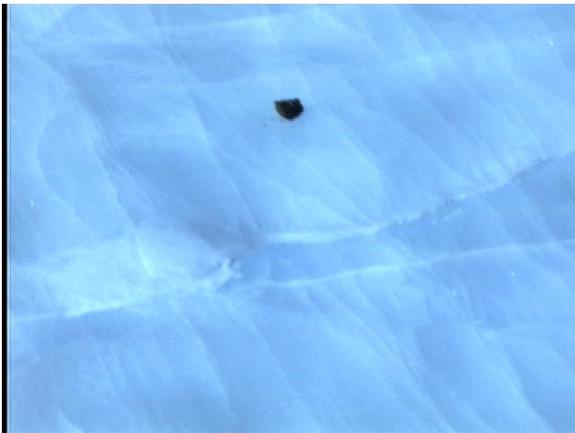


(iii) Basalt fragment Bad spectrum, partly of ice.

$$P(\text{meteorite} \mid \text{image}) = 0.24$$

$$P(\text{meteorite} \mid \text{spectrum}) = 0.36$$

$$P(\text{meteorite} \mid \text{image \& spectrum}) = 0.58$$

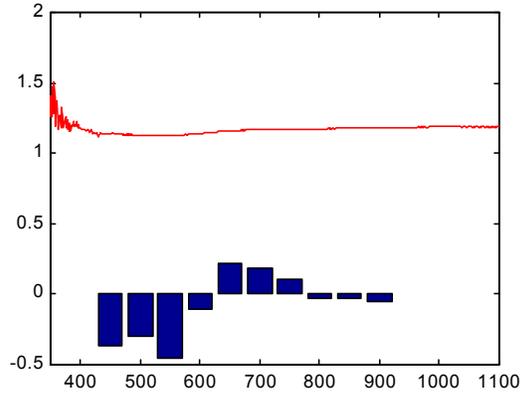
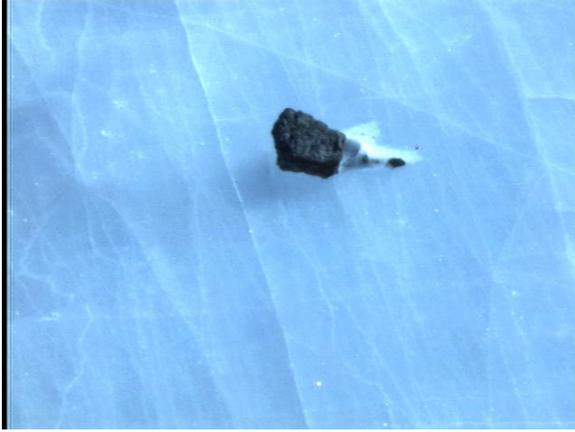


(iv) Basalt fragment (1. or 2.) Bad spectrum, but better than previous spectra.

$$P(\text{meteorite} \mid \text{image}) = 0.31$$

$$P(\text{meteorite} \mid \text{spectrum}) = 0.22$$

$$P(\text{meteorite} \mid \text{image \& spectrum}) = 0.33$$

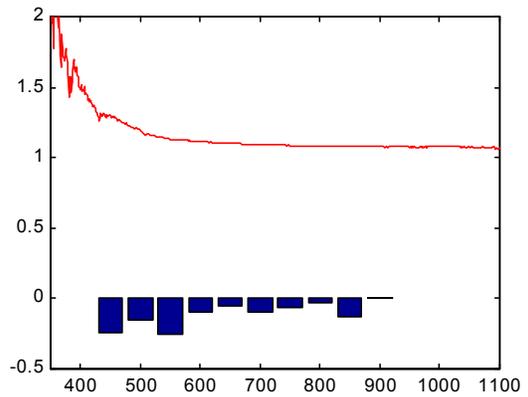


5. Tan siltstone. Good spectrum

$$P(\text{meteorite} \mid \text{image}) = 0.31$$

$$P(\text{meteorite} \mid \text{spectrum}) = 0.002$$

$$P(\text{meteorite} \mid \text{image \& spectrum}) = 0.006$$

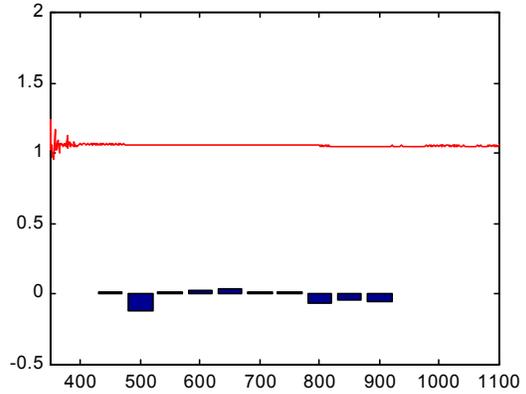
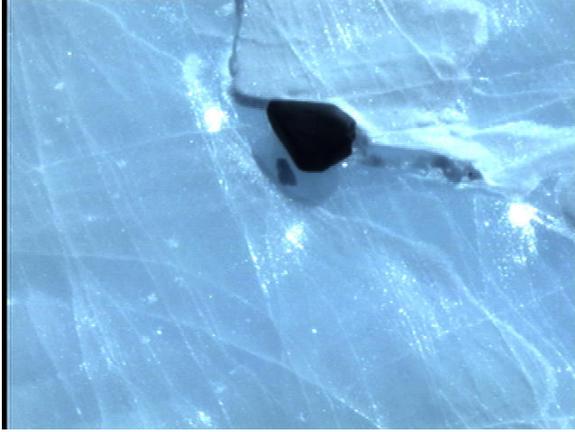


(vi) Dolerite. Spectrum is of the ice sheet.

$$P(\text{meteorite} \mid \text{image}) = 0.83$$

$$P(\text{meteorite} \mid \text{spectrum}) = 0.02$$

$$P(\text{meteorite} \mid \text{image \& spectrum}) = 0.08$$

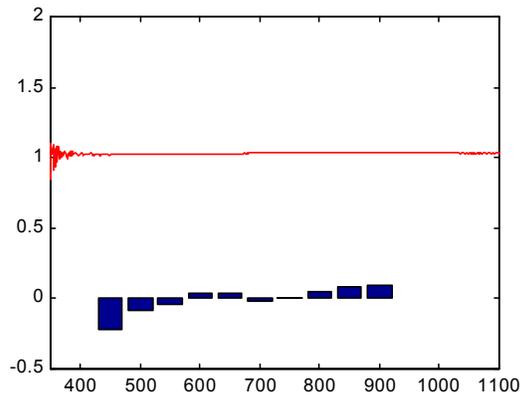
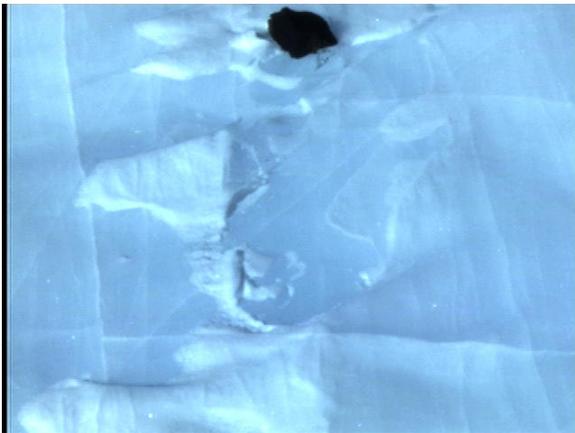


(vii) Black shale (easily confused with meteorite).

$$P(\text{meteorite} \mid \text{image}) = 0.14$$

$$P(\text{meteorite} \mid \text{spectrum}) = 0.28$$

$$P(\text{meteorite} \mid \text{image \& spectrum}) = 0.46$$



(viii) Meteorite! Ordinary Chondrite (ANSMET 2000 field no. 12067)

$$P(\text{meteorite} \mid \text{image}) = 0.23$$

$$P(\text{meteorite} \mid \text{spectrum}) = 0.71$$

$$P(\text{meteorite} \mid \text{image \& spectrum}) = 0.92$$

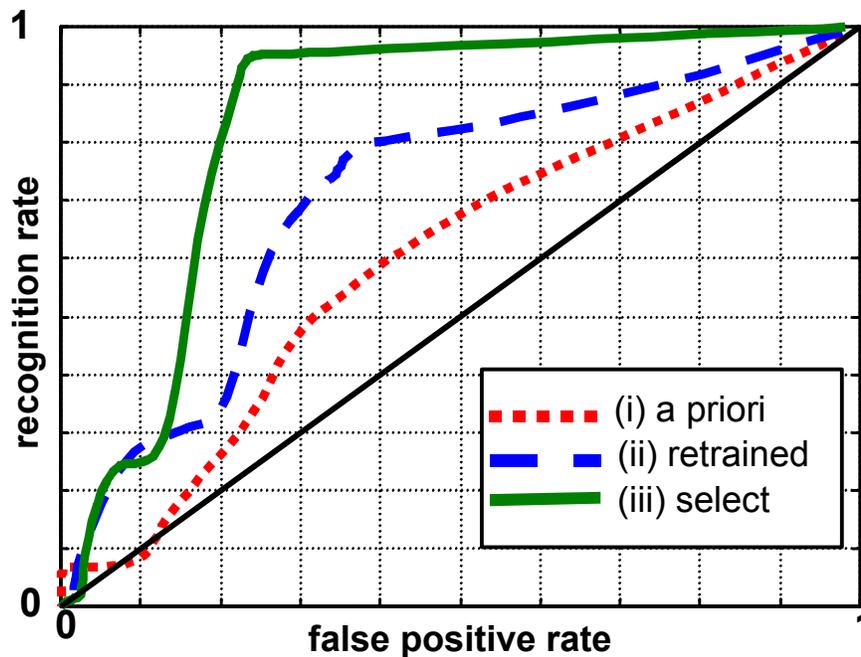
**Figure 8.5** Images, spectra and classifier computed probabilities of being a meteorite obtained using the classifier trained only with laboratory data obtained prior to arrival in Antarctica.

These tests show the classifier rejecting as meteorites those samples from which a good spectrum was obtained, including the black shale of Figure 8.5(vii), a rare specimen easily confused with meteorites but deliberately placed in the test area. A lower limit on the size of rocks that can be profitably examined was established. Furthermore, the camera images were shown to be unusable by the classifier based on close up higher quality images of rocks and meteorites in the laboratory, and so disengaged for all subsequent tests.

## 8.2. Overall classification performance

The previous tests, while showcasing the system, were too small to be statistically rigorous and indicated flaws in the system. Subsequent spectral readings of the rock samples were less encouraging and indicated the need to retrain the classifier with spectra autonomously acquired in the field by Nomad. The types of rocks found around the moraine were not identical to those in the sample sets used to train the classifier prior to Antarctic deployment.

To fully test the classification system Nomad acquired spectra from 5 other meteorites and 40 rock samples at the edge of the moraine. Multiple spectra were taken of each sample, and classified with the *retrained* Bayes network classifier. Note that the rocks used for retraining the classifier were *not* in this data set. The overall classification performance (using spectra only) is indicated by the ROC performance curves in Figure 8.6, generated from this test set. To ensure statistical fairness, each spectrum classification is weighted by the reciprocal of the number of spectra taken from the source rock or meteorite sample. Parzen density estimation [8.6] was used to smooth the curves. Figure 8.6(i) indicates performance using the *a priori* classifier built from the lab data prior to the Antarctic demonstration, which clearly does not match the performance on laboratory data (Chapter 4). There is a notable improvement (Figure 8.6(ii)) in classification after the aforementioned Bayes network classifier is augmented with field data acquired by Nomad from 8 rocks and 2 meteorites in the moraine (*not* in the test set used to generate the ROC curve)



**Figure 8.6** ROC curve demonstrating performance of Nomad meteorite classifier on rocks and meteorites at Elephant Moraine, Antarctica, in January 2000. This was generated from spectra acquired from 5 meteorites and 40 native terrestrial rocks by Nomad in Elephant Moraine. ROC curve (i) is the performance of the Bayes network classifier trained only with the laboratory data acquired earlier. Curve (ii) shows the improvement after this classifier is further trained with independent field data acquired by Nomad in the area. Curve (iii) is performance of the classifier in (ii) after all the hydrothermally altered iron oxide stained Basalt's and Dolerite's have been removed from the test set of rocks.

Even with retraining, classification is systematically bad for a particular class of rocks (hydro-thermally altered dolerites and basalts) that occurred in the Elephant moraine. These rocks are stained red with surface iron oxide (rust), which tends to dominate their reflectance spectra. Meteorites usually also contain iron and are therefore likely to have some surface iron oxide, giving rise to similar spectral features. Furthermore, these rocks were absent from the laboratory training set and not initially expected in the area. Excluding them from the test set shows that classification is significantly better for all the remaining rocks (Figure 8.6(iii)).

Choosing accept/reject meteorite probability thresholds for a 20% false positive rate the 3 curves in Figure 8.6 implies meteorite recognitions rates of 25%, 35% and 80% for curves (i), (ii) and (iii) respectively. On the 40 rock, 5 meteorite test set, examined over 4 days, 8 terrestrial rocks and 1, 2 or 4 meteorites would have been collected by the robot.

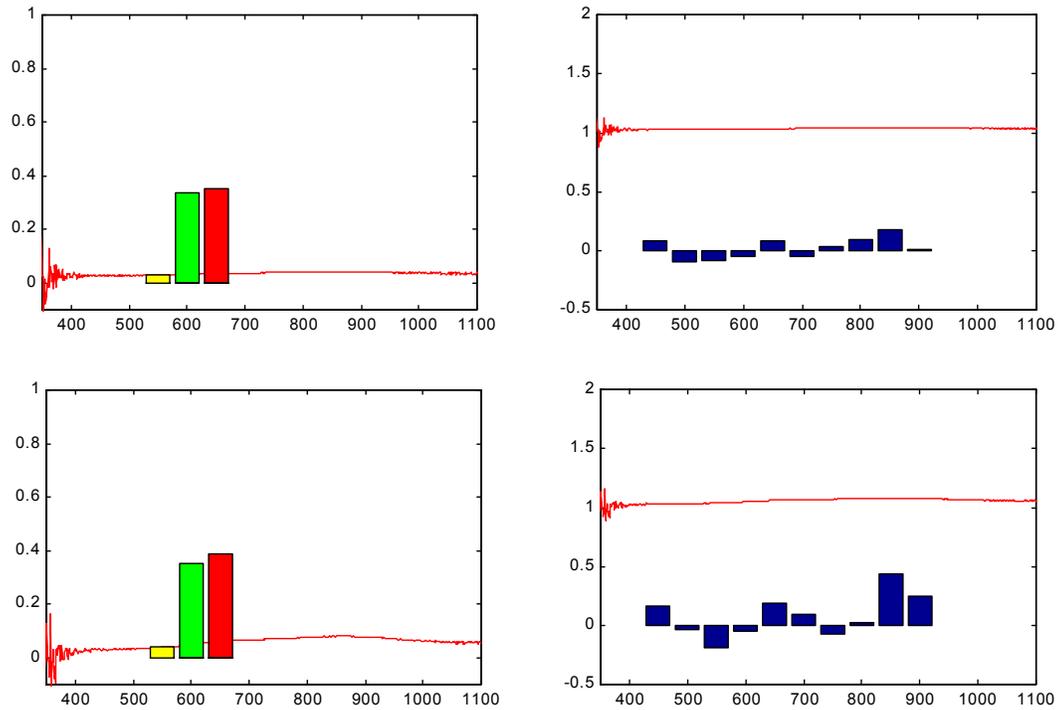
The following figures (Figure 8.7 to Figure 8.20) showcase the various rocks and meteorites that Nomad examined and classified in Elephant Moraine, using the *retrained classifier* (Figure 8.6(ii)). Spectra from all the 5 meteorites and select terrestrial rocks in the test set are shown. Along with spectra are the spectrum-template correlation coefficients used to detect the presence of peaks or troughs (right spectra) and their Red, Green and Intensity coefficients (left spectra). Refer to Chapter 5 for a complete description of these quantities. Images, although not used for classification, are either from Nomad's high-resolution camera or taken with a handheld digital camera for documentation purposes.

From these spectra, it is apparent that the main features distinguishing meteorite spectra are:

- Low overall visible intensity ( $I < 0.1$ ).
- A peak around 900 nm, indicated by a strong positive correlation feature.
- Flat spectra, except for the peak at 900 nm, as indicated by the small template correlation coefficients.
- Gradual and steady increase in reflectance with wavelength, up till 900 nm. This is indicated by a red spectral coefficient greater than the green spectral coefficient.

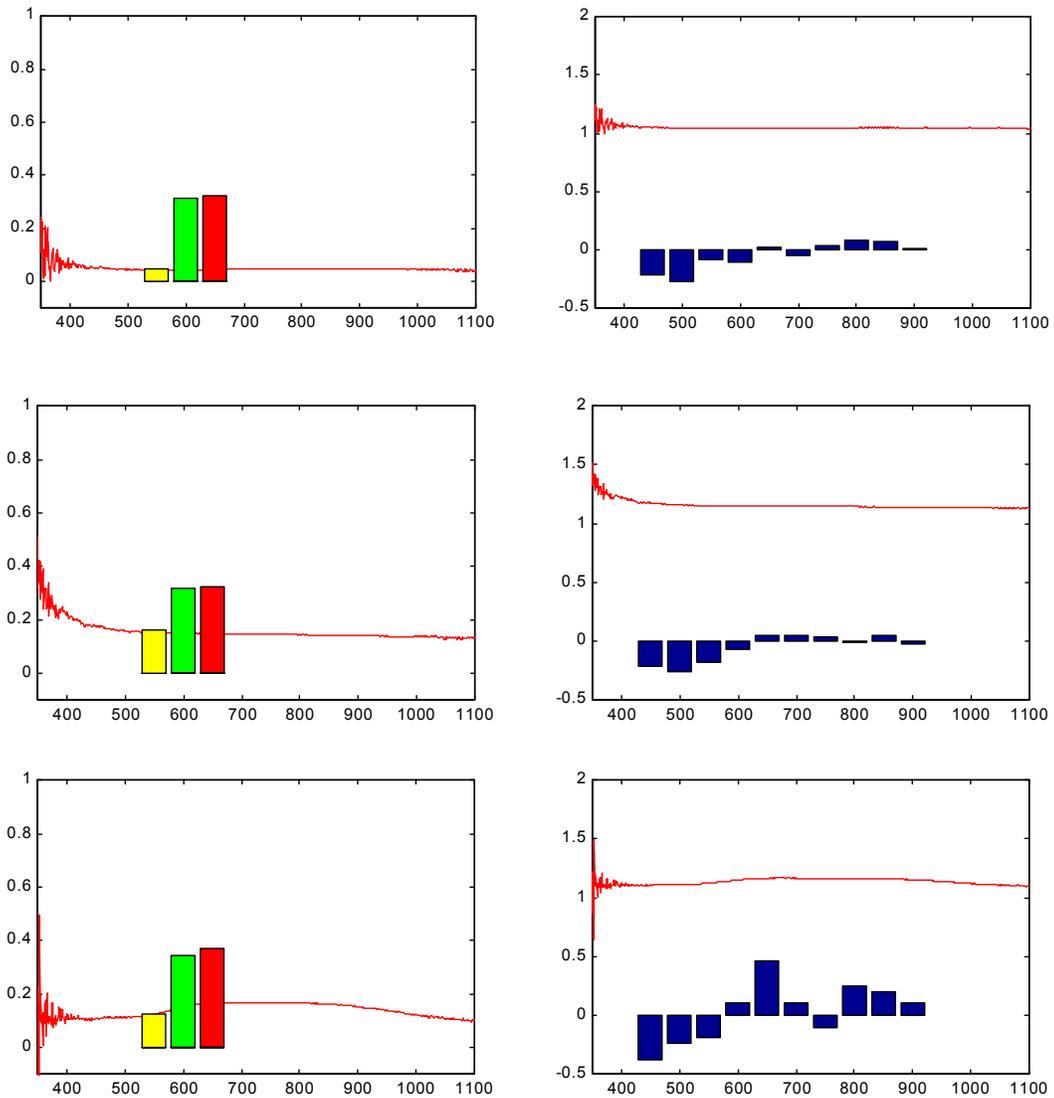
These features are exemplified by the meteorite in Figure 8.10, and to a lesser extent by Figure 8.11, a rare Achondrite meteorite. Unfortunately, some of these features are shared by the hydrothermally altered Basalts and Dolerites (Figure 8.15), particularly the strong peak at 900 nm and the steadily increasing reflectance with wavelength. This explains why they are so commonly misclassified as meteorites, and furthermore why otherwise clear meteorite spectra are sometimes not recognized as such (Figure 8.11 top). The spectra of the hydrothermally altered rocks are due mainly to surface iron oxide (rust), which tends to dominate the reflectance. Most meteorites contain metallic iron and so will also have traces of surface iron oxide.

**Elephant moraine meteorites:**



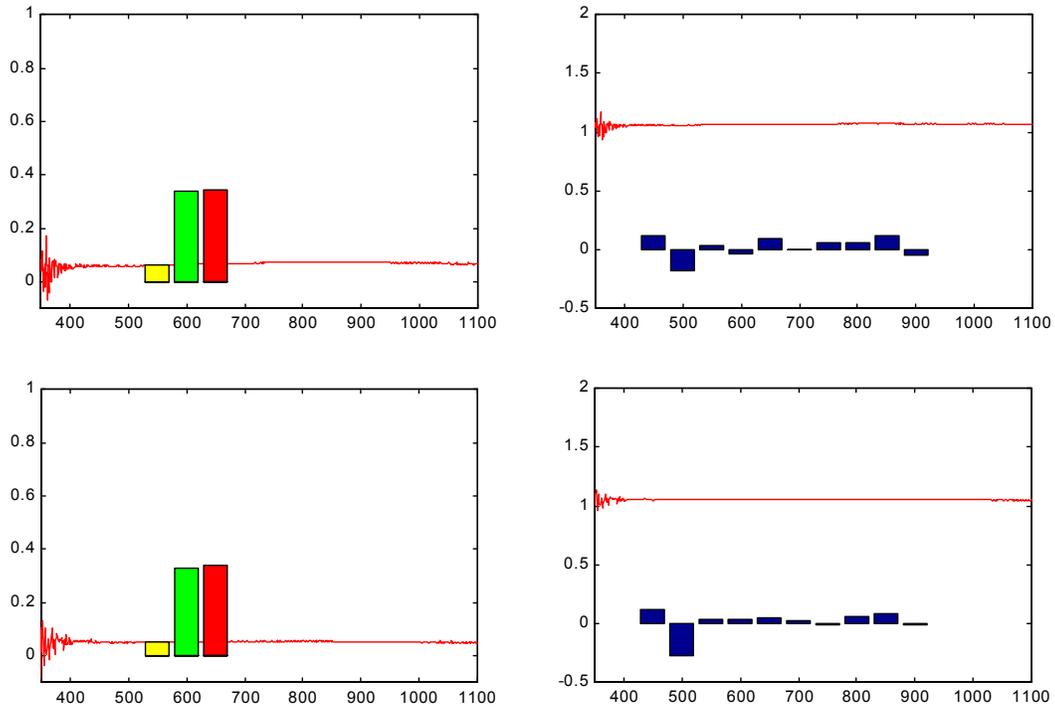
**Figure 8.7** Ordinary Chondrite with large tan inclusion (EET 99401, field no. 1131). Note how flat and black the spectrum is. These readings also appear to have a little noise, most likely due to the low signal level.

top: $P(\text{meteorite} \mid \text{spectrum}) = 0.23$ .
bottom: $P(\text{meteorite} \mid \text{spectrum}) = 0.39$



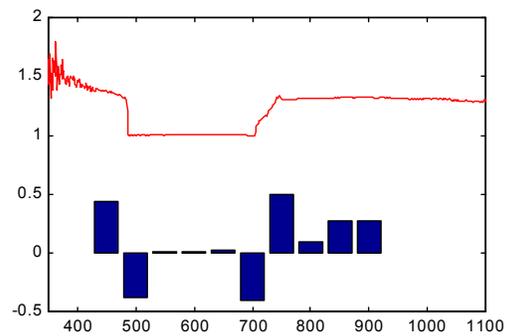
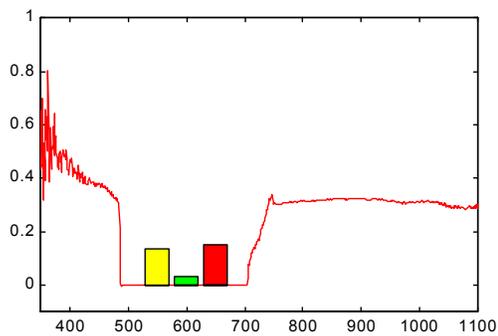
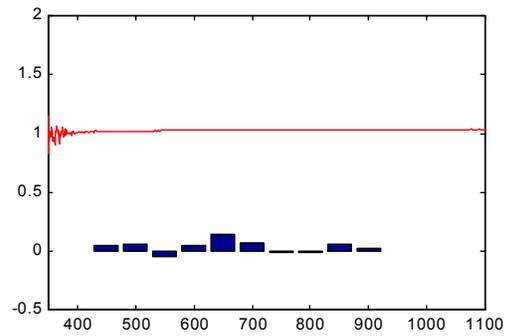
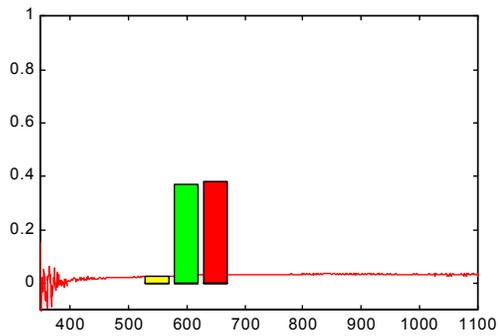
**Figure 8.8** Ordinary Chondrite meteorite (EET 99403, field no. 12097). Out of 5 spectra, the top is the only spectrum from this meteorite classified correctly. The other spectra give meteorite probabilities around 0.07 (middle and bottom spectra). The middle spectrum shows signs of snow, being excessively blue. The bottom spectrum appears very rocklike with the strong central peak at 650 nm, very unusual for meteorites.

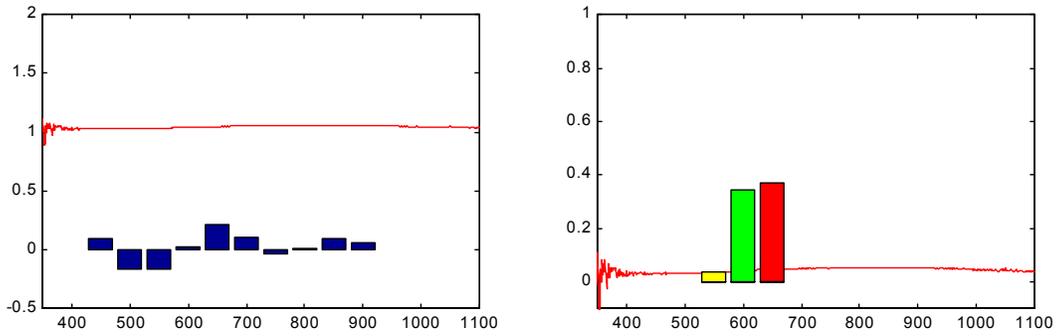
top:	$P(\text{meteorite} \mid \text{spectrum}) = 0.50.$
middle:	$P(\text{meteorite} \mid \text{spectrum}) = 0.06$
bottom:	$P(\text{meteorite} \mid \text{spectrum}) = 0.03$



**Figure 8.9** Ordinary Chondrite meteorite (EET 99404, field no. 12372). Of 3 spectra obtained, the top shows best results. The other two (bottom) appear consistent with a meteorite but are nonetheless improperly classified, perhaps because the 900 nm is not strong enough.

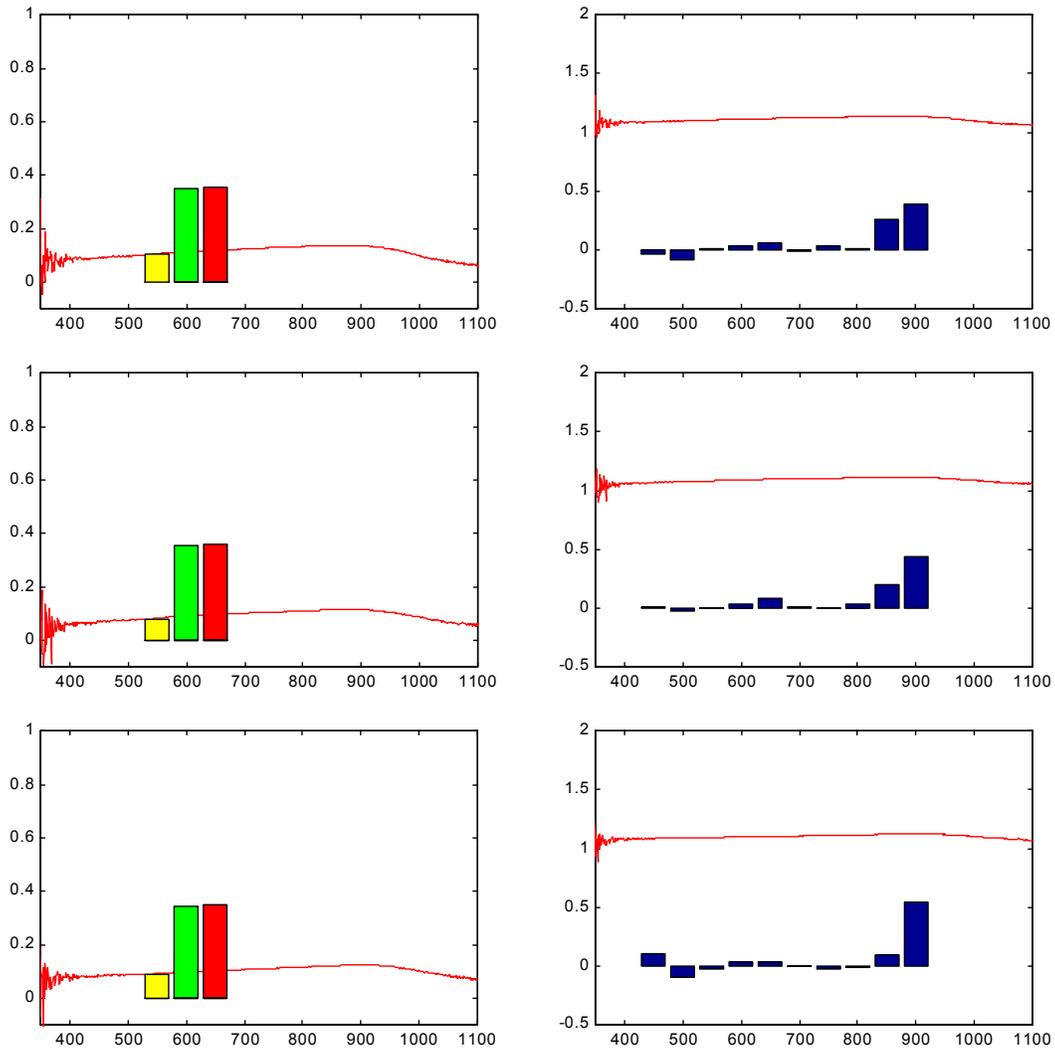
top: $P(\text{meteorite} \mid \text{spectrum}) = 0.33$ (1/3 spectra).
bottom: $P(\text{meteorite} \mid \text{spectrum}) = 0.05$ (2/3 spectra)





**Figure 8.10** Ordinary Chondrite meteorite (EET 99406, field no. 12387). The top spectrum is very well classified. Note the low intensity, generally flat spectrum with a slight peak at 9800 nm. The middle spectrum is a clear example of a measurement error. The bottom spectrum is the most typical of the spectra from this meteorite, which are all classified well.

top:	$P(\text{meteorite} \mid \text{spectrum}) = 0.73$ (best spectra)
middle:	$P(\text{meteorite} \mid \text{spectrum}) = 0.01$ (bad spectrum)
bottom:	$P(\text{meteorite} \mid \text{spectrum}) = 0.60$ (3/5 spectra)



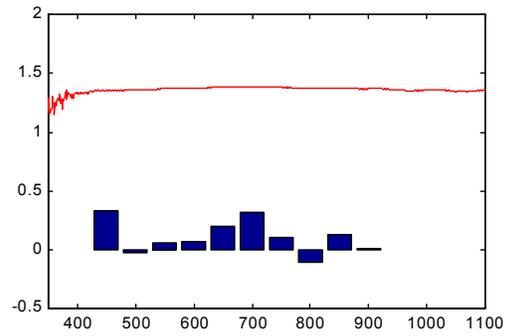
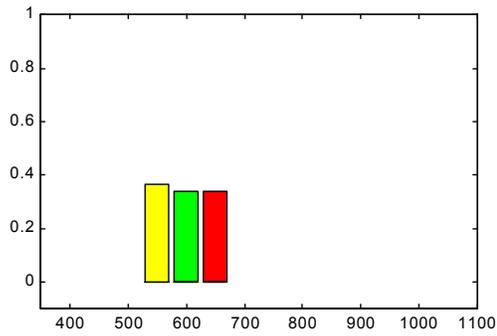
**Figure 8.11** Achondrite meteorite (EET 99400, field no. 12322). Achondrite meteorites are precious and unusual. This particular sample is especially so. Nonetheless, it is still thought to be a likely meteorite in 3/4 spectra obtained.

top:  $P(\text{meteorite} \mid \text{spectrum}) = 0.06$  (1/4 spectra)

middle:  $P(\text{meteorite} \mid \text{spectrum}) = 0.40$  (2/4 spectra)

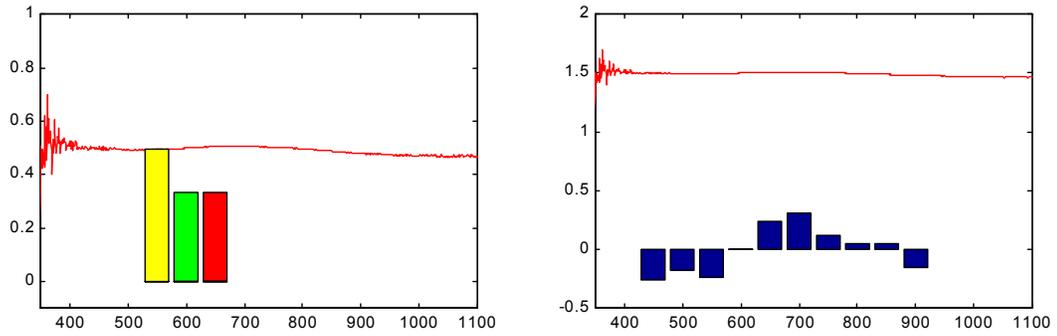
bottom:  $P(\text{meteorite} \mid \text{spectrum}) = 0.28$

**Elephant moraine native terrestrial rocks:**



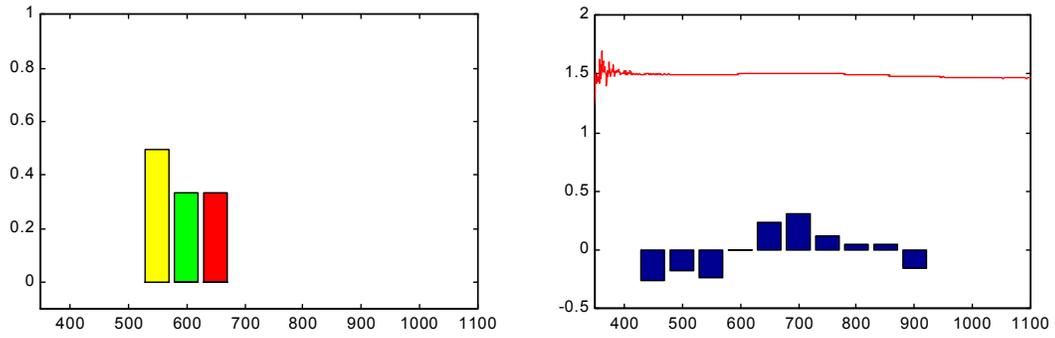
**Figure 8.12** Slightly hydrothermally altered basalt rock (M0).

$P(\text{meteorite} \mid \text{spectrum}) = 0.05$



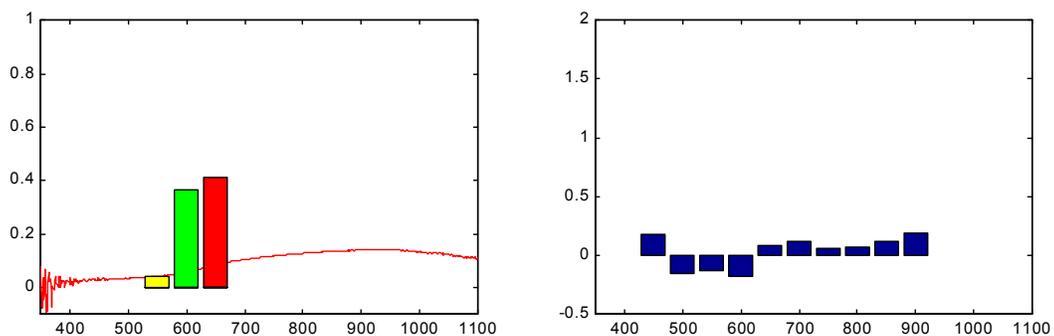
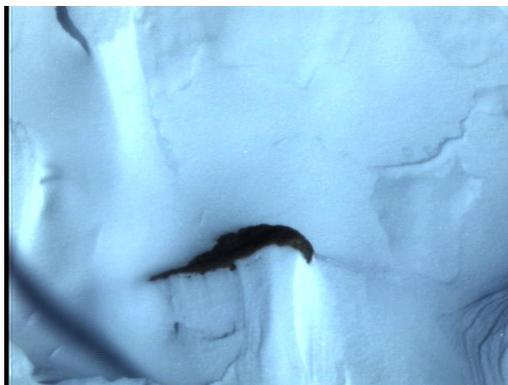
**Figure 8.13** Dolerite rock (M1).

$$P(\text{meteorite} \mid \text{spectrum}) = 0.02$$



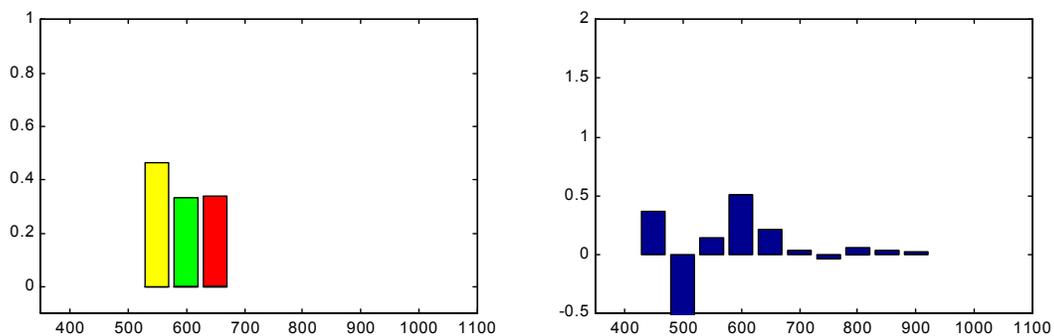
**Figure 8.14** Sandstone (M3).

$$P(\text{meteorite} \mid \text{spectrum}) = 0.05$$



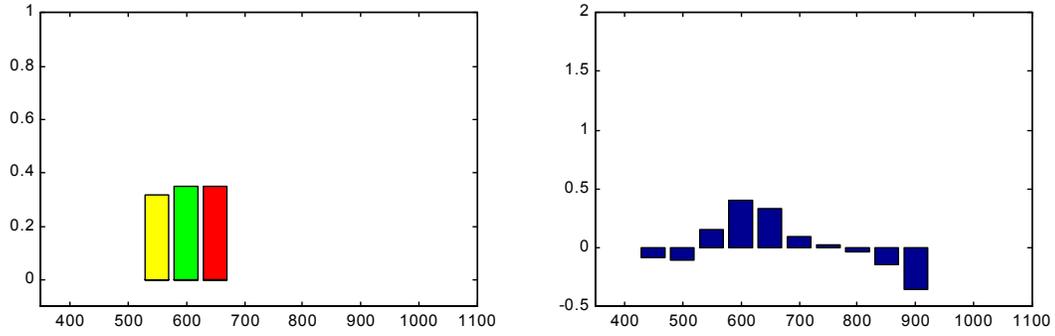
**Figure 8.15** Hydrothermally altered, Iron Oxide (red) stained Dolerite (M16). Note low measured albedo, relatively uniformly increasing reflectance with wavelength and peak at 900 nm. The spectrum is dominated by the Iron Oxide on the sample surface.

$$P(\text{meteorite} \mid \text{spectrum}) = 0.42$$



**Figure 8.16** Basalt rock (M21).

$$P(\text{meteorite} \mid \text{spectrum}) = 0.008$$

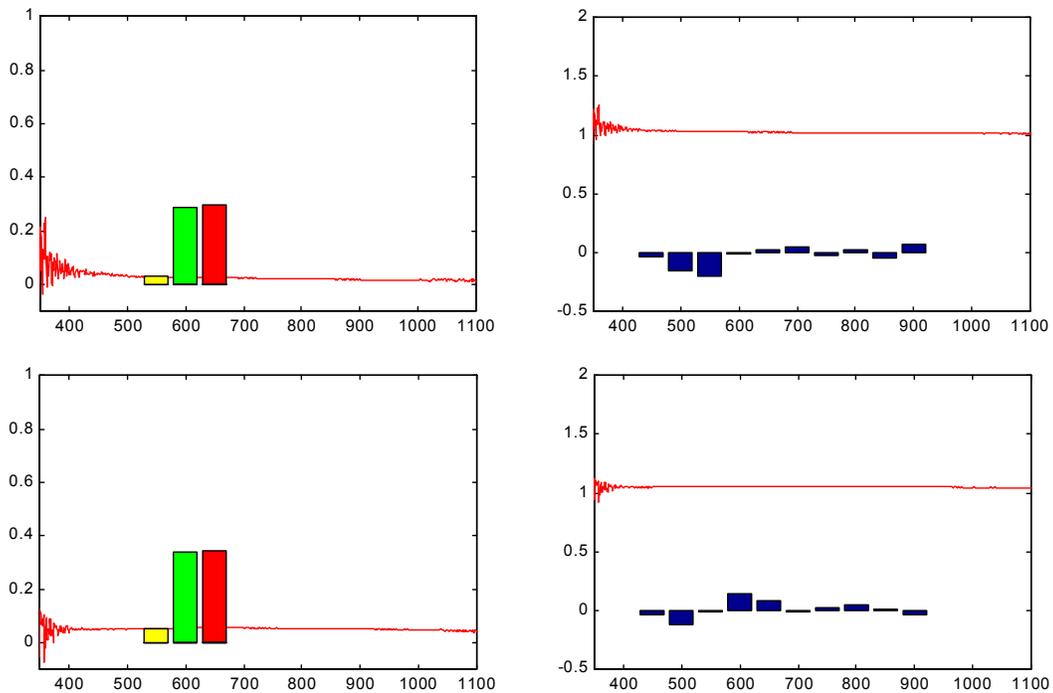


**Figure 8.17** Chert or hydrothermally altered lake sediment, or shale (M26).

$P(\text{meteorite} \mid \text{spectrum}) = 0.04$
---

**Elephant moraine ‘meteowrongs’:**

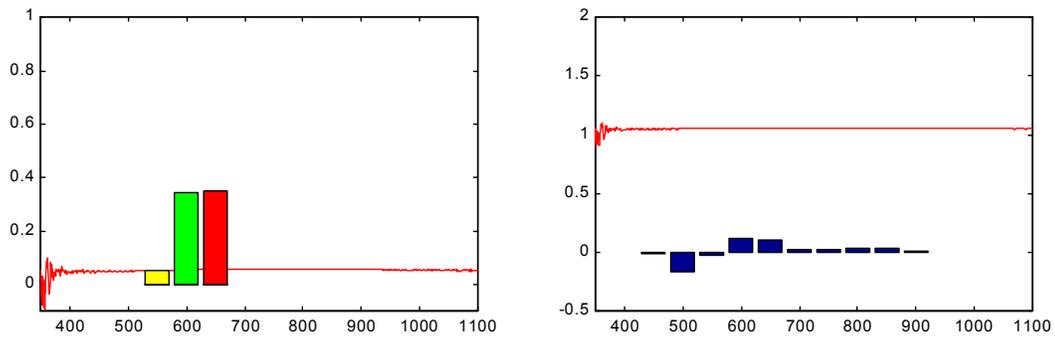
“Meteowrongs” are terrestrial rocks looking very much like meteorites to human searchers. They are also harsh tests for the robotic meteorite classifier. Two very unusual fused black sandstone rocks (Figure 8.18 and Figure 8.19) were found at Elephant Moraine, not far from where meteorites were found. When Nomad was directed to them it gave their spectra a higher than normal probability of being from meteorites. Although they have since been verified as terrestrial rocks, their spectra is indeed very much like that from a meteorite: black and flat, although they lack the peak at 900 nm.



**Figure 8.18** Peculiar black sandstone like rock with melted appearance (EET 99407, field no. 12346).

$$P(\text{meteorite} \mid \text{spectrum}) = 0.35 \text{ (mean of 6 spectra is 0.35)}$$

$$P(\text{meteorite} \mid \text{spectrum}) = 0.43$$



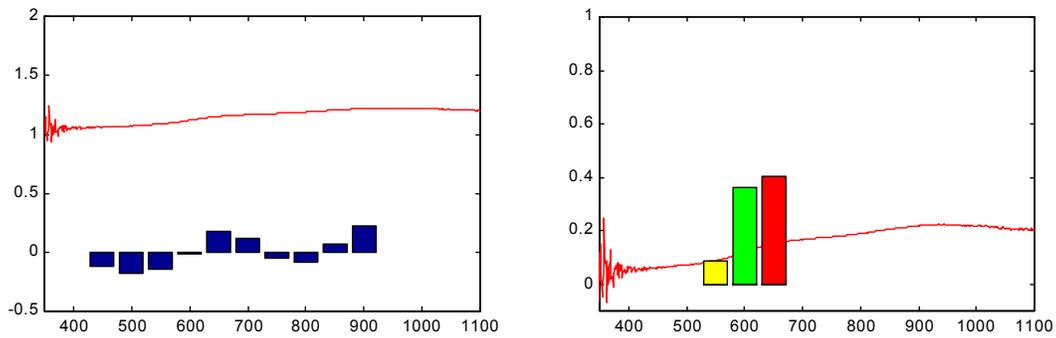
**Figure 8.19** Another peculiar black sandstone like rock with melted appearance (EET 99402, field no. 12341), possibly a fragment of EET 99407 (Figure 8.18)

$$P(\text{meteorite} \mid \text{spectrum}) = 0.41 \text{ (2 spectra, mean probability} = 0.34)$$

Being statistically anomalous, these rocks were not included in the test data set for the ROC performance curves in Figure 8.6.

### An Australian meteorite in Antarctica

For immediate testing of Nomad's classification system in the field, a meteorite<sup>3</sup> from Kulnine, Australia, was taken with the expedition to Antarctica. Figure 8.20 shows its spectrum, acquired by Nomad at Elephant Moraine. It well illustrates the features that most distinguish meteorite spectra from those of the terrestrial rocks commonly found in Antarctica.



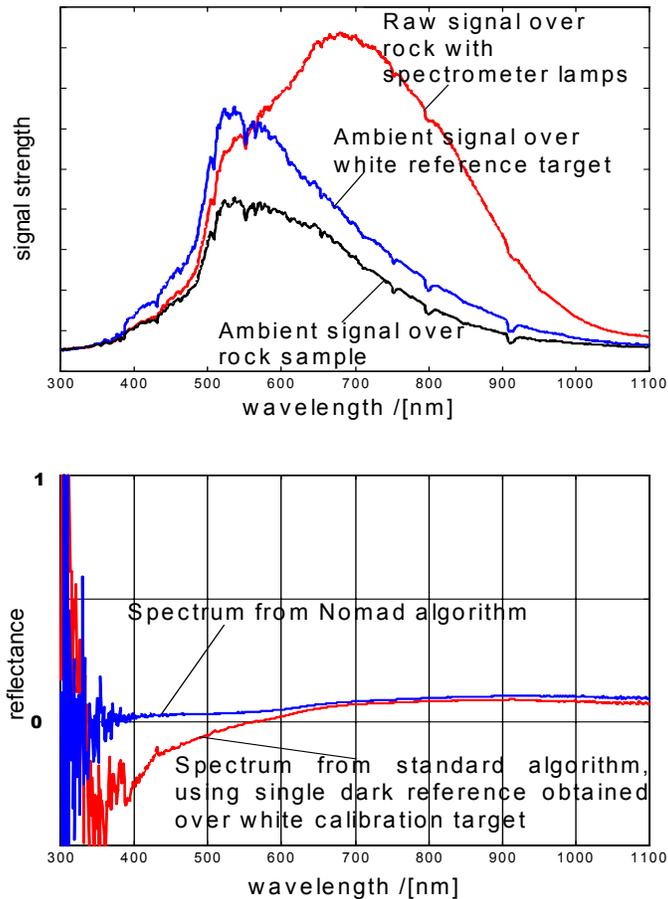
**Figure 8.20** Ordinary Chondrite meteorite (Kulnine, Australia), placed in front of Nomad for trouble shooting purposes. This result is NOT included in the ROC curve of Figure 8.6 as it was one of the meteorites used to field retrain Nomad's classifier.

$$P(\text{meteorite} \mid \text{spectrum}) = 0.65$$

<sup>3</sup> Kindly lent by Professor William Cassidy, University of Pittsburgh Department of Geology.

### 8.3. Spectrometer performance

The spectrometer was crucial to the successful classification of rocks and meteorites. The instrument on Nomad performed significantly better than what would have been possible using an unmodified off the shelf unit. Due to placement inaccuracies and the environment, the effect of ambient illumination was significant, but handled well by the active illumination system (**Figure 5.5**).



**Figure 8.21** Top: Unprocessed spectral signal from a sample rock with lamps illuminated, and spectral signals with lamps off when probe is over the white reflectance standard and over the rock. Note the strong signal due to ambient sunlight. Bottom: Computed reflectance spectra using Nomad's spectrometer algorithm and a traditional algorithm. Negative values for spectrum are nonsensical. The high noise below 400 nm is due to the low UV illumination from the spectrometer lamps.

The problems occurred when the probe was not adequately positioned over the sample. While the spectrometer head design allows considerable margins in probe placement accuracy, the sample must still fill the probe field of view. This was mostly a problem with small (approximately 3 cm or less) samples, for which the probe was sometimes placed over ice (Figure 8.5(iv)) or with a considerable amount of the sunlit ice and snow background in the field of view, causing detector saturation (Figure 8.5(i)-(iii)). Note that in the latter case the spectrometer design was still able to attain a reasonable spectrum at the longer wavelengths where the ambient lighting is not strong enough to saturate the detector.

As mentioned earlier, the cold caused the failure of the contact switch for sensing when to stop lowering the spectrometer onto a rock. Therefore, none of the spectral measurements was truly autonomous. A human was needed to for this task. Furthermore, experience proved that this arrangement with a contact switch is insufficient to reliably place the spectrometer probe at a set distance above a sample. With jagged rocks it is easy for the probe to be positioned at a very viewing angle with respect to the rock surface.

Occasionally some rock samples were partially covered in snow, but this was rarely a problem since the spectrometer autonomous placement algorithm attempts to visually servo the probe to those parts of an image most unlike the snow backgrounds. Rocks totally covered in snow were usually not found in the first place.

## **8.4. Camera performance**

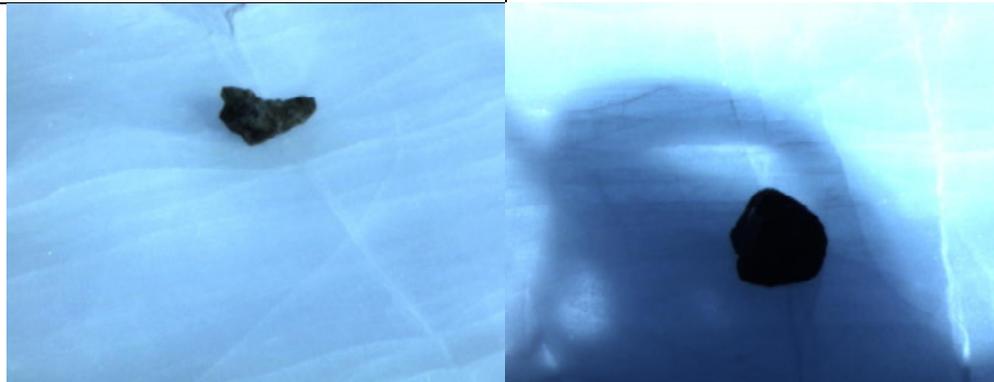
Nomad's camera was a disappointment for classification of rocks in the field. As indicated in 7.1 a problem prevented the acquisition of very high-resolution images of the rocks or meteorites. This resulted in a greater than expected part of the images being the bright ice or snow background. The camera auto-iris would attempt to compensate for this, optimizing the aperture for the bright background, not the rock samples, as had been the case for all the previous laboratory images of meteorites and rocks. This effect is clearly visible in all the robot images, including those taken for the image rock/ice segmentation calibration (Figure 8.22). Consequently all rocks looked much darker and without texture than the rocks and meteorites in the laboratory images used to train the classifier for visual classification. The unfortunate result is that almost every rock seen by Nomad at Elephant Moraine appeared to be a meteorite to the classifier.

Although visual classification with Nomad in Elephant Moraine was a failure, the algorithm for segmenting rocks and meteorites from the snow and ice background was a success, working well on all rocks on both the ice and snow. The background was usually either blue ice or clean fresh snow, and most of the rocks were distinctly darker than those seen in other parts of Antarctica.



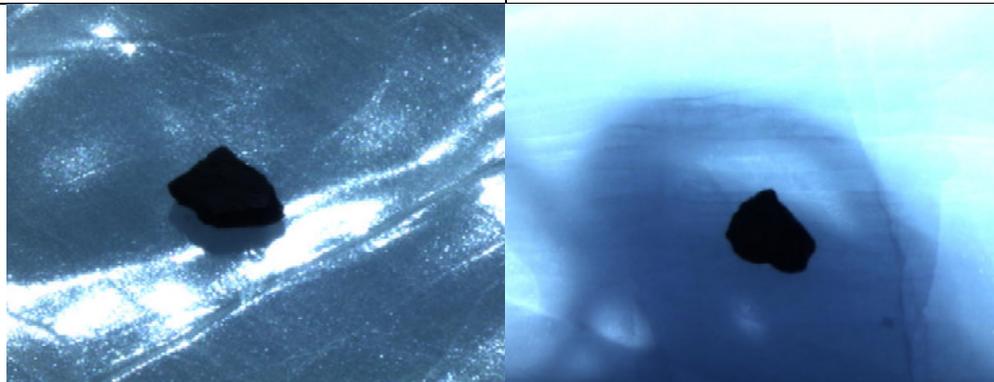
(i) Black Basalt on snow in direct sunlight.

(ii) Grey Dolerite on blue ice in direct sunlight



(iii) Sandstone on blue ice in direct sunlight.

(iv) Grey Dolerite on blue ice in shadow



(v) Meteorite (Kulnine) on blue ice, direct sunlight with specular highlights.

(vi) Meteorite (Kulnine) on blue ice, in shadow.

**Figure 8.22** Calibration images used to learn distribution of rock and ice pixel characteristics so that rocks may be segmented from the background during image recognition of meteorites.

## 8.5. References

- [8.1] K.S. Taylor, *Lithologies and Distribution of Clasts in the Elephant Moraine, Allan Hills South Victoria Land, Antarctica*, M.S. Thesis submitted to Kent State University Graduate College

- [8.2] Faure, G., Kallstrom, M.L., and Mensing, T.M., "Classification and Age of Terrestrial Boulders in the Elephant and Reckling Moraines", *Antarctic Journal of the U.S.*, 19(5), pp28-29, 1984
- [8.3] Faure, G., and Taylor K.S., "The Geology and Origin of the Elephant Moraine on the East Antarctic Ice Sheet", *Antarctic Journal of the U.S.*, 20(5), pp 11-12, 1985
- [8.4] Faure, G., Taylor, K.S., and Jones, L.M., "Hydrothermal Calcite in the Elephant Moraine", *Antarctic Journal of the U.S.*, 21(5), pp 21, 1986
- [8.5] Blatt, H., Jones, R.L., "Proportions of Exposed Igneous, Metamorphic, and Sedimentary Rocks", *Geological Society of America Bulletin*, v. 86, pp 1085-1088, August 1975
- [8.6] Fukunaga, K., *Introduction to Statistical Pattern Recognition*, 2<sup>nd</sup> ed, Academic Press, 1990.

## **Part IV**

# **Extensions of Bayes Network Generative Model**

## Autonomous Profiling of Unknown Environments

### 9.1. Introduction



**Figure 9.1** *Nomad investigates rocks while traversing the Antarctic ice sheet. Rocks near to each other tend to be similar, having come from the same source. Considering these statistical dependencies by learning the distribution of rock types as the robot explores the area and takes measurements boosts rock classification accuracy beyond what would be achievable if each rock sample is analyzed independently.*

For many tasks involving the robotic searching and exploration of an area to find and identify objects, it is necessary to characterize the operating environment. To classify objects using onboard sensors, the likely candidates and potential false targets must be known in advance. In addition, to optimally classify objects their prior probabilities of being encountered must be known.

Unfortunately, it is difficult to know *a priori* the relative chances of finding different objects in an unexplored area. Consider the problem of classifying rocks from a robot for the purpose of geological exploration. There are many possible rock types, some hard to distinguish from each other. A geological map, if available, only indicates the most common rock type over a very large area. It does not indicate all rock types or their relative probabilities, and ignores small-scale variations. The latter is very important, as the kinds of rocks present can change significantly over a short distance, such as when crossing geological strata.

The standard approach to classifying objects is to consider each one independently, and classify it based on observations. However, objects in an area may be correlated, exploiting this can significantly increase classification accuracy. It is common for objects of the same type to be clustered together.

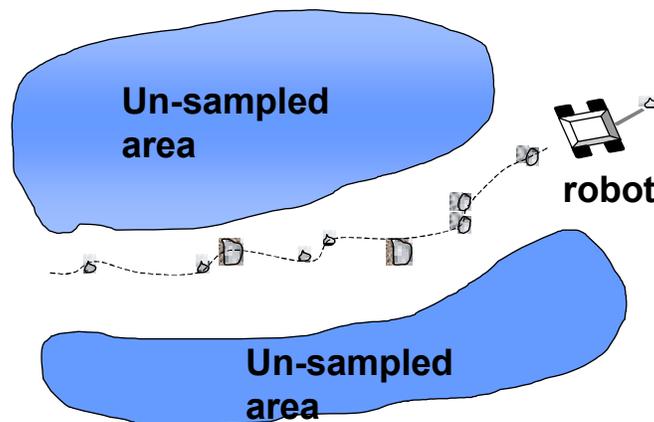
This chapter shows a Bayesian approach to using the dependencies between objects distributed over an area by learning the statistical prior probabilities for different objects as a function of position, and specifically from the perspective of exploring an area with a mobile robot. Classification is thus improved not only by exploiting spatial dependencies, but also through improved knowledge of the priors. Furthermore, a map of prior probabilities over a geographic area is itself a useful summary with which to characterize the region and recognize gross properties.

This work is applicable to a variety of tasks involving the classification of objects distributed across an area, such as geological exploration, landmine removal, or soil profiling for agriculture. However, this analysis focuses on the problem of identifying the abundances of different rock types in an area of the Antarctic ice sheets using a scientific robotic vehicle.

## 9.2. Robotic learning of the environment

Using a robotic vehicle to explore an area introduces unique issues.

- Sampling is not random. The robot examines rocks as it traverses a path through the terrain of interest, leaving areas unexplored (Figure 9.2). The uncertainty about what is found in these areas must be noted, and used to constrain subsequent changes in beliefs about the area as new data is added.



**Figure 9.2** *Modus operandi of a robotic explorer. Only selected samples are visited in a large area, and they are not randomly distributed, but lie along the robots path.*

- Relatively few rock samples are examined. The Nomad robot (Figure 9.1) obtained measurements of no more than 50 rocks in 2 days. To make matters worse, there are many possible rock classes [9.3]. The rock probabilities must therefore be initially coarsely defined and subsequently improved if and when data becomes available.
- Rock samples cannot usually be identified with complete certainty. Rather, when sensor data is obtained from a rock sample, only the likelihood's of different rock classes generating that data is known.

- The probabilities of different rock types are conditioned on geographical position, a continuous 2D (or 3D) quantity.

While machine learning and statistical estimation are mature fields, little prior work directly addresses the problem of characterizing a geographical area for the purposes of classification. The evidence grids of [9.4][9.5] are used to model the likelihood's of obstacles in an area. However, they fail to account for statistical dependencies between objects and require that space be discretized into a grid.

[9.6] and [9.7] survey strategies for either autonomously searching an area, given a prior description of the geographical distribution of targets, or exploiting the knowledge that targets tend to cluster together. They do not address how this information is obtained in the first place.

### 9.2.1. Representing the rock priors over an area

Consider the parameter  $\bar{\theta}_x = [\theta_{x1}, \dots, \theta_{xM}]$ , representing the relative proportions of each rock type present at a geographic location  $x$ , and the random variable  $\mathbf{R}_x$  the rock type (class labels 1, 2... M) of a rock sample found at there. Therefore

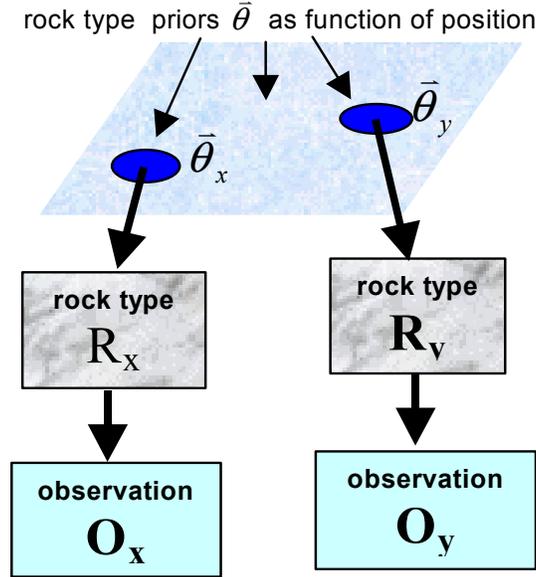
$$(Eqn 9-1) \quad \theta_{x1}, \dots, \theta_{xM} \in [0,1], \quad \sum_{i=1}^M \theta_{xi} = 1$$

$$(Eqn 9-2) \quad P(\mathbf{R}_x = k | \bar{\theta}_x) = P(\mathbf{R}_x = k | \theta_{xk}) = \theta_{xk}$$

$$(Eqn 9-3) \quad P(\mathbf{R}_x = k) = E\{\theta_{xk}\}$$

$\bar{\theta}_x$  is itself a random variable and depends on position. Furthermore, knowing its distribution allows the determination of the rock type priors at  $x$  (3). Therefore, the problem of learning these priors is solved by learning the distribution of  $\bar{\theta}_y$  at all positions  $y$ , given a sequence of robot sensor observations  $\{\mathbf{O}_{x_i} | i=1,2,\dots,N\}$  made at locations  $\{x_i | i = 1,2,\dots,N\}$ . To do this it is necessary to model the statistical relationships between  $\bar{\theta}_y$ , the rock types of each rock sample, and the associated observations. Consequently we will show how to compute  $p(\bar{\theta}_y | \mathbf{O}_{x_i})$ , the posterior representation of the rock priors over the area, and how to use it to improve classification.

## 9.2.2. Pseudo Bayes network generative model



**Figure 9.3** Pseudo Bayes network generative model of the statistical relationships between position  $(x,y)$ , relative proportions of different rock types at those positions  $(\bar{\theta}_x, \bar{\theta}_y)$ , type of rock samples found  $(R_x, R_y)$ , and the observations on those samples  $(O_x, O_y)$ .

It is reasonable to assume that the rock type of a sample at any given location  $x$  is conditionally dependent only on the local rock ratios  $\bar{\theta}_x$ , and sensor observations of a rock depend only upon its type (Figure 9.3).

Therefore, for any positions  $x,y$ , and rock type  $k \in \{1,2,\dots,M\}$

$$(Eqn 9-4) \quad P(R_x = k | \bar{\theta}_x, \bar{\theta}_y, R_y) = P(R_x = k | \bar{\theta}_x) = \theta_{xk}$$

No assumptions are yet made on the relationship between  $\bar{\theta}_x$  and  $\bar{\theta}_y$ , which is not usefully expressible in terms of a standard Bayes network diagram.

Furthermore, (Eqn 9-4) implies that

$$(Eqn 9-5) \quad \begin{aligned} P(R_x = k | \bar{\theta}_y) &= \iint P(R_x = k | \bar{\theta}_x, \bar{\theta}_y) p(\bar{\theta}_x | \bar{\theta}_y) d\bar{\theta}_x \\ &= \int \theta_{xk} \left\{ \int p(\bar{\theta}_x | \bar{\theta}_y) d\theta_{x1} \dots d\theta_{xk-1} \dots d\theta_{xk+1} \dots d\theta_{xM} \right\} d\theta_{xk} \\ &= \int \theta_{xk} p(\theta_k | \bar{\theta}_y) d\theta_{xk} \\ &= E\{\theta_{xk} | \bar{\theta}_y\} \end{aligned}$$

a result that will be used subsequently. See [9.8] for more on the Bayes network representation of statistical relationships.

## 9.2.3. Geographical models

Consider the case when rocks can be identified with complete certainty. Suppose a rock at position  $x$  is determined to be of type  $k$ . Then, using Bayes rule, the posterior density of  $\bar{\theta}_y$  is given by

$$(Eqn 9-6) \quad p(\bar{\theta}_y | R_x = k) = \frac{P(R_x = k | \bar{\theta}_y) p(\bar{\theta}_y)}{P(R_x = k)}$$

$$= \frac{\mathbf{E}\{\theta_{xk} | \bar{\theta}_y\} p(\bar{\theta}_y)}{\mathbf{E}\{\theta_{xk}\}}$$

To clarify the functional relationships it is convenient to define

$$(Eqn 9-7) \quad M_k(\bar{\theta}_y; x, y) := \begin{cases} \theta_{xk} & , x = y \\ \mathbf{E}\{\theta_{xk} | \bar{\theta}_y\} & , x \neq y \end{cases}$$

Then, the posterior density  $\bar{\theta}_y$  given the definitive observation is simply expressed in terms of  $M_k(\bar{\theta}_y; x, y)$

$$(Eqn 9-8) \quad p(\bar{\theta}_y | R_x = k) = \frac{M_k(\bar{\theta}_y; x, y) p(\bar{\theta}_y)}{\mathbf{E}\{\theta_{xk}\}}$$

The function  $M_k(\bar{\theta}_y; x, y)$  is fundamental to determine how a rock find at the location  $x$  affects the rock ratio's at all other locations  $y$  (including  $x$ ). In recognition of the functions importance it will be henceforth referred to as a *geographical model*, as it describes the statistical relationships between samples at different geographic positions.

#### 9.2.4. Properties of $M_k(\bar{\theta}_y; x, y)$

The geographical model  $M_k(\bar{\theta}_y; x, y)$  is a complicated multi-variate function, determined by the underlying geological (or other) mechanism by which different rock types are distributed across an area. Nonetheless, applying reasonable assumptions it is possible to constrain it and make it tractable without going into the details of the (unknown) underlying mechanism.

- *Small influence at a distance*

Observations at a distant location should have little or no relation to the rock ratios at the current location. This assumption can be formalized as

$$(Eqn 9-9) \quad M_k(\bar{\theta}_y; x, y) \rightarrow \mathbf{E}\{\theta_{xk}\} \text{ as } |x - y| \rightarrow \infty$$

Following this it also reasonable to assume some finite cut-off distance, beyond which measurements have no effect:

$$(Eqn 9-10) \quad \text{For some distance } D > 0 \quad |x - y| > D \Rightarrow M_k(\bar{\theta}_y; x, y) = \mathbf{E}\{\theta_{xk}\}.$$

- *Decreasing influence with distance*

Not only should observations at locations distant to each other be largely independent, but the dependence between observations should never increase when the distance between them is increased. While motivated by the previous statement, this is a stronger assumption. Formally:

$$\text{For any locations } x, y, z \text{ such that if } |x-y| > |x-z| \text{ then } |M_k(\bar{\theta}_y; x, y) - \mathbf{E}\{\theta_{xk}\}| \leq |M_k(\bar{\theta}_z; x, z) - \mathbf{E}\{\theta_{xk}\}|$$

- *Smoothness assumption*

$M_k(\bar{\theta}_y; x, y)$  should vary smoothly with  $x$  and  $y$ . This is consistent with natural laws and not very restrictive. However, relaxing this slightly to allow a discontinuity at the finite cut-off distance may be convenient.

- *Spatial invariance and isotropy*

This assumption requires that  $M_k(\bar{\theta}_y; x, y)$  be solely a function of the distance  $|x-y|$  between samples, and not depend on their individual positions or the direction from one to the other. However, this is inconsistent with the previous assumptions, which require explicit dependence on  $x$ . Nonetheless, if this dependence can be explicitly accounted for and functionally separated from the rest of  $M_k(\bar{\theta}_y; x, y)$ , then spatial invariance and isotropy are desirable and reasonable properties, provided that the geographic area of interest is not too large.

- *Conjugacy requirement*

This is the most restrictive and useful of the assumptions, requiring that for all positions  $x, y$ , and rock types  $k$ , the prior distribution of rock ratios  $p(\bar{\theta}_y)$  is the same class of distribution as the posterior  $p(\bar{\theta}_y | R_x = k)$ . Conjugate prior distributions are computationally convenient and the prior is easily interpretable as earlier measurements.

It is natural to represent  $p(\bar{\theta}_y)$  by a Dirichlet distribution [9.7] with parameters  $\alpha_{y1} \dots \alpha_{yM}$  :

$$(Eqn 9-11) \quad p(\bar{\theta}_y) = \text{Dirichlet}(\bar{\theta}_y; \alpha_{y1} \dots \alpha_{yM}) \\ = \frac{\Gamma(\alpha_{y1} + \dots + \alpha_{yM})}{\Gamma(\alpha_{y1}) \dots \Gamma(\alpha_{yM})} \theta_{y1}^{\alpha_{y1}-1} \dots \theta_{yM}^{\alpha_{yM}-1}$$

Ensuring that  $p(\bar{\theta}_y | R_x = k)$  is also a Dirichlet distribution requires that

$$(Eqn 9-12) \quad M_k(\bar{\theta}_y; x, y) = Z_k(x, y) \theta_{y1}^{\beta_1(x,y,k)} \dots \theta_{yM}^{\beta_M(x,y,k)}$$

which guarantees from (Eqn 9-8) that

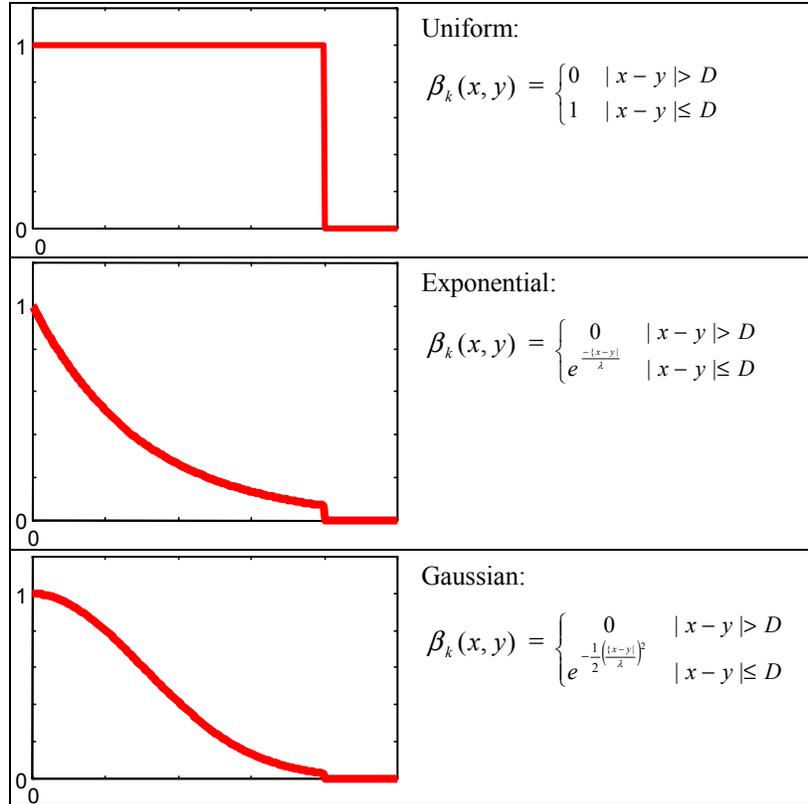
$$(Eqn 9-13) \quad p(\bar{\theta}_y | R_x = k) = \text{Dirichlet}(\bar{\theta}_y; \alpha_{y1} + \beta_1(x,y,k), \dots, \alpha_{yM} + \beta_M(x,y,k))$$

The small influence at a distance assumption (Eqn 9-9) implies that

$$(Eqn 9-14) \quad Z_k(x, y) \xrightarrow{|x-y| \rightarrow \infty} E\{\theta_{xk}\} \\ \beta_i(x, y, k) \xrightarrow{|x-y| \rightarrow \infty} 0$$

Introducing the finite cutoff of (Eqn 9-10) implies the limits are attained when  $x$  and  $y$  are separated by a finite distance. The assumption of decreasing influence with distance (Eqn 1-1) means that  $Z_k$  and  $\beta_1 \dots \beta_M$  never get further from their limits as  $|x-y|$  increases. They are smooth functions as  $M_k(\cdot; x, y)$  is assumed smooth.

Note that computation of the posterior distribution requires only that  $\beta_1(x,y,k) \dots \beta_M(x,y,k)$  be specified.  $Z_k(x,y)$  is implicitly defined by normalization of the posterior in (Eqn 9-8). Furthermore,  $Z_k(x,y)$  accounts for the spatial dependence of  $M_k(\cdot; x, y)$ , making it possible to assume spatial invariance and isotropy for  $\beta_1(x,y,k) \dots \beta_M(x,y,k)$ . That is,  $\beta_j(x, y, k) = \beta_j(|x-y|, k)$ .



**Figure 9.4** Possible formulae for the functions  $\beta_k(x,y)$ . While ad hoc, they satisfy all the restrictions on  $\beta_k(x,y)$ , including spatial invariance and isotropy. The uniform formula weights all samples within an area equally, whilst the others have steadily decreasing influence with distance. Choice of the constants  $D$  and  $\lambda$  involves a trade-off between spatial resolution and power to generalize. Large values allow the rapid learning of the rock ratios over a large area from a few samples but are less effective in learning regional variations, which will be smoothed over. Ideally, they should be comparable to the average separation between samples.

For co-located samples  $x = y$ , hence  $M_k(\bar{\theta}_y; x, y) = \theta_{yk}$ , implying that

(Eqn 9-15)  $\beta_k(x,x,k) = 1$  and  $\beta_j(x,x,k) = 0$  (for every  $j \neq k$ )

Because  $\beta_1(x,y,k) \dots \beta_M(x,y,k)$  approach zero and may not increase as distance  $|x - y|$  increases, it follows that  $\beta_j(x,y,k) = 0$  for all  $x, y$  and  $j \neq k$ . Therefore the notation for the  $\beta$ 's is redundant; it is sufficient to denote the remaining nonzero term  $\beta_k(x,y,k)$  as  $\beta_k(x,y)$ .

This implies that a *definitive* rock find (identified with 100% certainty) will not increase the assumed probabilities of finding other related but distinct rock types in the area.

To summarize

(Eqn 9-16)  $M_k(\bar{\theta}_y; x, y) = Z_k(x, y) \theta_{yk}^{\beta_k(x,y)}$

and

(Eqn 9-17)  $p(\bar{\theta}_y | R_x = k) = \text{Dirichlet}(\bar{\theta}_y; \alpha_{y1} + \beta_k(x,y), \dots, \alpha_{yM})$

To learn the distribution of rock ratios at a particular point given definitive observations in the area, sum the contributions of each observation using (Eqn 9-17) and an appropriate formula for the  $\beta$ 's (Figure 9.4).

### 9.2.5. Learning from uncertain observations

Regrettably, it is rarely the case that rock samples can be autonomously identified with certainty. Otherwise, there would be little reason to learn the priors.

Consider an observation  $O_x$  made on a rock sample  $R_x$  at position  $x$ , with likelihood's  $w_{xk}=P(O_x|R_x=k)$ . It can be shown that [Chapter 4]

$$(Eqn 9-18) \quad p(\bar{\theta}_y | O_x) \propto \sum_k [P_{xk} Z_k(x, y) \theta_{xk}^{\beta_k(x,y)}] p(\bar{\theta}_y)$$

where

$$(Eqn 9-19) \quad \begin{aligned} P_{xk} &= P(R_x = k | O_x) \\ &= \frac{P(O_x | R_x = k) E\{\theta_{xk}\}}{\sum_j P(O_x | R_x = j) E\{\theta_{xj}\}} \quad (\text{Bayes rule}) \\ &= \frac{w_{xk} \alpha_{xk}}{\sum_j w_{xj} \alpha_{xj}} \end{aligned}$$

As  $p(\bar{\theta}_y)$  is a Dirichlet distribution, it follows that

$$(Eqn 9-20) \quad p(\bar{\theta}_y | O_x) = \sum_{k=1}^M P_{xk} \text{Dirichlet}(\bar{\theta}_y; \alpha_{y1}, \dots, \alpha_{yk} + \beta_k(x,y), \dots, \alpha_{yM})$$

This is not a Dirichlet distribution, but a mixture model of Dirichlet distributions. It violates the conjugacy requirement and is intractable as more observations are made. Subsequent observations would produce mixture models with  $M^2$ ,  $M^3$ , and so on terms. To maintain a closed form model with a bounded number of parameters it is necessary to approximate

$$(Eqn 9-21) \quad P(\bar{\theta}_y | O_x) \approx \text{Dirichlet}(\bar{\theta}_y; \alpha_{y1} + P_{x1} \beta_1(x,y), \dots, \alpha_{yM} + P_{x1} \beta_M(x,y))$$

This is equivalent to computing the posterior distribution of  $\bar{\theta}_y$  given the definite observations  $\{R_x=k|k=1\dots M\}$  at  $x$ , each weighted according to their probability  $P(R_x=k|O_x)$  given the current observation and priors on  $\bar{\theta}_y$ . This exactly equals the true posterior whenever the observations identify the rock with complete certainty. The approximation is worst when the observations do not favor any rock type. The latter occurs when an observation is made that is equally likely for all possible rock types, in which case the probability of each rock type is equal to the prior. The approximation (Eqn 9-21) of the posterior distribution on  $\bar{\theta}_y$  then correctly has the same mean as the prior but a slightly decreased variance, which is incorrect, as nothing has been learned.

### 9.2.6. Multiple rock finds

Assuming the prior rock ratio density at any point  $x$  is Dirichlet, with parameters  $\alpha_{xk}^{prior}$ , and an observation is made on a particular sample. Then, the rock type probabilities for that sample are given by (Eqn 9-19), and the rock ratio densities at all other points given by (Eqn 4-17). If observations are made at all sample locations then (Eqn 9-19) and (Eqn 4-17) are combined to create the following coupled simultaneous equations, relating the rock type probabilities at each sample and the posterior rock ratio distribution at each sample given the rock type probabilities at the other samples:

$$(Eqn\ 9-22) \quad P_{xk} = \frac{w_{xk} \alpha_{xk}}{\sum_j w_{xj} \alpha_{xj}}, \quad \alpha_{xk} = \alpha_{xk}^{prior} + \sum_{y \neq x} P_{yk} \beta_k(y, x)$$

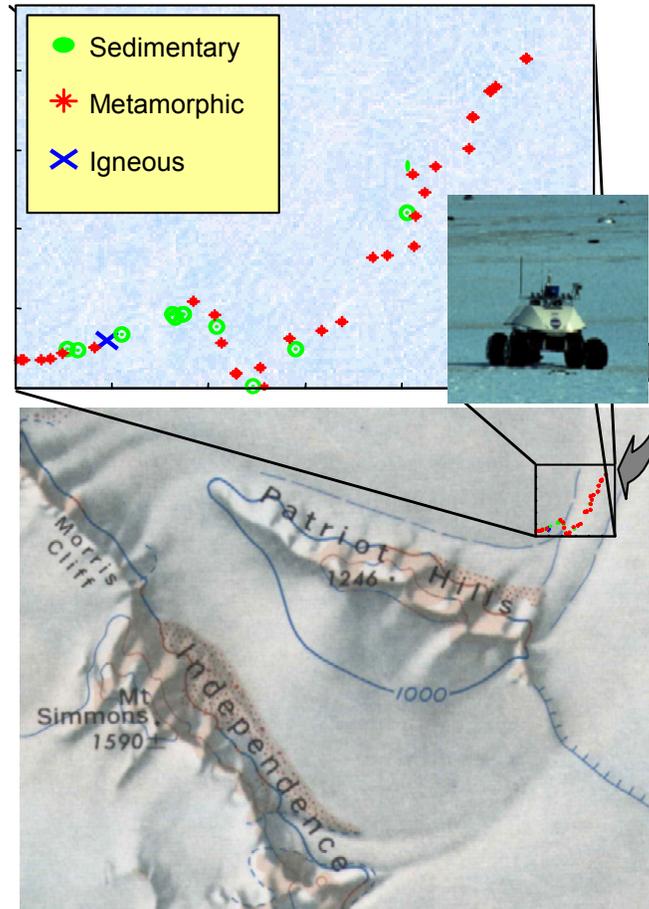
The  $\alpha_{xk}^{prior}$ 's define the assumed (Dirichlet) distributions on  $\bar{\theta}_x$  at all sample locations  $x$  prior to any measurements.

Note that the order in which samples are examined does not affect the computed probabilities and rock ratio distributions.

### 9.3. Experimental results

Results obtained by applying (Eqn 9-22), using Gaussian geographical models (Figure 9.4) to simulated data have shown statistically robust improvements in classification accuracy and generated consistent probability maps. However, there are too many arbitrary parameters to set when simulating data for it necessarily to be a good indicator of performance. There is no substitute for *real* data, gathered in the field by a robot.

In December 1998 the Nomad robot was dispatched to Patriot Hills, Antarctica where it traversed a moraine and acquired spectral data from 51 rocks in its path, along with their spatial coordinates (Figure 9.5). These rock samples were recovered and subsequently classified by a geologist to ascertain ground truth.

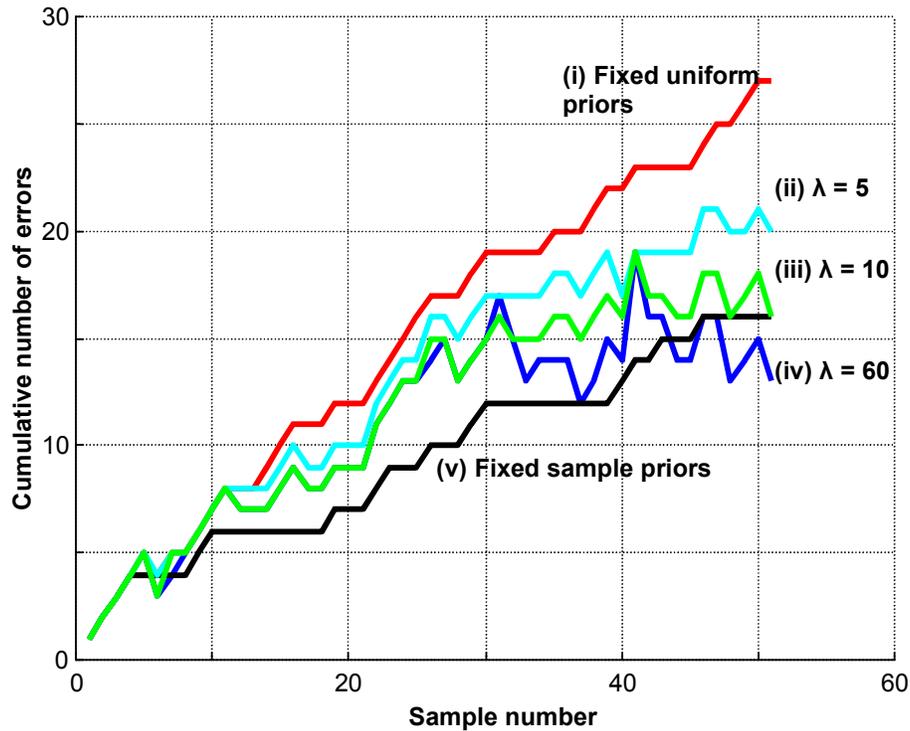


**Figure 9.5** Patriot Hills, Antarctica, data collection site. The Nomad robot traversed a 600m x 300m area, collecting spectroscopic data from each rock along its path. The samples have been grouped into according to their formation process: sedimentary, metamorphic, igneous and extraterrestrial (meteorites, although none are present in the sample), each of which encompasses many non-overlapping rock types.

Each sample spectrum was processed and run through the Bayes network classifier developed for rock and meteorite identification to determine the conditional likelihood of that spectrum for each of 25 different rock types. This data was sequentially entered into a statistical model (Eqn 4-17) of the possible rock ratios at each sample position, using a Gaussian geographical model (Figure 9.4) and assuming initial distributions with  $\alpha_{xk}^{prior} = 0.1$  at all locations. At each iteration the rock type probabilities for every sample entered so far were recomputed and the most likely formation (sedimentary, metamorphic, igneous and extraterrestrial) deduced. Comparing with the known formation processes, the cumulative number of classification errors Figure 9.6(ii)-(iv) as each sample is entered are determined.

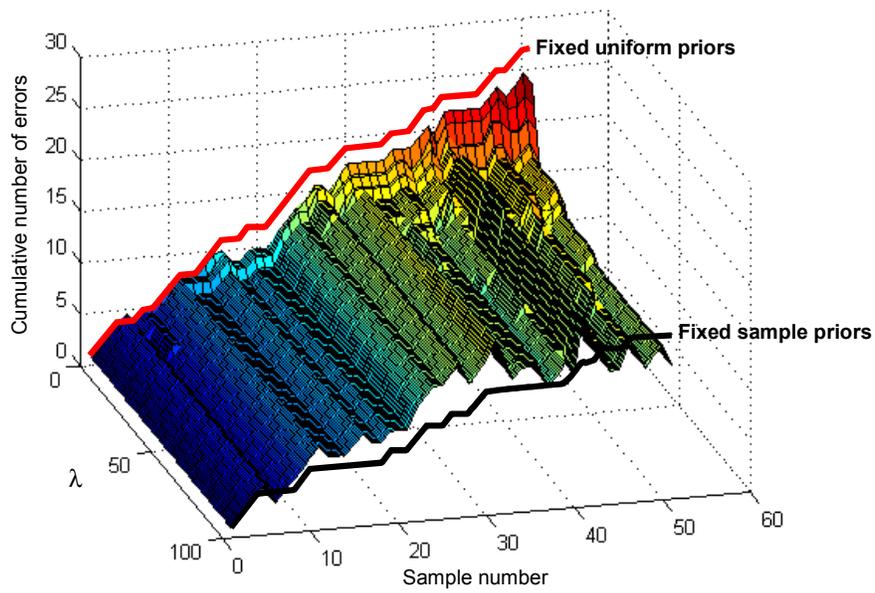
Note the occasional *reduction* in the total number of misclassifications as more samples are examined and the model becomes more precise. This would not be possible if each rock sample was examined independently, as in Figure 9.6(i) and Figure 9.6(v), where the rocks are respectively classified with the assumption of uniform fixed rock type priors everywhere, and the priors fixed to the known data set rock type ratios. These indicate the worst and the best that the rock classifier in [9.3] can do under different assumptions on the rock type priors over the entire region if each rock is classified independently and the

priors are assumed the same at each sample location. Note that performance of the classifier now can exceed the best performance attained by independent classification of objects.

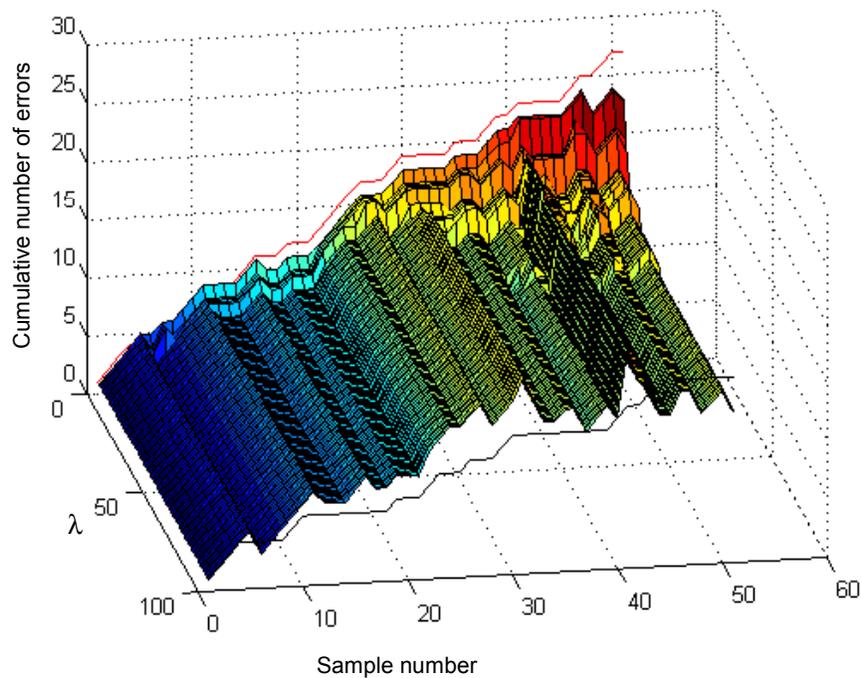


**Figure 9.6** Cumulative number of misclassifications as samples from the Antarctic field data are examined in the order they were encountered by the robot. In (i) samples are each independently classified assuming uniform rock type probabilities. In (ii) – (iv) the rock type probabilities are learned as samples are acquired, using Gaussian  $\beta$ 's with the indicated standard deviations  $\lambda$ ; and in (v) rocks are again classified independently, using the known fraction of each rock type in the data set as the priors.

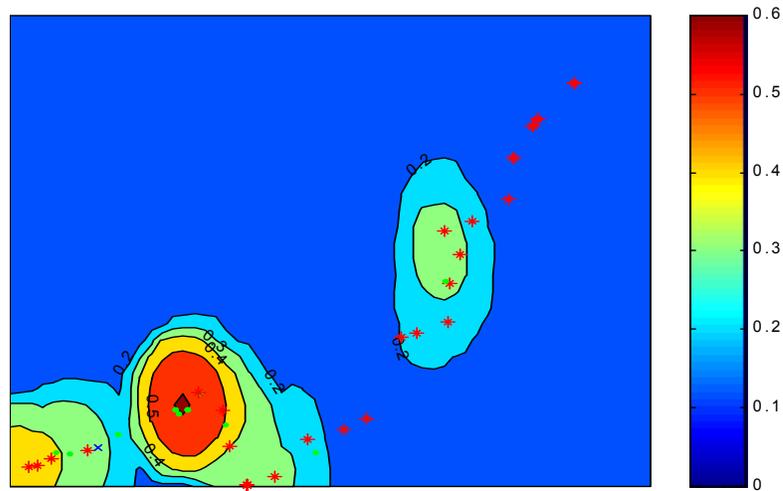
Empirical evidence suggests that the reduction in error rate is robust to changes in the precise form of the geographical model functions (Figure 9.7 and Figure 9.8).



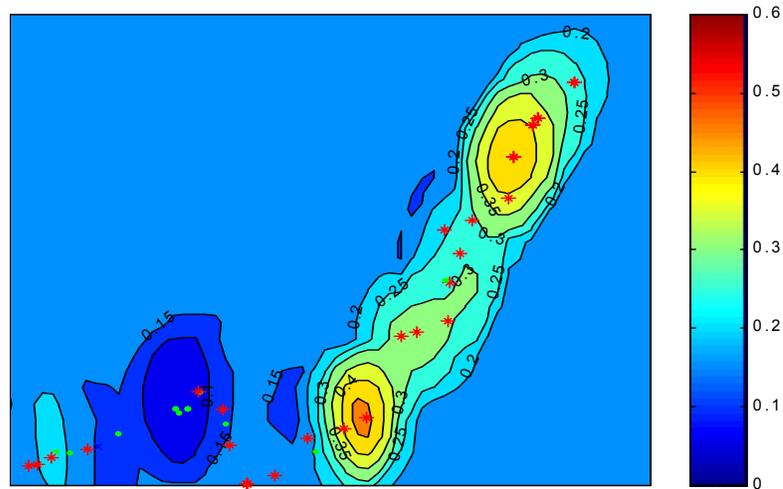
**Figure 9.7** Cumulative error functions using Gaussian  $\beta$ 's as the standard deviations  $\lambda$  vary. Note steady decrease in number of errors as  $\lambda$  increases. Increasing  $\lambda$  beyond 60 does not further affect error rates. Furthermore,  $\lambda = 1$  gives almost identical performance as independent classification using fixed uniform priors, as expected.



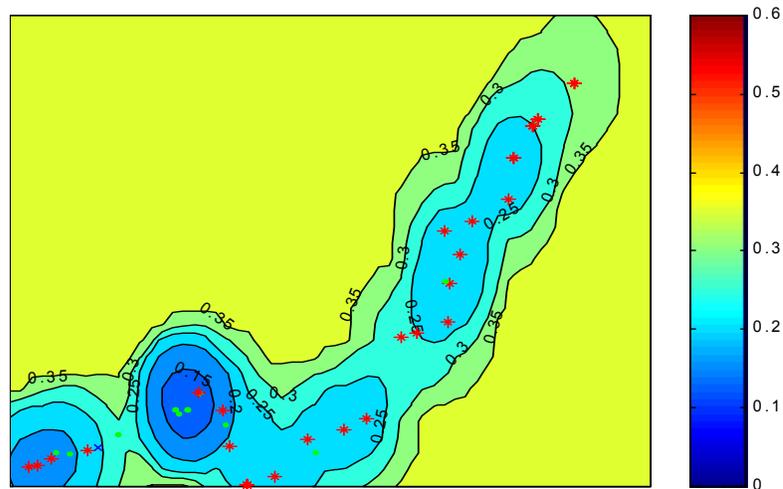
**Figure 9.8** Effect of using different forms of geographical models, this time modeling the  $\beta$  functions as uniform with a finite cut-off (Figure 9.4). Notice the similarity to the error functions generated using Gaussian  $\beta$  functions (Figure 9.7)



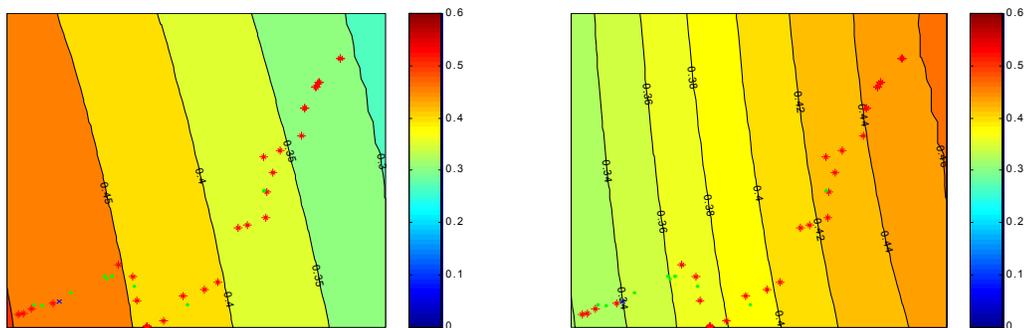
**Figure 9.9** *Learned probability of sedimentary rocks across the explored region of the data collection site in Antarctica. The original sample positions and types are indicated by the colored dots (c.f. Figure 9.5).*



**Figure 9.10** *Learned probability of metamorphic rocks.*



**Figure 9.11** *Learned probabilities of igneous rocks.*



**Figure 9.12** *Sedimentary and Metamorphic rock distribution maps generated using a very large kernel (Gaussian,  $\lambda = 60$ ).*

With all the rock data entered into the model it is also possible to compute the learned rock type probabilities everywhere in an area (Figure 9.9, Figure 9.10, and Figure 9.11 for medium sized kernel functions; and Figure 9.12 for the limit as the kernel functions become very wide), not just at the samples. Compare these to the actual distribution of the rock samples along the robots path in Figure 9.5 to verify that they are indeed consistent with what was seen by the robot. The regions dominated by sedimentary and metamorphic rocks are clearly identified. Conversely, nothing is learned about the areas distant from the robot path, since a geographic model with a finite cut-off is assumed. Furthermore, the low density of igneous rocks (and meteorites, which have a probability map almost identical to Figure 9.11) is learned, even in the area dominated by metamorphic rocks, which are often confused with igneous rocks.

## 9.4. Conclusions

The results of Figure 9.6 are very significant. They show that learning the probabilities gives a clear improvement in classification over assuming uniform priors everywhere and classifying each sample independently. In fact, performance approaches and possibly exceeds the optimum, achieved when the

average priors over the region are known beforehand. The latter might occur with more data and pronounced regional variations that can be exploited by this approach.

This is equivalent to a human geologist who looks at many rocks, constantly re-evaluating all the previous rocks seen every time another rock is looked at. Unfortunately, because of the lack of prior work in this exact area it is not possible to compare this with other methods.

A weakness of the approach proposed here is the ad hoc nature of the geographical models. While based upon reasonable assumptions they are still under-constrained. Choosing models with a wide footprint (large cut-off distance and low decay rate) results in faster convergence and generalizes over a larger area, but also less ability to capture and exploit small scale variations. Further work is needed to determine the optimal geographic models from the data.

All rock type probabilities are re-computed every time another rock sample is examined, and do not depend on the order in which they are found. Therefore, this method is robust to unlucky sequences of samples not representative of the area. Except for the approximation (Eqn 4-17), the learning algorithm is Bayesian, and should converge (in the probabilistic sense) to the correct probabilities as more data is added.

Computationally requirements are minimal. They increase with the number of samples squared. In practice, the matrices in (Eqn 9-22) are sparse and, depending on the robot path, complexity is order  $N^2$  each time a new sample is added, where  $N$  is the average number of samples within a circle whose radius is the geographical model cut-off distance.

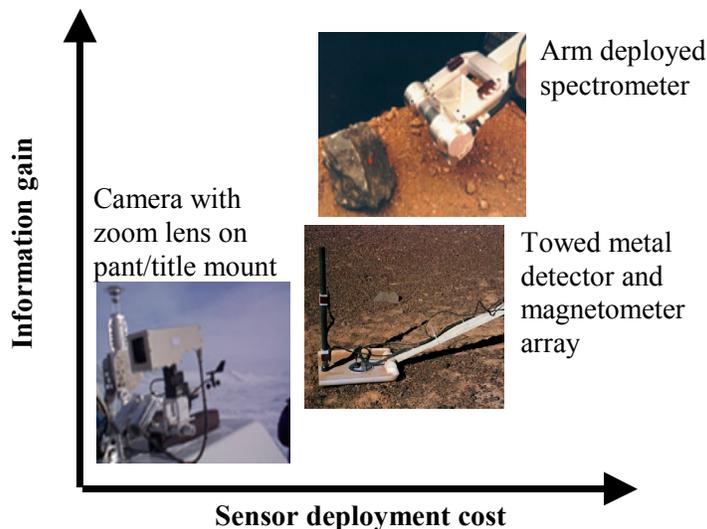
## 9.5. References

- [9.1] Apostolopoulos, D., Wagner, M., Shamah, B., Pedersen, L., K. Shillcutt, Whittaker, W., "Technology and Field Demonstration of Robotic Search for Antarctic Meteorites", *International Journal of Robotic Research*, special Field & Service Robotics issue, December 2000.
- [9.2] Pedersen, L., Apostolopoulos, A., Whittaker, W., Roush, T., and Benedix, G., "Sensing and data classification for robotic meteorite search" *Proceedings of SPIE Photonics East Conference*. Boston, USA, 1998.
- [9.3] Pedersen, L., Apostolopoulos, D., Whittaker, W., "Bayes Networks on Ice: Robotic Search for Antarctic Meteorites", *Neural Information Processing Symposium 2000*, Denver, Colorado, 2000.
- [9.4] Elfes, A., "Occupancy grids: a stochastic spatial representation for active robot perception", *Autonomous Mobile Robots: Perception, Mapping and Navigation*, S. Iyengar, A. Elfes, Eds. IEEE Comp. Soc. Press, 1991.
- [9.5] Elfes, A., "Robot Navigation: Integrating Perception, Environmental Constraints and Task Execution Within a Probabilistic Framework", *Reasoning with Uncertainty in Robotics*, L. Dorst, M. van Lambalgen and F. Voorbraak (Eds), Amsterdam, 1995.
- [9.6] Gelenbe, E., Schmajuk, N., Staddon, J., Reif, J., "Autonomous search by robots and animals: a survey", *Robotics and Autonomous Systems* 22, 1997, pp 23-34.
- [9.7] Gelenbe, E., Cao, Y., "Autonomous search for mines", *European Journal of Operational Research* 108, 1998, pp319-333.

- [9.8] Gelman, A., Carlin, J.B., Stern, H.S. and Rubin, D.B., *Bayesian Data Analysis*, Chapman & Hall, 1995.
- [9.9] Pearl, J., *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*, Morgan Kaufman, 1988.

## Active Sensor Selection

Multiple sensors are needed for rock classification. However, for any particular sample not all available sensors may be needed, or even useful. For example, an Antarctic rock that produces a strong signal in metal detector is very likely a meteorite. Obtaining a spectral reading will probably not gain much further information as to whether the rock is a meteorite or not. However, if no metal detector signal was noted, the rock could still be a meteorite and a spectral reading is called for.



**Figure 10.1** Cost – benefit trade-off when using different sensors. Different sensors have different benefits and costs associated with their deployment. A color camera is easy to deploy, requiring only a few moments to be pointed and focused on a target. An arm-mounted spectrometer requires the robot to maneuver until a rock target is within the arm workspace, after which the robot must halt while the arm is deployed and measurements taken. On the other hand, a camera image is a lot less effective than a spectral reading for investigating a rock sample to see if it is a meteorite. To rationally decide whether to use a sensor on a sample it is necessary to quantify its effectiveness and decide if the sample is worth investigating.

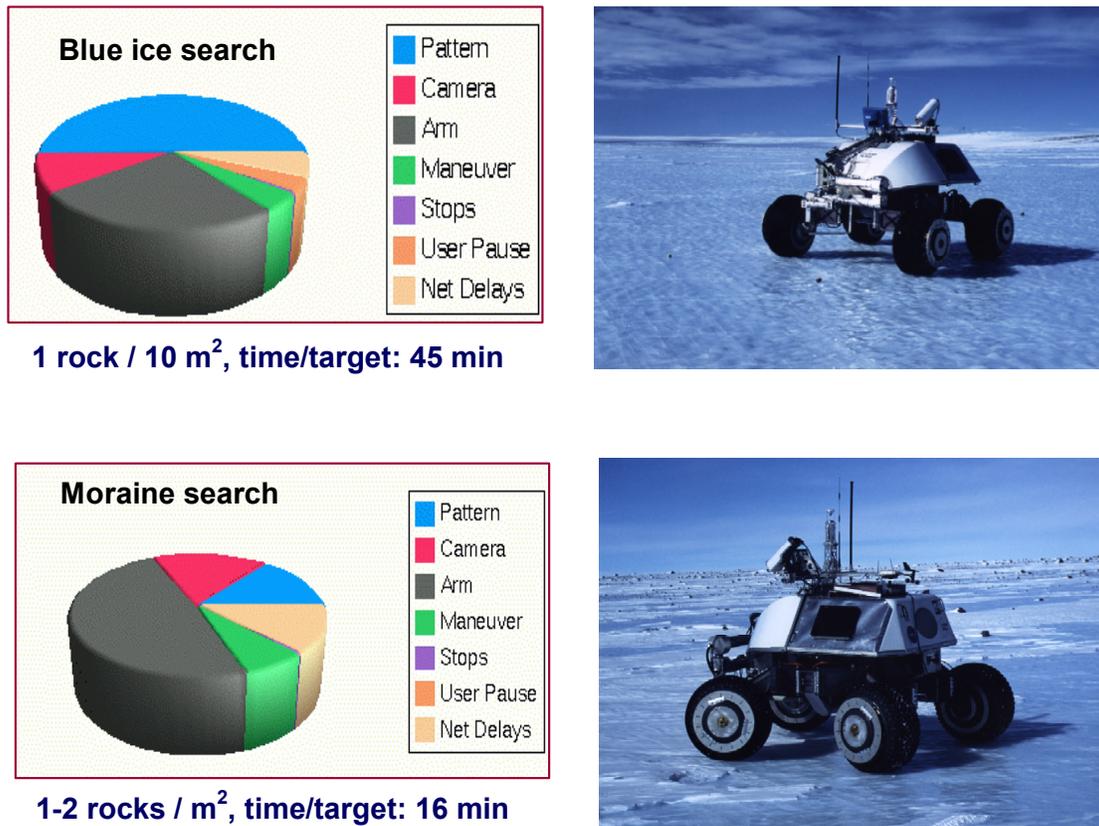
Furthermore, there are costs associated with using sensors (Figure 10.1). Time is required to deploy them, the robot may have to stop or maneuver, energy and computing are required. Quantifying both the costs and benefits of using a sensor allows a robot to better allocate sensing and other resources for an exploration or search task.

The benefits of selective sensor use are great, especially for a robotic vehicle with a variety of both long and short-range sensors. For example, in Antarctica, Nomad spent up to half its time deploying the spectrometer (Figure 10.2), with each spectral reading taking approximately 5-10 minutes to acquire. Using the color camera to identify those rocks clearly not meteorites, reserving the spectrometer for the more difficult cases, could potentially save much time and allow the robot to examine many more rocks.

The perceived benefit of a sensor for identifying a sample depends on what is known about the sample. Assuming appropriate cost models there is extensive literature [10.1][10.2] in the planning and operations research community about cost-benefit analysis and planning to maximize expected profit. Assuming the

perceived cost and benefits of using a sensor can be meaningfully compared, the sensor should be used whenever the cost is less than the benefit. A complication occurs when considering multiple targets and sensors, with benefits and costs that change as measurements are made, but this is beyond the scope of this thesis.

This chapter is concerned mainly with the quantification of sensor utility, and does not delve into cost models or how anything but the simplest cost-benefit analysis would be carried out.



**Figure 10.2** Breakdown of Nomad task times while searching for meteorites in Elephant Moraine, Antarctica. Most time is spent deploying the spectrometer sensor arm to acquire spectra, with some extra time spent maneuvering to bring Nomad into position. This illustrates the value of being able to dispense with spectral readings on those samples sufficiently well identified by the camera.

## 10.1. Information gain as measure of sensor utility

Information theory [10.3] defines the *entropy*  $H(\mathbf{X})$  of a random variable by

$$(Eqn 10-1) \quad H(\mathbf{X}) = -\sum_i P(\mathbf{X} = i) \log_2 P(\mathbf{X} = i)$$

Entropy is a measure of how well known  $\mathbf{X}$  is. If  $\mathbf{X}$  is known with certainty, then its entropy is zero. Otherwise, the entropy of  $\mathbf{X}$  is maximized if all states of  $\mathbf{X}$  are equally likely. Figure 10.3 shows the entropy plot for a binary random variable.

Strictly speaking, entropy is the expected information theoretic information (in bits) obtained from a measurement of a random variable. If the variable is known with certainty, no information is gained from a measurement on it, and therefore entropy is zero.

Consider another random variable  $\mathbf{O}$ , statistically related to  $\mathbf{X}$ . A measurement of  $\mathbf{O}$  will therefore change the distribution of  $\mathbf{X}$ . This change is quantified by the change in the entropy of  $\mathbf{X}$ . The entropy reduction for  $\mathbf{X}$  measures the information value of the measurement  $\mathbf{O}$ .

Subsequent to measuring  $\mathbf{O}$  the entropy of  $\mathbf{X}$  is given by

$$(Eqn 10-2) \quad H(\mathbf{X} | \mathbf{O}) = -\sum_i P(\mathbf{X} = i | \mathbf{O}) \log_2 P(\mathbf{X} = i | \mathbf{O})$$

Therefore the reduction in the entropy of  $\mathbf{X}$  given  $\mathbf{O}$  is

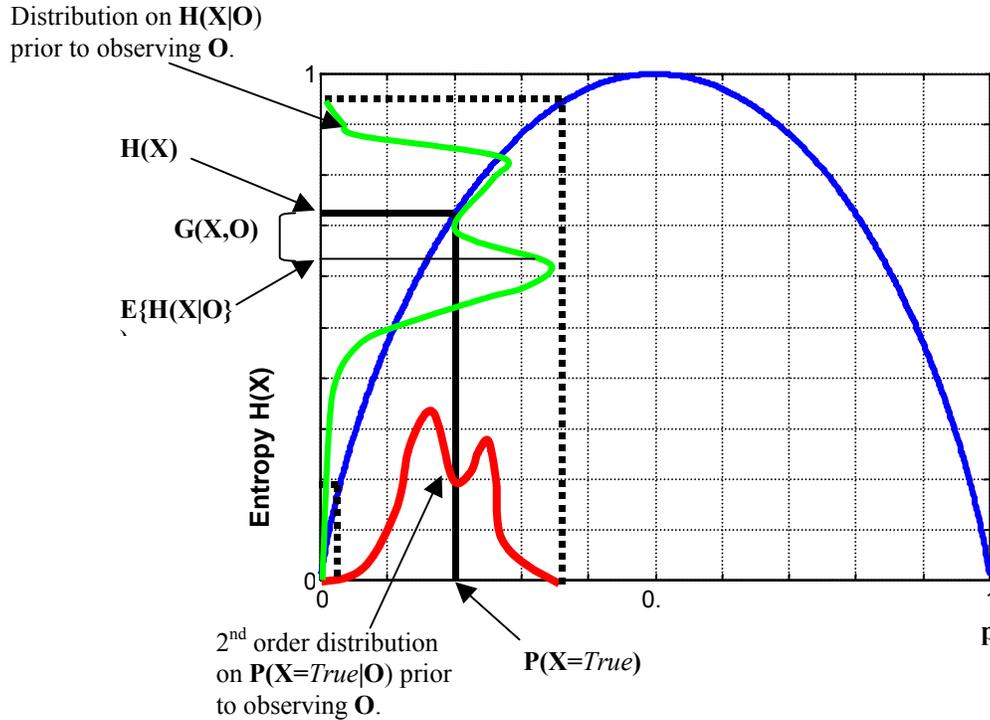
$$(Eqn 10-3) \quad -\Delta H(\mathbf{X}, \mathbf{O}) = H(\mathbf{X}) - H(\mathbf{X} | \mathbf{O})$$

*Information Gain*  $G(\mathbf{X}, \mathbf{O})$  is defined as the expected reduction in the entropy of  $\mathbf{X}$  if  $\mathbf{O}$  were to be measured.

$$(Eqn 10-4) \quad G(\mathbf{X}, \mathbf{O}) \equiv E_{\mathbf{O}}\{-\Delta H(\mathbf{X}, \mathbf{O})\} \\ = H(\mathbf{X}) - \sum_j P(\mathbf{O} = j) H(\mathbf{X} | \mathbf{O} = j)$$

If  $\mathbf{O}$  is continuous

$$(Eqn 10-5) \quad G(\mathbf{X}, \mathbf{O}) = H(\mathbf{X}) - \int p(\mathbf{O}) H(\mathbf{X} | \mathbf{O}) d\mathbf{O}$$



**Figure 10.3** Entropy and information gain on a binary random variable  $\mathbf{X}$  from a measurement of another random variable  $\mathbf{O}$ . Knowing the posterior entropy  $H(\mathbf{X}|\mathbf{O})$  for all possible values of  $\mathbf{O}$  allows the computation of the distribution of  $H(\mathbf{X}|\mathbf{O})$  from the distribution of  $\mathbf{O}$ , and thence the expected value of  $H(\mathbf{X}|\mathbf{O})$  once  $\mathbf{O}$  has been measured. Information gain  $G(\mathbf{X},\mathbf{O})$  is the expected reduction in the entropy of  $\mathbf{X}$  from a measurement of  $\mathbf{O}$ . (Note that 2<sup>nd</sup> order distributions do NOT come from real data).

In the context of robotic meteorite search,  $\mathbf{X}$  is a binary random variable representing whether a rock sample is a meteorite or not and  $\mathbf{O}$  is some sensor observation. When a sensor observation is made the belief as to whether a rock sample is a meteorite changes. The information gain  $G(\mathbf{X},\mathbf{O})$  is therefore a measure of the utility for meteorite identification of taking measurements with the sensor corresponding to  $\mathbf{O}$ .

Information gain is used for this purpose in many applications. In [10.5] entropy and information gain are used as measures of sensor effectiveness for mapping the environment. [10.6] use it in a military battlefield context to decide where to allocate reconnaissance resources. It has been used in medical diagnosis systems to determine what tests are necessary [10.10]. In decision tree classifiers, information gain is the criterion for determining the order of queries along each branch of the decision tree. It is therefore natural to adapt information gain as a measure of sensor utility for robotic exploration.

### 10.1.1. Computing Information Gain for sensor selection

Consider a Bayes network representation of the joint probability distribution of object classifications ( $\mathbf{X}$ ) and sensor observations ( $\mathbf{S}_1, \mathbf{S}_2, \dots, \mathbf{S}_N$ ). The information gain  $G(\mathbf{X}, \mathbf{S}_i)$ , indicating the utility of sensor  $\mathbf{S}_i$  for determining the classification  $\mathbf{X}$  of an object, is computed using (Eqn 10-4)(Eqn 10-2). All the probabilities needed for this computation are computed from the Bayes network classifier.

Sensor utility, and information gain, depend on the knowledge of the world encoded in the joint distribution of  $\mathbf{X}, \mathbf{S}_1, \mathbf{S}_2, \dots, \mathbf{S}_N$ . Making observations with a sensor changes this, and therefore changes the information gain from the remaining sensors. If  $\mathbf{S}_1 \dots \mathbf{S}_{i-1}$  have already been observed, then the information gain from  $\mathbf{S}_i$  is given by

$$(Eqn 10-6) \quad G(\mathbf{X}|\mathbf{S}_1, \dots, \mathbf{S}_{i-1}, \mathbf{S}_i) = H(\mathbf{X}|\mathbf{S}_1 \dots \mathbf{S}_{i-1}) - \sum_k P(\mathbf{S}_i = k) H(\mathbf{X}|\mathbf{S}_1 \dots \mathbf{S}_{i-1}, \mathbf{S}_i = k)$$

This indicates a big advantage of Bayes network based classifiers, for which computing the conditional probabilities above is trivial. No assumptions whatsoever need be made regarding statistical independence of the sensors. Doing this with non-Bayes network generative models is difficult. With the exception of decision tree classifiers, this analysis is meaningless to classifiers not based on Bayesian statistics.

In the context of robotic meteorite search, (Eqn 10-6) is a measure of the utility of obtaining a spectral reading from a sample *subsequent* to getting an image it and computing the likelihood it is a meteorite based on the image. If it exceeds some threshold, then the robot would be directed to acquire a spectrum from the rock. In this way, spectrometer utilization can be restricted to only those rocks requiring it.

### 10.1.2. Monte-Carlo computation of information gain

A single sensor measurement can corresponds to observing multiple random variables, single ones as it may appear above. For example, a spectral measurement may give a 15-element feature vector. The correct information gain for measuring the utility of a spectrometer is the information gain from jointly observing all 15 elements of the spectral feature vector, not each one individually.

Let  $\mathbf{V}_1 \dots \mathbf{V}_M$  represent the feature vector for a sensor  $\mathbf{S}$ . Then, the information gain for  $\mathbf{S}$  is given by summing over all states of the joint state space of  $\mathbf{V}_1 \dots \mathbf{V}_M$ .

$$(Eqn 10-7) \quad \begin{aligned} G(\mathbf{X}, \mathbf{S}) &= G(\mathbf{X}, \mathbf{V}_1 \dots \mathbf{V}_M) \\ &= H(\mathbf{X}) - \sum_{(i_1, \dots, i_M) \in \text{Joint state space}} P(\mathbf{V}_1 = i_1, \dots, \mathbf{V}_M = i_M) H(\mathbf{X}|\mathbf{V}_1 = i_1, \dots, \mathbf{V}_M = i_M) \end{aligned}$$

If a sensor observation  $\mathbf{S}$  has a “small” number of states, then equation (Eqn 10-7) is relatively straightforward. However, the joint state space of  $\mathbf{V}_1 \dots \mathbf{V}_M$  grows exponentially with  $M$ . For Nomad’s classifier, spectra are represented by 15-element feature vectors, each dimension quantized into 10 intervals, giving a spectral observation state space size of  $10^{15}$ . It would take a 100 MIPS computer approximately 100 days to evaluate (Eqn 10-7) in this case. This allows one measurement per Antarctic field season, which could displease the sponsors.

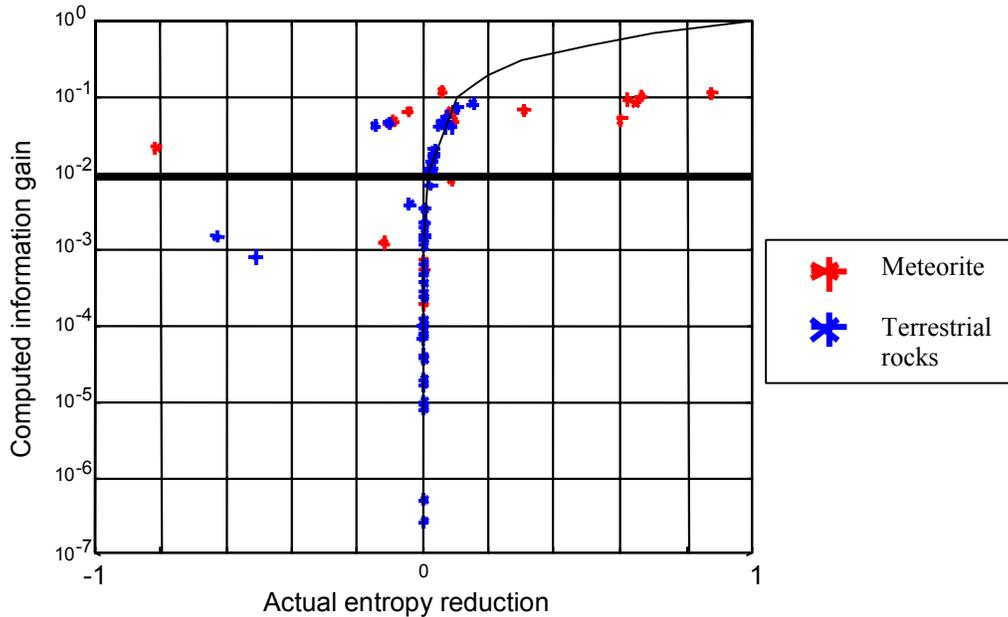
Fortunately, it is unnecessary to compute the exact information gain, which is itself only an *estimate* of sensor utility. Monte-Carlo estimation of information gain is feasible. Let  $\{s_1, \dots, s_N\}$  be independent identically distributed (iid) samples from the full distribution on the joint state space of  $\mathbf{S}$ . Then

$$(Eqn 10-8) \quad \hat{G}(\mathbf{X}, \mathbf{S}) = H(\mathbf{X}) - \frac{1}{N} \sum_{i=1}^N H(\mathbf{X}|s_i)$$

Assuming that  $\{s_1, \dots, s_N\}$  are unbiased, then (Eqn 10-8) is both an unbiased and consistent estimate of the information gain on  $\mathbf{X}$  from  $\mathbf{S}$ .

A Bayes network representation of the joint distribution of  $\mathbf{X}$  and the components of  $\mathbf{S}$  can be efficiently sampled in an unbiased manner.

### 10.1.3. Experimental results

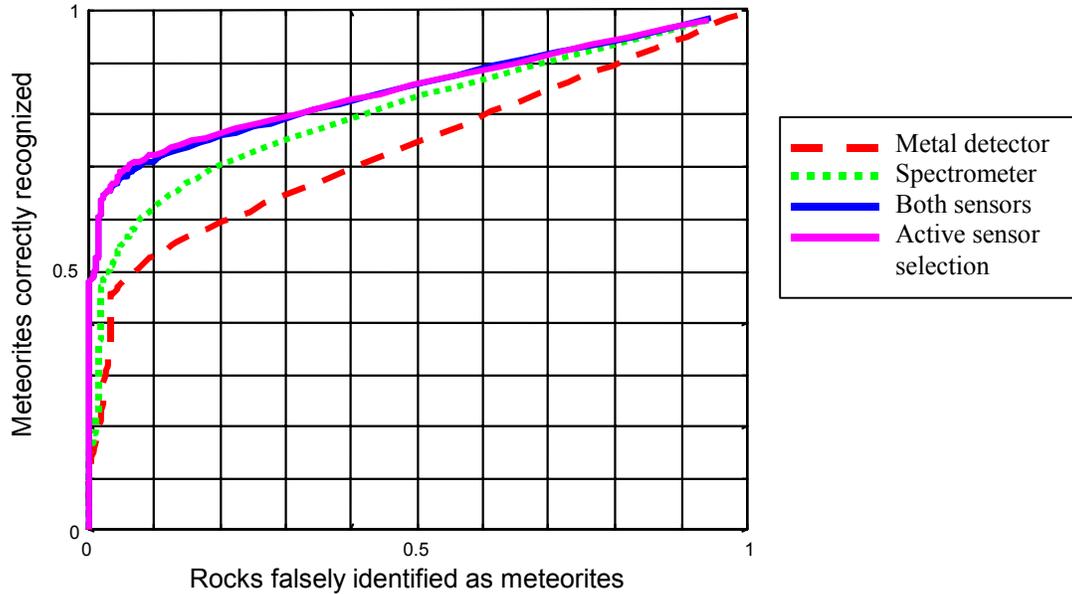


**Figure 10.4** Plot of information gain for meteorite identification given a metal detector reading versus the actual entropy reduction. This plot was generated by first measuring reflectance spectra from a set of Antarctic rocks and meteorites and computing the posterior likelihood of each sample being a meteorite. Following this, metal detector measurements were made and a new posterior computed. The actual entropy reduction is the difference between the entropies of these two posterior distributions, whereas the information gain is the estimated reduction in entropy if a metal detector reading were to be obtained. If the rocks had not been first analyzed with a spectrometer, then the same prior knowledge would be assumed for each sample, giving each the same information gain.

Note logarithmic-linear scale. The black line is a plot of  $y = x$ . Most terrestrial rock samples have information gains near zero. These have been identified as terrestrial with near certainty by the Bayes network classifier from the spectral observation; therefore only a small benefit is anticipated from a metal detector reading.

Figure 10.4 plots information gain versus actual entropy reduction for a test set of Antarctic rocks and meteorites. A spectral observation was entered into the Nomad Bayes network rock and meteorite classifier. Subsequently, the information gain from using a metal detector to determine if each rock is a meteorite was computed. Then the metal detector data was entered and the actual entropy reduction measured.

In this case, information gain appears to be a reasonable predictor of the actual reduction in entropy for terrestrial rock samples. Setting an information gain cut-off threshold of  $10^{-2}$ , above which samples are classified using both spectral and metal detector data, and below which only spectral data is used, the ROC classification curves in **Figure 10.5** is obtained. Classification performance using active sensor selection is only marginally below that using both sensors every time. However, on terrestrial rocks a 60% reduction in the use of metal detector data is obtained, and a 30% reduction on meteorites.



**Figure 10.5** Meteorite rate of classification curves using metal detector only, spectrometer only, both sensors all the time, and the spectrometer combined with metal detector data some of the time. An initial classification is made given the spectrometer data, and the estimated average entropy reduction given metal detector data completed. If greater than  $10^{-2}$  (Figure 10.4), metal detector readings are used to improve classification. With this threshold, get a 60% reduction in metal detector use on rock samples, and 30% reduction on meteorite samples is obtained. Within the sample size errors, classification accuracy is not affected.

### 10.1.4. Problems with Entropy and Information Gain

While Figure 10.4 and Figure 10.5 support the use of information gain to estimate sensor utility, some problems are apparent.

- Information gain is always positive, yet sometimes entropy increases after new data is obtained. Further data is always expected, on average, to reinforce the current belief. Occasionally contradictory data occurs and therefore increases the entropy.
- The information gain is significantly worse at predicting entropy decrease for one class of objects (meteorites) than for the other (rocks).

The reason for these becomes clear from Figure 10.3. The second order distribution on  $P(\mathbf{X}|\mathbf{O})$  prior to observing  $\mathbf{O}$  is centered about the current belief  $P(\mathbf{X})$ , and includes values of  $\mathbf{O}$  that both reinforce and contradict the current belief. If all states of  $\mathbf{X}$  are equally likely, then  $H(\mathbf{X})$  is maximized and the distribution on  $H(\mathbf{X}|\mathbf{O})$  consists exclusively of values less than or equal to  $H(\mathbf{X})$ . An observation on  $\mathbf{O}$  *must* reduce the entropy on  $\mathbf{X}$ . However, if some state of  $\mathbf{X}$  is more likely than others, then  $H(\mathbf{X})$  is not at a maximum, and the distribution of possible entropy reductions from observing  $\mathbf{O}$  contains positive and negative values, corresponding respectively to observations that reinforce or contradict the current belief.

In the case that  $H(\mathbf{X})$  is not a maximum, the expected entropy reduction in  $\mathbf{X}$  given  $\mathbf{O}$  will underestimate the magnitude of the change in the entropy (Figure 10.3). Furthermore, information gain and entropy

reduction depend strongly on the prior  $P(\mathbf{X})$ . If this is near an extreme (0 or 1) then information gain will be near zero. This explains the clustering near zero for the information gains and entropy reductions of terrestrial rocks in Figure 10.4. The prior probability of a rock being a meteorite is assumed to be 10% in this case, and spectral data from a terrestrial rock is likely to reinforce this belief. However, spectral data from meteorites tends to increase the belief that a sample is a meteorite to higher than the prior of 10%, closer to 50%. Therefore, there is a greater variation in the distribution of entropy reduction after a metal detector reading, explaining the greater spread in the entropy reductions for meteorites from the expected value.

Another problem is in the use of entropy reduction as a measure of how well a sensor did. Consider an observation that changes  $P(\mathbf{X})$  from 10% to 90%. In this case  $H(\mathbf{X})$  is unchanged and entropy reduction is zero, indicating that the observation had no effect. This is plainly not so, as some very important information was obtained.

A final problem is that prior probabilities  $P(\mathbf{X})$  may not be known well and may be assigned somewhat arbitrary values. Changing these then changes the perceived information gains from different sensors.

### 10.1.5. “Re-centered” entropy

When all the states of the random variable  $\mathbf{X}$  are equally likely, the aforementioned problems go away. This is because as the entropy of  $\mathbf{X}$  is maximized by this distribution. This motivates the idea of redefining the concept of information gain as the expected reduction in the entropy of  $\mathbf{X}$  when  $\mathbf{X}$  is initially in the state of maximum entropy (uniform probabilities).

Consider the posterior probability distribution on  $\mathbf{X}$  after an observation  $\mathbf{O}$ . Observe from Bayes rule that

$$(Eqn 10-9) \quad \log_2 \frac{P(\mathbf{X} | \mathbf{O})}{P(-\mathbf{X} | \mathbf{O})} = \log_2 \frac{P(\mathbf{O} | \mathbf{X})}{P(\mathbf{O} | -\mathbf{X})} + \log_2 \frac{P(\mathbf{X})}{P(-\mathbf{X})}$$

Or, if there had been prior observations  $\mathbf{O}^*$

$$(Eqn 10-10) \quad \log_2 \frac{P(\mathbf{X} | \mathbf{O}, \mathbf{O}^*)}{P(-\mathbf{X} | \mathbf{O}, \mathbf{O}^*)} = \log_2 \frac{P(\mathbf{O} | \mathbf{X}, \mathbf{O}^*)}{P(\mathbf{O} | -\mathbf{X}, \mathbf{O}^*)} + \log_2 \frac{P(\mathbf{X} | \mathbf{O}^*)}{P(-\mathbf{X} | \mathbf{O}^*)}$$

If  $P(\mathbf{X})$  (or  $P(\mathbf{X} | \mathbf{O}^*) = [p, \dots, p_N]$  and  $P(\mathbf{X} | \mathbf{O})$  (or  $P(\mathbf{X} | \mathbf{O}, \mathbf{O}^*) = [p^*, \dots, p_N^*]$ , then

$$(Eqn 10-11) \quad \text{logit}(p_k^*) = E + \text{logit}(p_k)$$

where

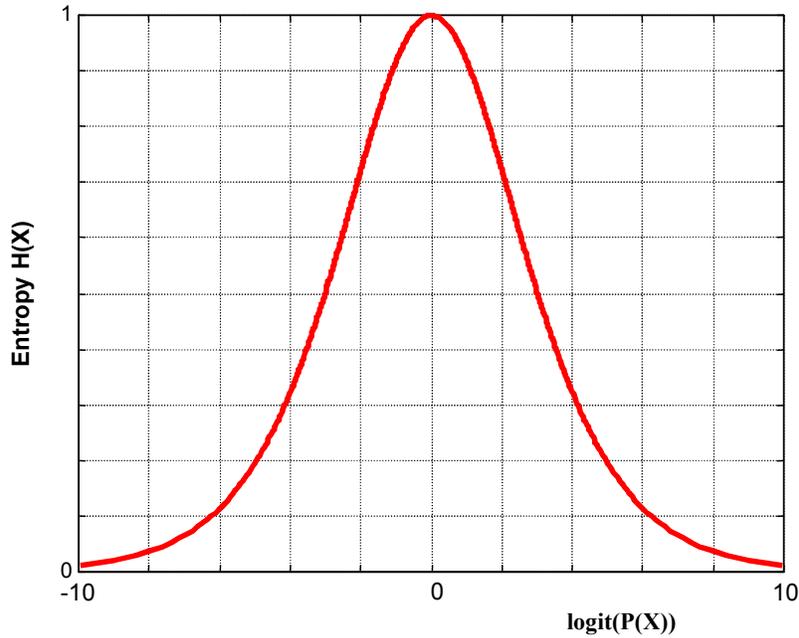
$$(Eqn 10-12) \quad \text{logit}(p) = \log_2 \frac{p}{1-p}$$

and  $E$  is the new *evidence* from the observation  $\mathbf{O}$ . If, prior to measuring  $\mathbf{O}$ ,  $\mathbf{X}$  has uniform probabilities, then

$$(Eqn 10-13) \quad \text{logit}(p_k^*) = E$$

$E$  is thus a measure of the effectiveness of the observation  $\mathbf{O}$  *independent* of the current beliefs (but not necessarily independent of previous observations).

Let  $H^*$  be entropy defined on the logit scale (Figure 10.6).



**Figure 10.6** Entropy plotted on the logit scale,  $\text{logit}(p) = \log_2\left(\frac{p}{1-p}\right)$

Define “re-centered” entropy

$$(Eqn\ 10-14) \quad H^*(\mathbf{X};P) = H^*(\text{logit}(\mathbf{X}) - \text{logit}(P))$$

This is equivalent to translating the entropy function in Figure 10.6 so that the maximum occurs when the random variable  $\mathbf{X}$  has the probability distribution  $P$ . If plotted against probability, this corresponds to a distorted entropy function Figure 10.7, with the maximum at  $P$ .

Re-define information gain as the expected reduction in the re-centered entropy of  $\mathbf{X}$  from an observation  $\mathbf{O}$ , with the entropy centered about the distribution of  $\mathbf{X}$  prior to observing  $\mathbf{O}$ . This is equivalent to the regular information gain but under the condition that  $\mathbf{X}$  has a uniform distribution prior to the observation.

$$(Eqn\ 10-15) \quad IG^*(\mathbf{X},\mathbf{O}) = E_{\mathbf{O}}\{H^*(\mathbf{X};P(\mathbf{X})) - H^*(\mathbf{X}|\mathbf{O};P(\mathbf{X}))\}$$

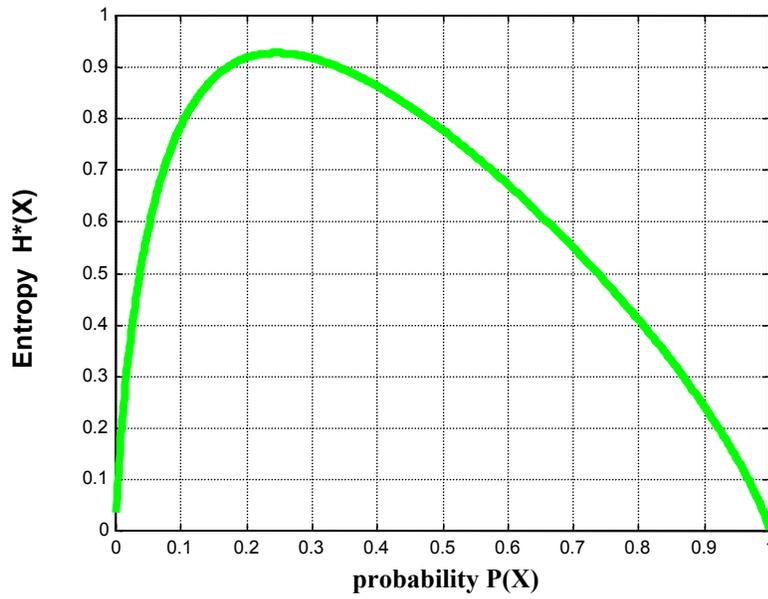


Figure 10.7 “Re-centered” entropy  $H^*(X;[0.25, 0.75])$ , peaking at  $P(X)=[0.25, 0.75]$ , instead of the uniform distribution  $[0.5, 0.5]$ .

### 10.1.6. Experimental results with re-centered entropy

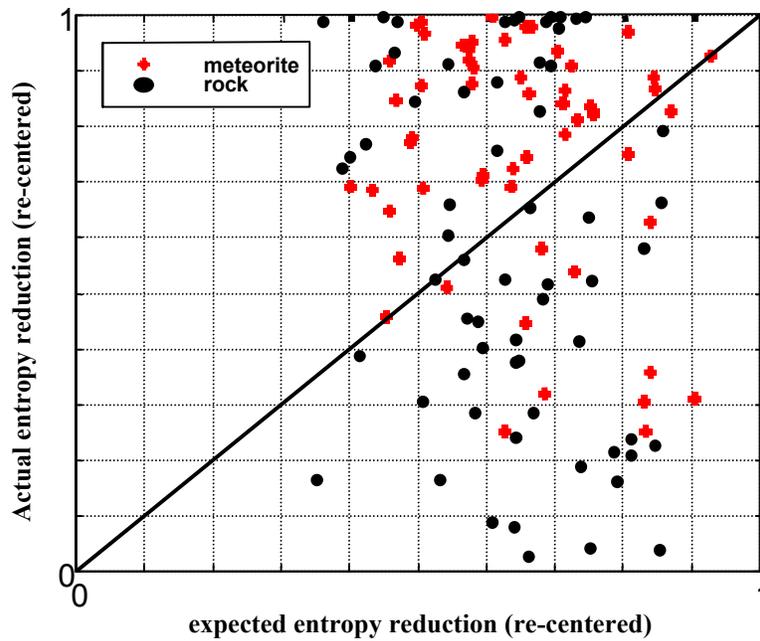
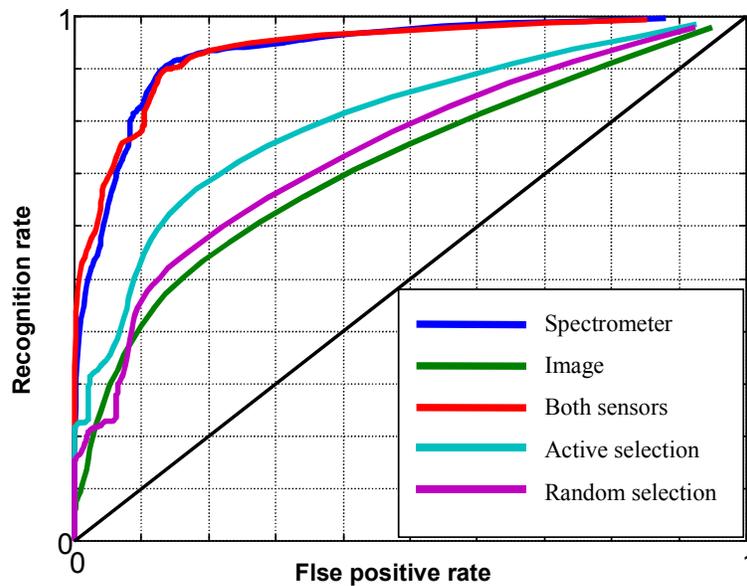


Figure 10.8 Actual decrease in entropy, re-centered about  $P(\text{meteorite} \setminus \text{image})$ , after entering spectral data into the Nomad Bayes network classifier, versus the expected decrease from a spectral reading, subsequent to classifying based on an image alone.

Using the Nomad classifier along with the laboratory image and spectral data used to test it, Figure 10.8 was generated. Each terrestrial rock and meteorite in the data set is first classified based on an image. The expected decrease in entropy from a spectral reading, re-centered about  $P(\text{meteorite} | \text{image})$ , is then calculated. The figure shows the actual decrease in re-centered entropy after entering spectral data into the Bayes network classifiers.

From Figure 10.8 it can be seen that (i) re-centered entropy always decreases as expected, and that (ii) the values expected and actual re-centered entropy reductions for the sample set range across the permitted values, instead of being clustered about zero as in Figure 10.4.

While Figure 10.8 shows a correlation between actual and expected reduction in re-centered entropy from a spectral reading, it is clear that the distribution of re-centered entropy reduction's for a spectral reading has a large variance; therefore, the expected value is not a good predictor of the actual entropy reduction. Improved image classification might narrow the distribution. This is certainly the case in Figure 10.4, where spectral classification is good enough that the distribution on entropy reductions from a metal detector reading is narrow and information gain is therefore a good predictor of the actual entropy reduction. In that case it was good enough that sensor selection could be achieved using a simple threshold on the expected reduction.



**Figure 10.9** ROC classification curves showing performance of Nomad classifier for recognizing meteorites. Best performance is achieved using the spectrometer, and the worst performance is with the camera only. Active sensor selection, using the spectral data only if the image data predicts a sufficient probability of the re-centered entropy reduction exceeding a threshold value, improves classification more so than random use of the spectrometer after an image would do so. These curves are computed for a 70% reduction in the use of the spectrometer.

A sensor selection strategy for the data in Figure 10.8 is to reserve use of spectral data for samples on which the chance of the re-centered entropy reduction exceeding some threshold is sufficiently large. This is simple to compute by sampling the perceived distribution of entropy reduction. This is already done for

the computation of the expected value. Figure 10.9 shows the effect of doing this on classification performance on laboratory data, with a 70% reduction in use of spectral data.

## 10.2. Conclusions

Active sensor selection would clearly enhance to capabilities of a robotic explorer. By economizing on sensor usage it speeds up the rate of exploration and conserves energy and other resources. These are serious considerations for a small power limited robotic vehicle, such as a solar powered planetary surface vehicle like Sojourner.

Information gain, from information theory, is in some cases, a good predictor of the utility of using a sensor. It works best on predicting sensor use on samples that have already been fairly well characterized, which somewhat defeats the purpose of trying to minimize use of the expensive short range sensors that can best characterize a sample in the first place.

By tinkering with the concept of entropy and information gain some of the problems with these metrics of sensor effectiveness can be overcome in a principled manner. This has worked well enough for a rudimentary demonstration of sensor selection using the Nomad Bayes network classifier for meteorite and terrestrial rock discrimination. Unfortunately, it appears that a significant improvement in the classifier's ability to interpret images is needed for this approach to work well.

An alternative approach to the aforementioned modifications to the concept of entropy is to use the concept of mutual information instead of entropy, and KL divergence as information gain. Given an initial uniform distribution these reduce to the former. More information can be found in [10.3].

A very different approach to sensor selection involves augmenting the Bayes network classifier with decision and utility nodes, creating an *influence diagram* [10.8][10.9]. From this one can compute *expected value of perfect information* (EVPI) and the *expected value of imperfect information* (EVII) [10.7] for a rock sample very time a decision must be made regarding sensor deployment. The advantage of these measures is that they incorporate the perceived (monetary or other) value of a meteorite and the cost of collecting terrestrial rocks. Combined with a sensor deployment cost model they permit an optimal sensor or near optimal allocation strategy in which a sensor is used if the EVII or EVPI from that sensor exceeds the perceived cost of using it. Compare this to using information gain, which relies on *ad hoc* thresholds and does not explicitly consider cost, but aims merely to reduce sensor use. However, it does have the advantage of relying less on poorly defined statistical priors and cost models. [10.10] suggest a hybrid approach that might be applicable to the robot sensor selection problem.

## 10.3. References

- [10.1] Dean, T., K. Basye, M. Lejter, "Planning and Active Perception", *Proceedings of the DARPA Workshop on Innovative Approaches to Planning, Scheduling and Control*, DARPA, 1990.
- [10.2] Raiffa, H., *Decision Analysis, Introductory Lectures on Choices Under Uncertainty*, McGraw-Hill, 1997.
- [10.3] Cover, T., and Thomas, J., *Elements of Information Theory*, New York: Wiley, 1991.
- [10.4] Tom Mitchell, *Machine Learning*, McGraw-Hill Companies, Inc. 1997

- [10.5] Elfes, A., "Dynamic Control of Robot Perception Using Stochastic Spatial Models", *Information Processing in Mobile Robots*, Schmidt (ed.), Springer-Verlag, 1991.
- [10.6] Barr, D.R., and Sherrill, E.T., "Measuring information gain in tactical operations", *Technical report of the Operations Research Center United States Military Academy*, 1996.
- [10.7] Clamen, R.T., *Making Hard Decisions: an introduction to decision analysis*, 2<sup>nd</sup> edition. Chapter 12: Value of Information.
- [10.8] Shachter, Ross D., "Probabilistic inference and influence diagrams", *Operations Research*, vol. 36, No. 4. July – August 1988.
- [10.9] Shachter, Ross D., "Evaluating influence diagrams", *Operations Research*, vol. 34, No. 6. November-December 1986.
- [10.10] Carenini, G., Monti, S., Banks, G., "An information-based Bayesian approach to History Taking", *5th Conference of AI in Medicine Europe (AIME '95)*, pages 129-138. Pavia, Italy, June 1999.

## Accomplishment and conclusions

The rock and meteorite classification system on Nomad made the first autonomous identification of a meteorite. Nomad's classification system is unique in the way it addresses robotic concerns, handles multiple sensors and can be extended to integrate data taken from many samples over an area.

Classification performance for rocks not stained with iron oxide is reasonable given the quality of the data and the difficulty of the task. Assuming a 20% false positive rate is acceptable, an 80% meteorite recognition rate is achieved. For the 5 meteorites and 35 non-stained rocks in the test set this translates to collecting 4/5 meteorites and 7/35 terrestrial rocks.

This is hard to compare directly with human meteorite collection rates, although humans are still clearly superior. During the 2000 Antarctic field season, U.S. researchers collected 626 meteorites. It is not known how many they missed, or the fraction of terrestrial rocks they brought back as meteorites. However, experience indicates that high recognition rates and low false positive rates (less than 1%) are the norm. The limiting factor for humans is how many rocks they can examine, not their ability to distinguish meteorites from terrestrial rocks that have been brought to their attention.

There are no other robotic systems that can be directly compared to the Nomad classifier. [11.1] demonstrated a rule-based classifier for recognizing carbonates from reflectance spectra obtained using the Marsokhod rover. They report an astonishing 80% recognition rate with a 0% false positive rate. However, apart from the fact that they are attempting to recognize a different class of objects, there are a number of key differences between the Marsokhod and the Nomad systems that complicate a direct comparison:

- The Marsokhod classifier incorporated spectral data quality metrics. Sensor data considered too noisy was rejected and not classified. In fact, 43% of the spectra acquired from carbonates, and 55% of the spectra acquired from non-carbonates were rejected. The Nomad classifier attempts to classify all data.
- An extended range spectrometer, measuring wavelengths up to 2500 nm (mid infrared), was used on the Marsokhod rover. Carbonates have distinctive spectral absorption features around 2300 nm. By comparison, the Nomad spectrometer gives good data from 400 – 1000 nm. Meteorites and many terrestrial rocks do not have very distinctive spectral features in this wavelength range. Furthermore, meteorites contain similar minerals to most of the rocks in Elephant Moraine, and therefore have similar spectra. Carbonates are a class of minerals, not a rock type. There is less variation in their spectra.
- Sensor data acquisition is not fully autonomous. Targets are designated by human operators, and are larger than the typical rocks encountered by Nomad.

Other spectral classification systems [11.2][11.3][11.4] that exist do not attempt to distinguish meteorites from rocks. They are concerned with recognizing the spectra of pure minerals in powdered form.

### *Robotics and scientific discovery*

This thesis demonstrates the combined potential of robotics and artificial intelligence for scientific exploration. Up until now, scientific discovery has been a very human centered process. Even with the advent of robotic spacecraft and rovers to explore areas inaccessible to human explorers, human scientists

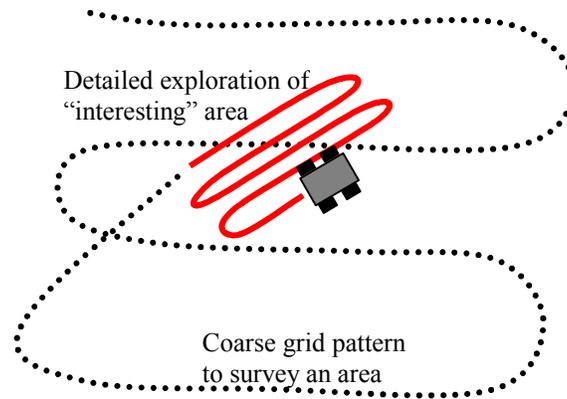
have completely dominated the game. Robots have been relegated to the role of passive data collectors or platforms for carrying instruments.

As robotic technology has improved, mobile machines capable of autonomously taking care of themselves have evolved. Obstacle avoidance, path planning and locomotion technologies are mature enough that a robotic vehicle can autonomously take itself to a scientifically interesting location and deploy sensors. However, it is still human scientists who decide what a robot should investigate and how.

The logical next step in this evolution is a robot that can itself decide what to investigate and how to do so. Bandwidth and other considerations make it impossible for modern robotic explorers in remote areas, such as Mars, to return all the data sensor data they are capable of acquiring. By interpreting sensor data themselves they can investigate an area more thoroughly, communicating the most interesting discoveries back to mission control.

The Nomad rock and meteorite classification system is a step in this direction. It is the first complete robotic system with explicit autonomous onboard scientific data interpretation capabilities. Looking for, and recognizing meteorites amongst terrestrial rocks is a beginning. Robotically recognizing a variety of rocks in an area is the next step.

It is my belief that this thesis establishes the viability of autonomous robotic exploration and onboard scientific analysis. Furthermore, the Bayes network generative model described in this thesis is a principled framework in which to interpret data from the many sensors a robotic explorer carries. It tolerates ambiguity, noise and works with whatever data is available. Very importantly, it provides a means for the robot to make decisions on how to proceed with its scientific exploration. At the very least, this means deciding what sensors to use in order to identify or otherwise characterize an object, a demonstrated capability of the Nomad classifier. More advanced is the ability for the robot to decide where to explore and what to look for.



**Figure 11.1** An application of using the autonomous profiling capability to allow a robot to itself decide where to explore or search. An area is first coarsely surveyed by the robot to identify the general distributions of rocks types (or whatever else is of interest) over the area. Having done so it can identify those regions most worthy of a more detailed search. In the case of meteorite search, the best places to search are those where the terrestrial rocks are clearly different from meteorites, making it much easier to find meteorites amongst them. A robot looking for fossils on Mars would benefit from searching areas with carbonate deposits.

The classification system of Nomad extends in a natural way to the autonomous profiling of an area [11.5](Chapter 9). This ability is particularly valuable for robotic exploration. It improves classification by relating sensor readings from many objects in the area and is scientifically valuable in its own right. Moreover, it is a tool that a robot could use to decide where next to explore, or what areas may be most profitably searched (Figure 11.1).

## **Lessons – “Perspectives and Wisdom”**

### *Importance of field training*

The experience in Elephant Moraine illustrated how difficult it can be to anticipate field conditions. The Nomad classifier initially performed poorly at the Elephant Moraine compared to its performance on the laboratory data used to train it. Not only were field conditions different, but also the rocks were different from those in the laboratory data sets and the autonomous placement of the sensors had a significant effect on the data.

Because of the poor performance, it was necessary to obtain more training data in the field, and use it to update the classifier. Even though only a limited number of field training samples could be obtained, they significantly improved the classifier performance.

This indicates the importance of realistic training data, preferably obtained in *by the robot in field conditions*. As this is not always possible, future robotic systems may require a period of *in situ* training with expert supervision.

### *Sensor selection*

In areas with many rocks, Nomad spent approximately 17% of its time deploying the camera and 49% of its time deploying the manipulator arm. Away from the moraine, where the rock density was lower, Nomad spent 9% and 26% of its time on these activities respectively. These figures show the potential timesavings if active sensor selection were available, whereby the spectrometer is only used on those samples the classifier deems it to be useful based on a close-up image. Unfortunately, without proper camera images it was not possible to evaluate this capability. Results using information gain on laboratory data have shown limited promise and are a subject of ongoing research.

### *Quality of sensor data*

Bad sensor data is, unfortunately, a common problem when sensors are autonomously deployed in the field. The care and effort that a human researcher will take to ensure good measurements is considerable. An autonomous system in the field is not nearly as capable, and may be subject to unforeseen systematic errors.

The Nomad spectrometer design proved to be fundamental to the project success. It performed well in the extreme field conditions and was tolerant of sensor placement errors due to the autonomous deployment procedure. Nonetheless, the spectrometer probe would often be placed so badly that nonsense spectra were obtained. Sometimes it was placed in such a manner as to not get spectral data from the target rock at all. In these cases, the classifier output could be quite arbitrary.

Recognizing bad data would be a significant improvement. The Marsokhod system [11.1] had this capability, partly explaining its excellent ability to recognize carbonates.

Modeling the failure modes of a sensor can do this. Many failures leave a distinctive signature in the returned data. For example, the obviously bad spectra obtained by Nomad are characterized by sensor saturation. Explicit sensors to verify correct data collection might be needed. If Nomad had a better system to verify spectrometer placement, the better data might have been obtained.

Nomad's Bayesian classifier could conceivably recognize bad data in a simple way. Because it models the complete distribution of sensor data for many known rock types, it can recognize if a measurement is a statistical outlier. In other words, the computed conditional likelihood of that data measurement is too low to have been generated by any rock types it knows. This has the advantage of being sensitive to unanticipated sensor failure modes. Unfortunately, it can also exclude data from novel objects, reducing the systems ability to recognize an exceptional sample.

#### *Need for explicit photometric calibration of camera*

The correct measurement of albedo and color are necessary if rocks are to be classified with any success from color images. The failure of Nomad's camera system to obtain images accurately portraying the brightness and color of the rock samples on the ice meant that it could not be used for classification. Explicit use of a standard calibration target, like those used on some spacecraft, to determine the correct color balance and the aperture that optimizes the images of rocks given the lighting would have corrected this.

#### *Difficulty of rock classification from spectra*

Recognizing a plethora of different rock types using visible light reflectance spectroscopy alone is difficult. The most diagnostic features usually occur out of the visible range. While Nomad's classifier can distinguish meteorites from terrestrial rocks by their spectra with reasonable success, it lags behind the ability of humans to tell the two apart by looking at them. Nomad's ability to distinguish amongst many different terrestrial rocks is worse than it's ability to recognize meteorites.

Nomad's ability to distinguish meteorites from terrestrial rocks would have been much better if it could have used as metal detector in addition to the camera and spectrometer. However, a metal detector is not very useful for the more general problem of recognizing the many varieties of terrestrial rock types, nor would it recognize the rare and valuable Achondrite meteorites.

A diversity of sensors is essential for recognizing rock types. No single one sensor is good for recognizing all the possible types. Attempting to do so with a visible light reflectance spectrometer is hard. Better performance is possible with more sensors such as a metal detector [11.5].

Laser Induced Breakdown Spectroscopy (LIBS)[11.6], is a very promising sensing modality for this task. A small section of rock is vaporized, and the vapor ionized, by an intense laser pulse. As the ionized gas returns to the ground state, a sharp and well-defined emission spectrum is produced. This allows the precise determination of elemental abundances in the rock sample.

Most meteorites and Antarctic rocks can be discriminated from each other based on knowledge of elemental abundances, as typical meteorites contain significant amounts of iron and nickel. However, LIBS does not necessarily determine the mineral composition of a rock sample, as different minerals can contain similar amounts of the same elements, but arranged in a different molecular or crystalline structure. Having multiple minerals and contaminants in rock sample further complicates the situation. Nor can LIBS

distinguish between rocks composed of identical minerals but with different grain sizes (Dolerites and Basalts for example). Nonetheless, knowledge of elemental abundances greatly constrains the possible minerals in a rock, their relative proportions, and therefore the rock type. Combined with other data, using the Bayesian formalism although not necessarily a traditional Bayes network, LIBS might be the tool of choice for the robust identification of rock types and autonomous geology.

The biggest advantage of LIBS, from the robotic operational perspective, is that measurements can be done remotely (up to 10m distance), removing the need for expensive and time-consuming placement of sensors with a robot arm. Furthermore, the initial laser pulse can sufficiently ablate the rock surface to obtain measurements from un-weathered material beneath the surface.

Unfortunately, because LIBS burns a small pit into the samples being analyzed, it may cause concern to the meteorite science community for the purpose of robotic meteorite search. Although the damage to a sample is minor, and hardly compares to that suffered during atmospheric entry, it is nonetheless a potential factor.

## **Other applications and extensions**

This thesis work clearly has applications beyond robotic meteorite search in Antarctica. It establishes a paradigm for the onboard interpretation of scientific sensor data by a robotic explorer, enabling it to recognize objects of interest, characterize a geographic area and make decisions about where to explore next and what sensors to use.

The obvious application of this work is, as discussed earlier, the robotic scientific exploration of remote environments, such as Mars or the oceans of Europa. Immediate communication of information back to mission control is not feasible from these places and a robot investigator needs to exercise considerable autonomy in the way it conducts a scientific investigation.

An important terrestrial application is landmine clearance. Currently, this is a very labor intensive and dangerous task. A robotic solution could potentially save many lives. The issues in robotic landmine detection and removal are almost identical to those faced by a robotic meteorite searcher, except that there is a greater cost for not finding a landmine.

There are many possible sensors for finding landmines, each with advantages and disadvantages. Sensitive metal detectors can detect most mines, even the supposedly non-metallic varieties that still have metallic firing pins. Unfortunately, these sensors also pick up many other things. Beer bottle tops, cans, nails, bullets, certain soils and other things all register with these metal detectors. Other sensors, such as ground penetrating radar, magnetometers, infrared cameras, and thermal neutron emission detectors can also detect some landmines and other things. They are also very expensive, and in the case of the thermal neutron detector, very time consuming to deploy.

Compounding evidence from these different sensors to find landmines and discriminate them from the various possible false positives is identical to meteorite search, except that different sensors are used. Use of the more expensive sensors must be curtailed where possible. Landmines are not distributed randomly but are often clustered. The autonomous profiling algorithm can exploit this, both to improve classification and to locate areas where landmines are most likely to be found. Furthermore, knowing the distribution of false targets over an area can influence the choice of sensors and search strategy.

### *Extension to further sensors*

The obvious extension of this work would be to include sensor data from a Laser Induced Breakdown Spectrometer (LIBS), as indicated above. Indeed, had LIBS been available earlier in the project, and assuming that the fears of the scientific community could be assuaged, it would almost certainly have been used in addition to reflection spectroscopy.

Doing this may require an explicit closed-form model for computing the distribution of elemental abundances for each rock type, as this could be intractable to represent using the standard Bayes network representation. Further research is necessary to evaluate this claim.

Nonetheless, assuming contaminants are not present, it is reasonable to assume that LIBS measurements are conditionally independent of reflection spectra given the mineral composition of a rock. This is because they operate in different bond energy regimes, and are affected by different mechanisms. This makes it significantly easier to fuse these two sensor modalities.

Extending the Bayes network approach to handle many more closely related sensors is a lot harder and is perhaps the biggest bottleneck suffered by this approach, since each sensor must be explicitly modeled using expert knowledge, and sensor interactions accounted for. As more sensors are added, the additional statistical dependencies between them will eventually make the Bayes network approach intractable.

### *Extension to multiple vehicles*

Multiple robotic vehicles are needed order to thoroughly search or explore an area within a given time frame. Each might be equipped with different sensors and have different locomotion capabilities. High-speed vehicles with long-range sensors would rapidly search an area, locating targets of interest and characterizing the area as best as they can. Subsequently, other robots with short-range but more precise sensors would be directed to investigate selected targets.

For example, consider the problem of searching for fossils on Mars. An aerial vehicle equipped with an imaging spectrometer could be used to locate areas with high concentrations of carbonate deposits, and therefore a greater likelihood of fossils, assuming life ever existed in the area. A ground vehicle would then investigate those areas to locate the most promising rocks using its onboard suite of sensors. Samples of these rocks would be collected and taken to a centralized “mothership” for further analysis and possible return to Earth.

The Bayes network generative model paradigm might be profitably extended to handle the multi-vehicle scenario. The initial survey of an area could be handled by the autonomous profiling algorithm. Deciding which vehicles to send to which targets is an extension of the sensor selection problem.

Nevertheless, extending the rock classification system of this thesis to handle multiple vehicles presents significant challenges. The vehicles need to behave in a coordinated fashion, yet operate autonomously from each other, sharing only the minimal amount of information needed to accomplish the task. Treating each vehicle as an extended sensor of the mothership, building a centralized model of the area may require too much bandwidth [11.8]. Furthermore, it is not easily scalable. Adding more vehicles with new capabilities requires a redesign of the mothership.

Another problem is registering the measurements from different vehicles. Because of navigation errors it is possible that a robot commanded to examine a sample looked at earlier by another vehicle might

erroneously investigate a different sample. This problem is significantly worsened if the robots are operating in a dynamic environment, such as an undersea vent for example.

## References

- [11.1] Roush, T. and Gazis, P., “Autonomous Identification of Carbonates Using Near-IR Reflectance Spectra During the February 1999 Marsokhod Field Tests”, in *Journal of Geophysical Research*, vol. 106, April 2001.
- [11.2] Glymour, C., Ramsey, J., Roush, T., Scheines, R., Richardson, T., Bailer, H., “Automated Identification of Mineral Composition from Reflectance Spectra”, *Draft interim report*, July 24, 1998.
- [11.3] Clark, R.N. and Swayze, G.A., “Mapping Minerals, Amorphous Materials, Environmental Materials, Vegetation, Water, Ice and Snow, and Other Materials: The USGS Tricorder Algorithm”, *Summaries of the Fifth Annual JPL Airborne Earth Science Workshop*, January 23- 26, R.O. Green, Ed., JPL Publication 95-1, p. 39-40, 1995.
- [11.4] Clark, R.N., G.A. Swayze, A. Gallagher, N. Gorelick, and F. Kruse, “Mapping with Imaging Spectrometer Data Using the Complete Band Shape Least-Squares Algorithm Simultaneously Fit to Multiple Spectral Features from Multiple Materials”, *Proceedings of the Third Airborne Visible/Infrared Imaging Spectrometer (AVIRIS) Workshop*, JPL Publication 91-28, 2-3, 1991.
- [11.5] Pedersen, Liam, “Robotic Deployment of Electro-Magnetic Sensors for Meteorite Search”, IEEE International Conference on Robotics and Automation, Leuven, Belgium, 1998.
- [11.6] Blacic, J.D., Pettit, D.R., and Cremers, D.A., “Laser-Induced Breakdown Spectroscopy for Remote Elemental Analysis of Planetary Surfaces”, *Proceedings of the International Symposium on Spectral Sensing Research*, Maui, HI, November 15-20, 1992.
- [11.7] Pedersen, Liam, “Autonomous Characterization of Unknown Environments”, in *IEEE International Conference on Robotics and Automation*, Seoul, Korea, 2001.
- [11.8] Estlin, T., Gray A., Mann, T., Rabideau, G., Castaño, R., Chien, S., and Mjolsness, E., “An Integrated System for Multi-Rover Scientific Exploration”, *AAAI-99*, 1999.

## Master Reference

- [1] Apostolopoulos, D., Wagner, M., Shamah, B., Pedersen, L., K. Shillcutt, Whittaker, W., “Technology and Field Demonstration of Robotic Search for Antarctic Meteorites”, *International Journal of Robotic Research*, special Field & Service Robotics issue, December 2000.
- [2] Apostolopoulos, D., Wagner, M., Whittaker, W., “Technology and Field Demonstration Results in the Robotic Search for Antarctic Meteorites”, *Field and Service Robotics*, Pittsburgh, USA, 1999.
- [3] Barr, D.R. and Sherrill, E.T., “Measuring information gain in tactical operations”, Technical report of the Operations Research Center United States Military Academy, 1996.
- [4] Bishop, C. M., *Neural Networks for Pattern Recognition*, Oxford University Press, 1995.
- [5] Blacic, J.D., Pettit, D.R., and Cremers, D.A., “Laser-Induced Breakdown Spectroscopy for Remote Elemental Analysis of Planetary Surfaces”, *Proceedings of the International Symposium on Spectral Sensing Research*, Maui, HI, November 15-20, 1992.
- [6] Blatt, H., Jones, R.L., “Proportions of Exposed Igneous, Metamorphic, and Sedimentary Rocks”, *Geological Society of America Bulletin*, v. 86, pp 1085-1088, August 1975.
- [7] Burges, Christopher J.C., “A Tutorial on Support Vector Machines for Pattern Recognition”, *Knowledge Discovery and Data Mining*, 2(2), 1999.
- [8] Carenini, G., Monti, S., Banks, G. “An information-based Bayesian approach to History Taking”, *5th Conference of AI in Medicine Europe (AIME '95)*, pages 129-138. Pavia, Italy, June 1999.
- [9] Cassidy, W., Harvey, R., Schutt, J., Delisle, G., and Yanai, K., “The Meteorite Collection Sites of Antarctica”, *Meteoritics*, 27(5), pp 490 – 525, 1992.
- [10] Castillo, Enrique Jose Manuel Gutierrez and Hadi, Ali S., *Expert Systems and Probabilistic Network Models*, Springer, New York, 1997.
- [11] Charniak, Eugene “Bayesian networks without tears” in *AI Magazine*, 12(4), 50-63, Winter 1991.
- [12] Chesney, R.H., Das, Y., McFee, J.E., and Ito, M.R., “Identification of Metallic Spheroids by Classification of Their Electromagnetic Induction Responses”, *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. PAMI-6, no. 6, November 1984.
- [13] Clamen, Robert T., *Making Hard Decisions: an introduction to decision analysis*, 2<sup>nd</sup> edition. Chapter 12: Value of Information.
- [14] Clark, R.N. “Spectroscopy and Principles of Spectroscopy”, *Manual of Remote Sensing*, A. Rencz, Editor, John Wiley and Sons, Inc. 1999.
- [15] Clark, R.N. and Swayze, G.A., “Mapping Minerals, Amorphous Materials, Environmental Materials, Vegetation, Water, Ice and Snow, and Other Materials: The USGS Tricorder Algorithm”, *Summaries of the Fifth Annual JPL Airborne Earth Science Workshop*, January 23- 26, R.O. Green, Ed., JPL Publication 95-1, p. 39-40, 1995.

- [16] Clark, R.N., A.J. Gallagher, and G.A. Swayze, "Material Absorption Band Depth Mapping of Imaging Spectrometer Data Using a Complete Band Shape Least-Squares Fit with Library Reference Spectra", *Proceedings of the Second Airborne Visible/Infrared Imaging Spectrometer (AVIRIS) Workshop*. JPL Publication 90-54, 176-186, 1990.
- [17] Clark, R.N., G.A. Swayze, A. Gallagher, N. Gorelick, and F. Kruse, "Mapping with Imaging Spectrometer Data Using the Complete Band Shape Least-Squares Algorithm Simultaneously Fit to Multiple Spectral Features from Multiple Materials", *Proceedings of the Third Airborne Visible/Infrared Imaging Spectrometer (AVIRIS) Workshop*, JPL Publication 91-28, 2-3, 1991.
- [18] Clark, R.N., G.A. Swayze, C. Koch, A. Gallagher, and C. Ager, "Mapping Vegetation Types with the Multiple Spectral Feature Mapping Algorithm in both Emission and Absorption", *Summaries of the Third Annual JPL Airborne Geosciences Workshop, Volume 1: AVIRIS Workshop*. JPL Publication 92-14, 60-62, 1992.
- [19] Clark, R.N., T.V.V. King, and N.S. Gorelick, "Automatic Continuum Analysis of Reflectance Spectra", *Proceedings of the Third Airborne Imaging Spectrometer Data Analysis Workshop*, JPL Publication 87-30, pp 138-142, 1987.
- [20] Clark, R.N., T.V.V. King, M. Klejwa, G. Swayze, and N. Vergo, "High Spectral Resolution Reflectance Spectroscopy of Minerals", *Journal Geophysical Research*, 95, pp12653-12680, 1990.
- [21] Cover, T., and Thomas, J., *Elements of Information Theory*, New York: Wiley, 1991.
- [22] Das, Y., McFee, J., Toews, J., Stuart, G.C., "Analysis of an Electromagnetic Induction Detector for Real-Time Location of Buried Objects", *IEEE Transactions on Geoscience and Remote Sensing*, vol. 28, no. 3, May 1990.
- [23] Das, Y., McFee, J.E., and Chesney, R.H., "Determination of Depth of Shallowly Buried Objects by Electromagnetic Induction", *IEEE Transactions on Geoscience and Remote Sensing*, vol. GE-23, no. 1, January 1985.
- [24] Dean, T., K. Basye, M. Lejter, "Planning and Active Perception", *Proceedings of the DARPA Workshop on Innovative Approaches to Planning, Scheduling and Control*, DARPA, 1990.
- [25] Dempster, A.P., "A Generalization of Bayesian Inference", *Journal of the Royal Statistical Society, Series, B* 30:pp205-247, 1968.
- [26] Dietrich, R.V. and Skinner, B.J., *Rocks and Rock Minerals*, John Wiley & Sons, Inc, 1979.
- [27] Dodd, R.T., *Meteorites: a Petrologic-Chemical Synthesis*, Cambridge University Press.
- [28] E. Gelenbe, N. Schmajuk, J. Staddon, J. Reif, "Autonomous search by robots and animals: a survey", *Robotics and Autonomous Systems* 22, pp 23-34. 1997.
- [29] E. Gelenbe, Y. Cao, "Autonomous search for mines", *European Journal of Operational Research* 108, pp319-333. 1998.
- [30] Elfes, A., "Dynamic Control of Robot Perception Using Stochastic Spatial Models", *Information Processing in Mobile Robots*, Schmidt (ed.), Springer Verlag, 1991

- [31] Elfes, A., "Occupancy grids: a stochastic spatial representation for active robot perception", *Autonomous Mobile Robots: Perception, Mapping and Navigation*, S. Iyengar, A. Elfes, Eds. IEEE Comp. Soc. Press, 1991.
- [32] Elfes, A., "Robot Navigation: Integrating Perception, Environmental Constraints and Task Execution Within a Probabilistic Framework", *Reasoning with Uncertainty in Robotics*, L. Dorst, M. van Lambalgen and F. Voorbraak (Eds), Amsterdam, 1995.
- [33] Estlin, T., A. Gray, T. Mann, G. Rabideau, R. Castaño, S. Chien, and E. Mjolsness, "An Integrated System for Multi-Rover Scientific Exploration", *AAAI-99*, 1999.
- [34] Faugeras, Olivier, "Fundamentals in Computer Vision: an advanced course", Cambridge University Press, 1983.
- [35] Faure, G. and Taylor, K.S. "The Geology and Origin of the Elephant Moraine on the East Antarctic Ice Sheet", *Antarctic Journal of the U.S.*, 20(5), pp 11-12, 1985.
- [36] Faure, G., M.L. Kallstrom, and Mensing, T.M., "Classification and Age of Terrestrial Boulders in the Elephant and Reckling Moraines", *Antarctic Journal of the U.S.*, 19(5), pp28-29, 1984.
- [37] Faure, G., Taylor, K.S., and Jones, L.M., "Hydrothermal Calcite in the Elephant Moraine", *Antarctic Journal of the U.S.*, 21(5), pp 21, 1986.
- [38] Fukunaga, K., *Statistical Pattern Recognition*, 2<sup>nd</sup> ed, Academic Press, 1990.
- [39] Gelman, A., Carlin, J.B., Stern, H.S. and Rubin, D.B., *Bayesian Data Analysis*, Chapman & Hall, 1995.
- [40] Glymour, C., Ramsey, J., Roush, T., "Comparison of Human and Machine Identification of Mineral Classes from Reflection Spectra", *Draft interim report*, January, 1999.
- [41] Glymour, C., Ramsey, J., Roush, T., Scheines, R., Richardson, T., Bailer, H., "Automated Identification of Mineral Composition from Reflectance Spectra", *Draft interim report*, July 24, 1998.
- [42] Healey, G., Slater, D., "Image Understanding Research at UC Irvine: Automated Invariant Recognition in Hyperspectral Imagery".
- [43] Heide, Fritz, *Meteorites*, University of Chicago Press, 1964.
- [44] Henrion, Max, John S. Breese and Eric J. Horvitz (1991) "Decision Analysis and Expert Systems" in *AI Magazine*, 12(4), 64-91. Winter 1991.
- [45] Hunt, G., "Spectral Signatures of Particulate Minerals in the Visible and Near Infrared", *Geophysics*, vol. 42, no.3, p. 501-513, April 1977.
- [46] Jain, A. K., *Fundamentals of Digital Image Processing*, Prentice-Hall, 1989.
- [47] Jensen, Finn V. *An Introduction to Bayesian Networks*, Springer-Verlag, New York. 1996.
- [48] Jensen, Frank, Finn V. Jensen and Soren L. Dittmer "From influence diagrams to junction trees" in Proceedings of 10<sup>th</sup> conference on *Uncertainty in Artificial Intelligence*, Seattle, 1994.
- [49] Jordan, M.I. and Bishop, C.M., "Neural Networks", *CRC Handbook of Computer Science*, A. Tucker (ed.), CRC Press, 1996.

- [50] Jordan, M.I., “Why the Logistic Function? A Tutorial Discussion on Probabilities and Neural Networks”, *Computational Cognitive Science Technical Report 9503*, August 1995.
- [51] Krotkov, E., “ Focusing”, *International Journal of Computer Vision*, vol. 1, no. 3, October 1987, pp. 223-237. (Reprinted in *Physics-Based Vision*, Jones and Bartlett, 1992.)
- [52] Lauritzen, Steffen L. and Spiegelhalter, David J., “Local computations with probabilities on graphical structures and their application to expert systems” in *Journal of the Royal Statistics Society B*, 50(2), 157-194. 1998.
- [53] Lauritzen, Steffen L., “The EM algorithm for graphical association models with missing data” in *Computational Statistics and Data Analysis*, 19(2), 191-201. 1995.
- [54] McFee, J. and Das, Y., “Development of an Intelligent System to Locate and Identify Ferrous Objects”, *Artificial Intelligence, Expert Systems and Languages in Modeling and Simulation*, C.A. Kulikowski, R.M. Huber and G.A. Ferrate (eds.), 1998.
- [55] McFee, J.; and Das, Y.; “Determination of the Parameters of a Dipole by Measurement of the Magnetic Field”, *IEEE Transactions on Antennas and Propagation*, Vol. AP-29, No. 2, March 1981.
- [56] McFee, J.; Das, Y.; and Ellingson, R., “Locating and Identifying Compact Ferrous Objects”, *IEEE Transactions on Geoscience and Remote Sensing*, Vol. 28, No. 2, March 1990.
- [57] McSween, H.Y., *Meteorites and Their Parent Planets*, Cambridge University Press, pp 6-7, 1987.
- [58] Mitchell, T., *Machine Learning*, McGraw-Hill Companies, Inc. 1997.
- [59] Neapolitan, Richard E. *Probabilistic Reasoning in Expert Systems: Theory and Algorithms*, John Wiley & Sons, New York, 1990.
- [60] Netica; <http://www.norsys.com>
- [61] Papoulis, A., Probability, *Random Variables, and Stochastic Processes*, 3rd edition, McGraw-Hill, Inc., 1991.
- [62] Pearl, J., *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*, Morgan Kaufman, San Mateo, CA. 2<sup>nd</sup> edition, 1991.
- [63] Pedersen, L., “Autonomous Characterization of Unknown Environments”, in *IEEE International Conference on Robotics and Automation*, Seoul, Korea, 2001.
- [64] Pedersen, L., “Robotic Deployment of Electro-Magnetic Sensors for Meteorite Search”, *IEEE International Conference on Robotics and Automation*, Leuven, Belgium, 1998.
- [65] Pedersen, L., Apostolopoulos, D., Whittaker, W. “Bayes Networks on Ice: Robotic Search for Antarctic Meteorites”, *Neural Information Processing Symposium 2000*, Denver, Colorado, 2000.
- [66] Pedersen, L., Apostolopoulos, D., Whittaker, W., Roush, T. and G. Benedix “Sensing and data classification for robotic meteorite search” *Proceedings of SPIE Photonics East Conference*. Boston, USA, 1998.

- [67] Pedersen, L., Wagner, M., Apostolopoulos, D., Whittaker, W., “Autonomous Robotic Meteorite Identification in Antarctica”, in *IEEE International Conference on Robotics and Automation*, Seoul, Korea, 2001.
- [68] Pomerleau, D., “ALVINN: An Autonomous Land Vehicle In a Neural Network”, in *Advances in Neural Information Processing Systems 1*, Morgan Kaufmann, 1989.
- [69] Raiffa, H., *Decision Analysis, Introductory Lectures on Choices Under Uncertainty*, McGraw-Hill, 1997.
- [70] Rich, E., and Knight, K., *Artificial Intelligence*, McGraw-Hill, 1991.
- [71] Roush, T. and Gazis, P., “Autonomous Identification of Carbonates Using Near-IR Reflectance Spectra During the February 1999 Marsokhod Field Tests”, in *Journal of Geophysical Research*, vol. 106, April 2001.
- [72] Roush, T., V. Gulick, V., Morris, R., Gazis, P., Benedix, G., Glymour, C., Ramsey, J., Pedersen, L., Ruzon, M., Buntine, W., and Oliver, J., “Autonomous Science Decision Making for Mars Sample Return”, in *Lunar and Planetary Science Conference XXX*, Houston, TX, 1999.
- [73] Shachter, Ross D., “Evaluating influence diagrams”, *Operations Research*, vol. 34, No. 6. November-December 1986.
- [74] Shachter, Ross D., “Evidence absorption and propagation through evidence reversals” in Proceedings of the Fifth Workshop on *Uncertainty in Artificial Intelligence*, 303-308. 1989.
- [75] Shachter, Ross D., “Probabilistic inference and influence diagrams”, *Operations Research*, vol. 36, No. 4. July – August 1988.
- [76] Shafer, G., *A Mathematical Theory of Evidence*, Princeton, NJ: Princeton University Press, 1976.
- [77] Smith, James E., Samuel Holtzman and James E. Matheson, “Structuring conditional relationships in influence diagrams” in *Operations Research*, 41(2), 280-297. 1993.
- [78] Spiegelhalter, D.J., Philip, D.A., Lauritzen, S.L. and Cowell, R.G., “Bayesian analysis in expert systems” in *Statistical Science*, 8(3), p219-283., 1993.
- [79] Spirtes, P., “Directed Cyclic Graphical Representation of Feedback Models”, Proceedings of the Eleventh *Conference on Uncertainty in Artificial Intelligence*, ed. by Philippe Besnard and Steve Hanks, Morgan Kaufmann Publishers, Inc., San Mateo, 1995.
- [80] Spirtes, P., and Meek, C., “Learning Bayesian Networks with Discrete Variables from Data”, in *Proceedings of The First International Conference on Knowledge Discovery and Data Mining*, ed. by Usama M. Fayyad and Ramasamy Uthurusamy, AAI Press, pp. 294-299. 1995.
- [81] Spirtes, P., Glymour, C., Meek, C., “Learning Bayesian Networks from Data”, *AISB Quarterly*, 1996
- [82] Spirtes, P., Meek, C., and Richardson, T., “Causal Inference in the Presence of Latent Variables and Selection Bias” , in *Proceedings of the Eleventh Conference on Uncertainty in Artificial Intelligence* , ed P. Besnard and S. Hanks, Morgan Kaufmann Publishers, Inc., San Mateo, pp 499-506. 1995.

- [83] Taylor, K.S., *Lithologies and Distribution of Clasts in the Elephant Moraine, Allan Hills South Victoria Land, Antarctica*, M.S. Thesis submitted to Kent State University Graduate College.
- [84] Wagner, M., Apostolopoulos, D., Shillcutt, K., Shamah, B., Whittaker, W., “The Science Autonomy System of the Nomad Robot”, in *IEEE International Conference on Robotics and Automation*, Seoul, Korea, 2001.
- [85] Wasson, John T., *Meteorites: classification and properties*, Springer-Verlag, 1974.
- [86] Zadeh, L., “A Theory of Approximate Reasoning”, *Machine Intelligence 9*, ed. J. Hayes, D. Michie, and L.I. Mikulich, pp. 149-194, New York: Halstead Press, 1979.