

Autonomous Robotic Meteorite Identification in Antarctica

Liam Pedersen[†], Michael Wagner, Dimitrios Apostolopoulos, and William "Red" Whittaker
The Robotics Institute
Carnegie Mellon University
Pittsburgh PA 15213, USA

Abstract

This paper describes the development and implementation of a Bayes network based system for distinguishing terrestrial rocks from meteorites from onboard the mobile robotic rover Nomad. Equipped with a color camera and spectrometer, Nomad autonomously made the first robotic identification of a meteorite, in January 2000 at the Elephant Moraine, Antarctica. This paper discusses rock classification from a robotic platform and the challenges of autonomously obtaining good sensor data in the field, while focusing on Nomad implementation.

Introduction

Antarctica is an excellent place to find meteorites. The pristine, dry and cold environment ensures that meteorites deposited there are preserved for long periods. Subsequent glacial flow of the ice sheets in which they land concentrates them in particular areas. Furthermore, they are less likely to be contaminated by terrestrial compounds. To date, most meteorites ever recovered were found in Antarctica in the last 20 years.

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. In certain blue ice fields, the only rocks visible are meteorites. In the moraines, where the ice flow brings rocks to the surface, searchers have to contend with a multiplicity of terrestrial rocks

The Robotic Antarctic Meteorite Search (RAMS) project of Carnegie Mellon University has built and demonstrated a robot, Nomad (Figure 1)[1], to autonomously search for meteorites in Antarctica. It is capable of long duration missions in harsh environments and equipped with suitable sensors to locate and identify meteorites amongst terrestrial rocks on the ice sheet.

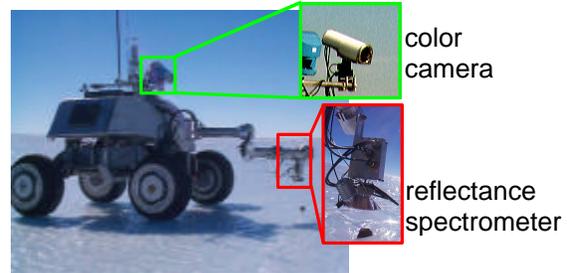


Figure 1 The Nomad robot examining potential meteorite with color camera and visible light spectrometer.

In January 2000, Nomad was deployed to Elephant Moraine in Antarctica where it made the first autonomous identification of a meteorite by a robot. The meteorite identification task goes beyond standard classification. Not only is it difficult to distinguish rocks from meteorites, autonomously taking sufficiently high quality data from a mobile robot is a serious challenge.

The eventual goal beyond Antarctic meteorite search is to develop technologies for extended robotic exploration of remote areas, including planetary surfaces. Of particular interest is autonomous onboard interpretation of scientific data [3] 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.

This paper is primarily concerned with Nomad's rock and meteorite classification system and its performance. For further details on the Nomad robot itself, including its performance in Antarctica, see the companion paper [2].

Nomad meteorite identification system

The meteorite identification system onboard Nomad consists of sensors, a classifier and a planning system to integrate these with the rest of the robot. 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

[†] Email: pedersen+@ri.cmu.edu

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.

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.

Nomad's planning system, referred to as the *Science Autonomy System* (SAS) [8] 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.

Sensors for meteorite identification

Several sensor modalities for discriminating meteorites from rocks have been explored, including metal detectors and magnetometers [4], visible and near IR reflectance spectroscopy [5] and high-resolution color imagery [5]. These are good for reasonable discrimination. However, it is also essential to consider their efficacy from onboard a robot, and whether they can be deployed autonomously.

Nomad is equipped with a color camera on a pan-tilt platform and an arm mounted reflectance spectrometer (Figure 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.

The relevant features extracted from these sensors for rock classification are detailed in [5]. They include cross sectional area and average color from color images, and the presence or absence of absorption peaks in a spectrum.

High-resolution camera

The visual classification of rocks requires images with a spatial resolution of at better than 0.5 mm/pixel. Nomad's camera 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 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 later proved unnecessary for the 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 angular resolution by the distance from the lens to the rock sample, deduced by assuming the world is flat and projecting a ray from the camera to the sample. 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 and the actual focal length f (mm). Assuming a logarithmic variable resistor as the lens zoom transducer it follows that $f = \alpha e^{-\beta Z}$ and therefore the angular scaling factor Fov is given by $Fov = A e^{\beta Z}$ radians/pixel. The constants A and β were determined by measuring Fov on a set of calibration images taken with different zoom settings and the best fit determined (Figure 2).

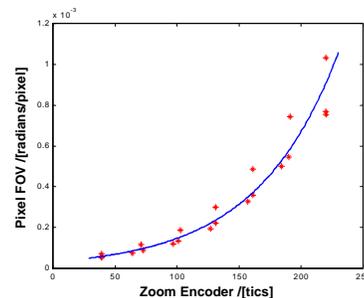


Figure 2 Camera zoom encoder calibration.

Auto-focus software was developed because 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

$$FocusQuality = \sum_{c=C_0 \dots C_n} \sum_{r=R_0 \dots R_m} |Intensity(r, c) - Intensity(r, c+1)|$$

The optimum focus setting is that which maximizes *FocusQuality*. The focus space is non-monotonic so a binary search is inappropriate. A two-step search was used that performed one coarse linear search over the 256 focus encoder values followed by a refined search around the coarse optimum focus setting.

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 standard and in theory independent of the illumination.

Spectrometer

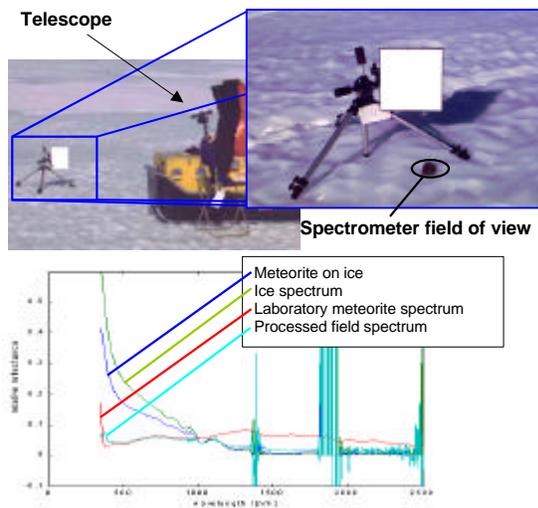


Figure 3 The problems with remotely acquired spectra of rocks on the ice sheet. A telescope is used to collect light for spectral analysis. The signal from the ice overwhelms that from the rock or meteorite sample. If the background ice spectrum is known, it can be subtracted from the combined signal to deduce the rock spectrum with only limited success.

The 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 [6]. However, previous experiments [5] 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.

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 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.

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 (Figure 1) with suitable sensors and software to autonomously place the probe, a difficult task to do with sub-centimeter precision.

A special sensor head (Figure 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.

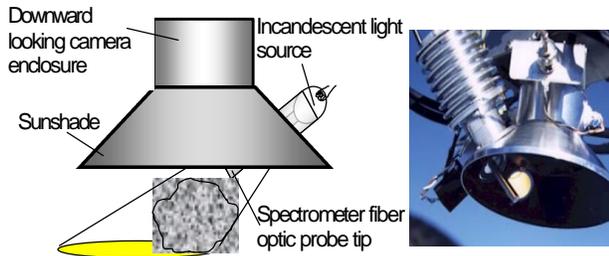


Figure 4 Spectrometer sensor head, deployed by manipulator arm onto rock samples on ice.

Rock and meteorite classifier

Many signal classification techniques exist. However, when classifying rocks from onboard a mobile robot using data from its sensors several issues must be addressed.

- The classifier must learn from examples. Rock classes are ambiguous and human experts cannot always explain how they identified a rock or meteorite. According to a veteran Antarctic meteorite searcher¹ “First you find a few meteorites, then you know what to look for”.
- It is difficult to obtain a large set of example data *under realistic field conditions*. Anticipating the quality of autonomously acquired data is even tougher. Therefore it is desirable to encode any prior knowledge or constraints into the classifier so that it can learn to distinguish rocks and meteorites from a minimum number of examples.
- The classifier must accept incomplete data, and compound evidence for different hypotheses as more becomes available. The robot has multiple sensors, and there is a cost associated with using each one. Sensors such as the spectrometer are expensive to use because the robot must be maneuvered to bring samples into the sensor manipulator workspace. Therefore, it is desirable that initial classifications be made using data from cheap long range sensors, such as a color camera, before final

¹ Professor William Cassidy, University of Pittsburgh Department of Geology.

verification using expensive sensors.

- As corollary of the above, the classifier should accept prior evidence from other sources, such as expert knowledge on what to expect in a particular region.
- The classifier must handle the ambiguity inherent in trying to classify rocks, and indicate several likely hypotheses if a definite classification cannot be achieved.

These requirements argue for a Bayes network based approach to classification. A Bayes network [7] is an acyclic directed graphical model of the statistical relationships between various quantities. Each node represents a random variable, such as the rock type of a particular sample or some observation made on it. Links between nodes represent direct statistical dependencies, with each node conditioned only on its parent nodes. Bayes networks are efficient representations of the joint probability density of all their nodes and allow the computation of the conditional probabilities of any random variables given any of the others. Therefore, from a network encoding the relationships between rock types and sensor observations one can compute the likelihood of a rock being a meteorite given the sensor observations available so far. As more observations become available the new likelihood of a sample being a meteorite can be computed by Bayesian update.

The intuitive graphical structure of a Bayes network makes it easier to encode physical constraints into the network topology, reducing the intrinsic dimensionality. Bayesian update is a principled way to compound evidence, and prior information is naturally represented by prior probabilities. The uncertainty in the classifications is naturally represented by a probability vector rather than a binary classification. Finally, given the Bayes network topology for a set of random variables it is possible to learn the conditional densities between adjacent nodes from example data.

Other classification methods cannot do all of these. Neural networks and nearest neighbor classification can learn from examples and express ambiguity, but not compound evidence. Dempster-Shafer theory is not amenable to learning. Decision trees and predicate logic do not handle ambiguity and are vulnerable to noise.

A Bayes network can also compute the likelihood of any new data given the current knowledge, and thus conceivably recognize bad sensor readings as outliers. Furthermore, the network can be queried to estimate the information gain of further sensor readings, possibly enabling active sensor selection.

Bayes network rock model

Learning the structure of a Bayes network from data is possible but requires substantial data. A network model (Figure 5) for rocks and meteorites was developed by hand based on domain specific knowledge. It is a compromise between a fully connected network with no constraints but computationally intractable, and a naïve Bayes classifier that can be efficiently evaluated but lacks representational power. Sensor features are only weakly directly dependent on each other because of the way they are calculated [5] and because the intermediate node **Rock type**, whose states include all possible rock and meteorite types likely to be encountered by the classifier.

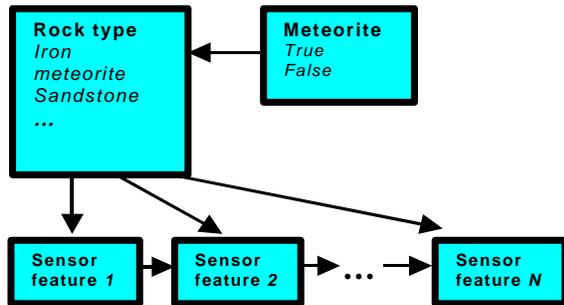


Figure 5 Simplified diagram of Bayes network rock and meteorite classifier.

A complication is that the sensor features are continuous quantities, yet the Bayes network implementation used on Nomad can only handle discrete variables. Therefore the continuous variables need to be suitably quantized.

Network training

The network conditional probabilities are learned from example data using the robust Bayesian learning algorithm in [10]. Because the continuous feature vectors are quantized, there is a trade-off between resolution and having enough example data to populate the conditional probability histograms. To solve this an error model was developed to perturb sample spectra and generate multiple training examples from each one. If $x(\lambda)$ is a sample spectrum then the perturbed spectrum is given by multiplying by a random scale factor Λ , centered at 1, and adding Gaussian noise:

$$x^*(\lambda) = \Lambda x(\lambda) + N(0, \sigma_r(\lambda)), \text{ where } \sigma_r^2 \approx \left(\frac{\Lambda s}{x_{ref}}\right)^2 \left(1 + \frac{1}{x_{dr}^2}\right)$$

and $x_{ref}(\lambda)$ a typical white reference spectrum.

The network was trained using laboratory data², and using 20% test / 80% training cross validation the ROC

² Obtained from laboratory studies of rock and meteorite samples done at Ohio State University and NASA JSC using the Nomad spectrometer and a handheld digital camera.

classification performance Figure 7 is created. This indicates the trade-off between correctly recognizing meteorites and the false positive rate. Each point on the curve corresponds to a particular choice of threshold on the classifier probability output for deciding whether to classify a sample as a meteorite or not. 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.

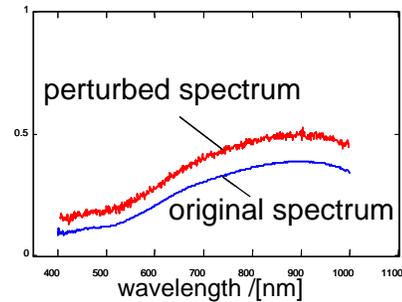


Figure 6 Perturbed sample spectrum for network training purposes. Note slight increase of noise at tail end of spectrum.

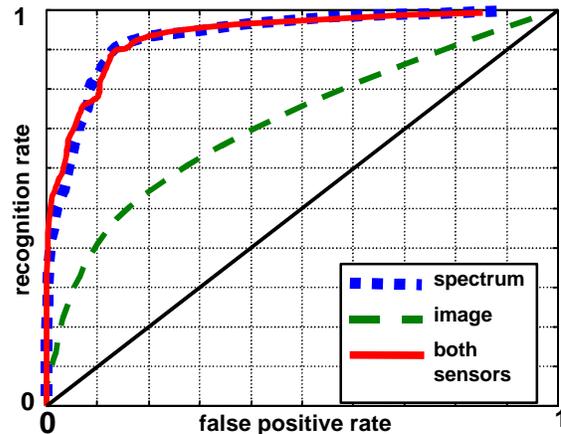


Figure 7 ROC curve of classifier performance recognizing meteorites in laboratory training data, using 20% cross validation.

System operation

Nomad's meteorite searches are carried out as follows: a command is given to search an area designated by differential GPS coordinates. Nomad's mission planner executes a search pattern to exhaustively cover the area while avoiding obstacles.

During the search pattern Nomad processes images from the high-resolution camera to detect new samples using a

threshold-based color segmentation algorithm that distinguishes dark rocks from the generally bright snow and ice. New samples are placed into a database and the mission planner selects the best target to examine using a set of deployment cost parameters, such as distance. Nomad is subsequently maneuvered to take a close-up image that is passed to the classifier for initial assessment. The classifier estimates the information gain of a spectrum from that sample for determining if it is a meteorite. If it exceeds a preset threshold Nomad is commanded to maneuver and deploy the spectrometer. When complete, the arm is stowed and the spectral data passed to the classifier for assessment, after which Nomad resumes the search. The computed likelihood of a rock being a meteorite is passed on to the users, who may decide to collect the sample if the likelihood exceeds some threshold determined by the presumed classifier ROC curve.

Antarctica 2000 field results

Ten meteorite search demonstrations were performed at Elephant Moraine, during which a total area of approximately 2500 m² 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 the SAS autonomously acquired new 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.

The spectrometer data was 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 8). The problems occurred when the probe missed the sample altogether or there was snow on the sample. Furthermore, the contact switch proved

inadequate for autonomously lowering the spectrometer onto targets and so requiring human intervention.

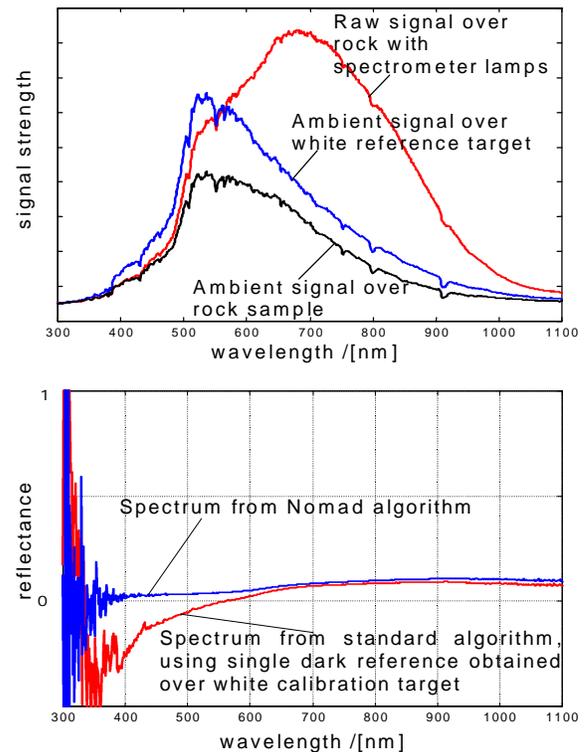


Figure 8 Top: Unprocessed spectral signal from a sample rock with lamps illuminated, and spectral signals with lamps off when probe is over the white reflectance standard and over the rock. Note the strong signal due to ambient sunlight. Bottom: Computed reflectance spectra using Nomad's spectrometer 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.

Overall classification performance (using spectra only) is indicated by the ROC performance curves in Figure 9, generated from a test set of rocks and meteorites (40 and 4 samples respectively, with multiple readings of each) in a particular area of the moraine. Figure 9(i) is using the *a priori* classifier built from the lab data (Figure 7). Performance on the field data clearly does not match that on the lab data (a false positive rate of 20% gives a 25% field recognition rate versus the laboratory recognition rate of 95%). There is a notable improvement in Figure 9(ii), the ROC curve for the same classifier further trained with field data acquired by the robot in the area (from 8 rocks and 2 meteorites *not* in this test set).

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 iron oxide (rust) whose spectrum peaks at 900 nm where many meteorite spectra also have a peak. This is unsurprising, given that most meteorites contain metallic iron, and therefore some surface rust. These rocks were absent from the laboratory training set and not initially expected in this area. Performance is improved when they are excluded from the test set and the retrained classifier is used (Figure 9(iii)).

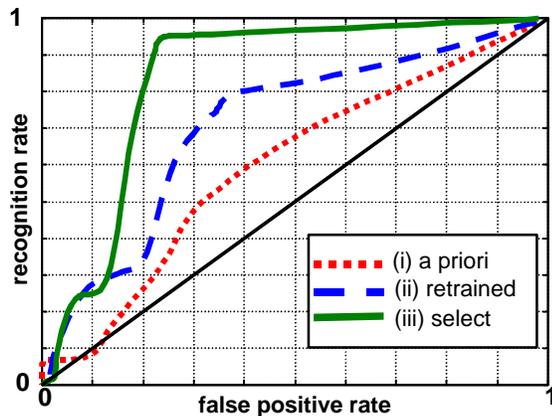


Figure 9 ROC curve demonstrating performance of Nomad meteorite classifier on rocks and meteorites at Elephant Moraine, Antarctica, in January 2000.

Conclusions and Lessons Learned

Classification performance for rocks not stained with iron oxide is reasonable given the quality of the data and the difficulty of the task. However, there is no comparable prior work to calibrate this. Better performance is possible with more sensors such as a metal detector [4]. However, a metal detector or magnetometer sensor is inappropriate for the more general problem of identifying different rock types, not just discriminating meteorites from terrestrial rocks. Furthermore, they would not detect the rarer and more valuable meteorite samples, such as achondrites and SNC's, that do not contain any metallic iron whatsoever.

The significant classification improvement after field retraining 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.

Correct aperture control is an important for rock classification from imagery. 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 is a possibility.

The spectrometer design performed well in the extreme field conditions and was usually tolerant of sensor placement errors due to the autonomous deployment procedure. Occasionally however the probe would be placed such that it did not view a rock at all. In these cases, a clearly erroneous spectrum was measured and the classifier returned absurd results. A sensor to automatically verify the correct placement of the spectrometer probe is essential, as is a modification to the classifier to allow it to recognize nonsense data. Nonetheless, it is reasonable to say that the modifications to the spectrometer were fundamental to the success of the RAMS project, as without them it would not have been possible to autonomously obtain spectra.

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.

A further improvement in classification rates is attainable by autonomously learning the distributions of different rock types over an area and exploiting the statistical dependencies amongst samples, as explained in [9]. This has proven effective in dramatically reducing error rates even given ambiguous sensor data.

Acknowledgements

Special thanks are due to William Cassidy of University of Pittsburgh, Gunter Faure of Ohio State University, Marilyn Lindstrom at NASA Johnson Space Center, Ted Roush of NASA Ames Research Center, and Martial Hebert of Carnegie Mellon University for their generous support, geological expertise and access to rock and meteorite specimens and equipment.

The RAMS project was funded by a grant from the National Aeronautics and Space Administration, and supported in Antarctica by the NSF Office of Polar Programs and the ANSMET program.

References

- [1] D. Apostolopoulos, M. Wagner, W. Whittaker, "Technology and Field Demonstration Results in the Robotic Search for Antarctic Meteorites", *Field and Service Robotics Conference*, Pittsburgh, USA, 1999
- [2] Dimitrios Apostolopoulos, Liam Pedersen, Ben Shamah, Kimberly Shillcutt, Michael D. Wagner, William "Red" Whittaker, "Robotic Antarctic Meteorite Search: Outcomes", in *IEEE International Conference on Robotics and Automation*, Seoul, Korea, 2001.
- [3] T. Roush, et al, "Autonomous Science Decision Making for Mars Sample Return", in *Lunar and Planetary Science Conference XXX*, Houston, TX, 1999.
- [4] L. Pedersen, "Robotic Deployment of Electro-Magnetic Sensors for Meteorite Search", *IEEE International Conference on Robotics and Automation*, Leuven, Belgium, 1998.
- [5] L. Pedersen, D. Apostolopoulos, W. Whittaker, T. Roush, G. Benedix, "Sensing and Data Classification for Robotic Meteorite Search", *Proceedings of SPIE Photonics East Conference*, Boston, 1998.
- [6] G. Hunt, "Spectral Signatures of Particulate Minerals in the Visible and Near Infrared", *Geophysics*, vol. 42, no.3, p. 501-513, April 1977.
- [7] J. Pearl, *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*, Morgan Kaufman, San Mateo, CA. 2nd edition 1991.
- [8] M. Wagner, et al, "The Science Autonomy System of the Nomad Robot", in *International Conference on Robotics and Automation*, Seoul, Korea, 2000
- [9] Liam Pedersen, "Autonomous Characterization of Unknown Environments", in *IEEE International Conference on Robotics and Automation*, Seoul, Korea, 2001
- [10] D.J. Spiegelhalter, D.A. Philip, S.L. Lauritzen and R.G. Cowell, "Bayesian analysis in expert systems" in *Statistical Science*, 8(3), p219-283., 1993.