

Creation and Testing of an Artificial Neural Network Based Carbonate Detector for Mars Rovers

Benjamin Bornstein, Rebecca Castaño
Jet Propulsion Laboratory
California Institute of Technology
4800 Oak Grove Drive
Pasadena, California 91109
818-393-5349, 818-393-5344
{ben.bornstein, rebecca.castano}@jpl.nasa.gov

Martha S. Gilmore, Matthew Merrill, James P. Greenwood
Dept. of Earth and Environmental Sciences
Wesleyan University
265 Church St.
Middletown, Connecticut 06459
860-685-3129, 860-685-4477, 860-685-2545
{mgilmore, mmerrill, jgreenwood}@wesleyan.edu

Abstract—We have developed an artificial neural network (ANN) based carbonate detector capable of running on current and future rover hardware. The detector can identify calcite in visible/NIR (350–2500 nm) spectra of both laboratory specimens covered by ferric dust and rocks in Mars analogue field environments. The ANN was trained using the Backpropagation algorithm with sigmoid activation neurons. For the training dataset, we chose nine carbonate and eight non-carbonate representative mineral spectra from the USGS spectral library. Using these spectra as seeds, we generated 10,000 variants with up to 2% Gaussian noise in each reflectance measurement. We cross-validated several ANN architectures, training on 9,900 spectra and testing on the remaining 100. The best performing ANN correctly detected, with perfect accuracy, the presence (or absence) of carbonate in spectral data taken on field samples from the Mojave desert and clean, pure marbles from CT. Sensitivity experiments with JSC Mars-1 simulatant dust suggest the carbonate detector would perform well in aeolian Martian environments^{1,2}.

geologically relevant information and prioritizing its downlink have the potential to dramatically increase the science return of future rover missions. Carbonate minerals are of particular geologic importance because they signal the presence of water and can be associated with life on Earth. For these reasons, we have developed an artificial neural network (ANN) based carbonate detector capable of running on current and future rover hardware.

The carbonate detector was developed using a supervised learning method. A brief introduction to neural networks is provided in Section 2. In Section 3, we describe the source of data used to train the detector. We then describe the development of the neural network based carbonate detector in Section 4. In Section 5, two tests on spectral measurements of representative samples are described. One data set is of a field site in Silver Lake, California and the other data set consists of measurements of a calcite (carbonate) sample covered with layers of increasing thickness of dust. A key aspect of the selection of a neural network classifier for the detector is the low computational requirements to analyze a measurement. In Section 6, we detail the computational requirements of the detector and explain how it might be used on a rover as part of an onboard science data analysis system for analyzing data during rover traverses. Finally, in Sections 7 and 8, we draw conclusions and offer directions for future work.

TABLE OF CONTENTS

1. INTRODUCTION.....	1
2. NEURAL NETWORKS	1
3. GENERATIVE MODEL.....	2
4. DETECTOR CONSTRUCTION.....	3
5. DETECTOR TESTING.....	4
6. FLIGHT READINESS.....	5
7. CONCLUSIONS	6
8. FUTURE DIRECTIONS	6
9. ACKNOWLEDGEMENTS	6
10. REFERENCES	6
11. BIOGRAPHIES	7

2. NEURAL NETWORKS

In this section, we give a brief overview of artificial neural networks, beginning with their biologically inspired roots. Our goal is to give readers unfamiliar with ANNs an intuitive understanding of what they are and how they work. As such, we omit any formal mathematical descriptions of ANNs. Readers familiar with the topic may wish to skip to the next section.

1. INTRODUCTION

Data collection capabilities of Mars rovers far surpass downlink rates, which can result in a critical loss of valuable science data. Onboard techniques for identifying

Artificial neural networks are a class of machine learning techniques loosely based on theories of how biological neural networks store and retrieve information, or more generally, compute. At a coarse level, a biological neural network is comprised of neurons and the synapses that connect them. Synapses carry electrochemical signals to

¹ 0-7803-8870-4/05/\$20.00© 2005 IEEE

² IEEEAC paper #1527, Version 7, Updated December 27, 2004

and from neurons. Synapses feeding into neurons are called *dendrites*, while those fanning out from neurons are *axons*. A stimulus received at a dendrite is either amplified or diminished as it is propagated to the connected neuron. Since a neuron may have multiple dendrites all providing stimuli, all dendrite signals taken together form the neuron’s *total input*. If the total input signal exceeds some threshold the neuron is *excited* and, as a result, propagates its electrochemical response to neighboring neurons via its axon. Otherwise, the neuron’s response is *inhibited*. In general, the greater the total input, the greater the neuron’s response. The information contained in a neural network is embodied in its collection of synapses. Finally, *plasticity* is the ability of synapses to change the degree to which they amplify or diminish stimuli over time, based on the frequency and strength of the stimuli; it is plasticity that allows a neural network to both learn and forget.

An artificial neural network is similar to its biological counterpart. Stimuli and synapses are modeled as real numbers. In the case of synapses, these numbers, often called *weights*, represent the degree to which the synapse amplifies (positive value) or diminishes (negative value) the input stimuli. A signal propagated along a synapse is simply the value of the stimuli multiplied by the synapse’s weight value. The total input to a neuron is the sum of all incoming propagated signals. Neurons are modeled as functions that map the total input stimulus to a real number response. The function’s output is the degree to which a neuron is excited or inhibited. Plasticity is modeled as a specific algorithm (e.g. Backpropagation [9]) for tuning synapse weight values to produce the desired responses for a set of input stimuli.

Training an artificial neural network proceeds as follows: We present the network with a pattern to learn, usually represented as a set of input stimuli. The network propagates the pattern through its synapses and neurons until it produces a final output pattern or value. The training algorithm compares the output to an expected output and if they do not match, the algorithm slightly adjusts the network’s synaptic weights so that if we again presented the pattern to the network, the network’s output would be closer to the expected output. It’s important the synapse weight adjustments do not produce exactly the expected output. Otherwise, the network learns to identify only this pattern, whereas the typical goal is to have a network learn to recognize many patterns and further generalize to new, unseen (yet similar) patterns. We repeat this pattern presentation and weight adjustment process with new input patterns, over and over again, cycling through the training patterns, until the network produces acceptable outputs for all inputs. To bring this discussion full circle, note that this process is very similar to how humans learn through many (similar) repetitions.

3. GENERATIVE MODEL

A potential drawback of ANNs is that large numbers of training examples are often required for the network to converge (learn). Our initial experiments showed that tens of spectra were insufficient to train an ANN to predict the presence or absence of mineral end-members. This presented a problem since collecting only a few dozen samples often requires days of fieldwork followed by lab work to verify the modal mineralogy of each sample. While it is possible to use spectral libraries to alleviate the labor burden of sample collection and analysis, for our purposes, spectral libraries have two drawbacks: i) they often contain ten or fewer spectra per sample and ii) the spectrum is usually measured in the laboratory, which yields data cleaner than that collected in the field. For an ANN detector to perform well in field, it must be trained on field data or something that closely resembles it.

In an effort to inexpensively obtain spectral data with many of the characteristics of field data, we have begun to develop a *generative model* for spectra. The model starts with a database populated from two well-know spectral libraries: JPL’s Advanced Spaceborne Thermal Emission and Reflection Radiometer (ASTER) library [7] and the U.S. Geological Survey’s (USGS) speclib04 [2]. To this database we add *virtual rocks*, described in terms of either absolute mineral constituent percentages or percentage ranges. For data generation, all percentages are chosen using a uniform pseudorandom number generator.

In addition to a percentage (range), for each virtual rock description, we classify each mineral as *essential*, *accessory*, or *accidental*. Essential minerals are required to occur in a rock at the percentage (or within the range) specified. Accessory minerals are added to a rock if the chosen percentages for essential minerals do not sum to 100. Finally, accidental minerals are added to rocks with low probability (usually less than 10% of the rocks generated). The compositions of the virtual rocks were selected to represent rocks known (from meteorite specimens and orbital spectroscopy) and predicted to occur on Mars. This includes basic volcanic rocks (i.e., basalts) and dust that comprise the majority of the surface as well as minerals of interest (minerals associated with water such as carbonates and evaporates). The classification of minerals as *essential*, *accessory*, or *accidental* and their percentages are constrained to be geologically reasonable.

To generate spectra for each virtual rock, we take the rock’s constituent mineral spectra and corresponding percentages and apply a *mixing model*. The simplest is a linear mixing model. Let $r(m, b)$ be the reflectance for end-member m at bandwidth b . Then, given the mix percentage for each end-member w_m , the mixed reflectance for band b , R_b , is:

$$R_b = \sum_m w_m r(m, b) \tag{1}$$

Thus, under the linear mixing model, mixed reflectances are simply weighted, linear combinations of end-member spectra. Recall, however, the spectra are drawn from spectral libraries, which contain predominately laboratory spectra. To simulate instrument and other noise encountered when taking field measurements, we add pseudorandom Gaussian noise with mean zero and variance σ :

$$R_b = \sum_m w_m r(m,b) + N[0,\sigma] \quad (2)$$

While such a mixing model is simplistic, it has allowed us to setup and test our generative model framework and it can provide a wealth of spectra with many subtle variations to train ANNs and other machine learning techniques. We are currently developing a set of richer, nonlinear mixing models based on the reflectance and refraction models of Hapke [6].

4. DETECTOR CONSTRUCTION

We used our generative model to create 10,000 spectra to train several ANNs (see Figure 1). The spectra were based on nine carbonate and eight non-carbonate minerals (see Table 1). The carbonate set includes those carbonates that are most abundant on the Earth and those for which library spectra were available in the ASTER spectral library [7]. The non-carbonate set includes the major rock-forming minerals of igneous rocks on the Earth and the martian meteorites [e.g., 8], and one mineral (clinocllore) produced during the interaction of mafic igneous rocks with water in a hydrothermal environment. For the sake of simplicity, in our initial experiments we turned off mineral mixing and instead added only up to 2% Gaussian noise to each reflectance measurement. Initially we generated spectral data at the full instrument resolution of 2,151 measurements (0.35 to 2.5 μm in .001 μm increments), but after visual inspection we found we could reduce the number of network inputs to 215 by averaging every 10 μm without appreciable information loss. Further, we were concerned about noise at 1.4 and 1.9 μm due to atmospheric water. To avoid these bands altogether and focus on the absorption regions most characteristic of carbonates, we eliminated all bands except the 2.0 to 2.4 μm region. Averaging every 10 μm reflectance values in this region yielded a more manageable 41 input values. This significantly reduced our network training times, allowing us to test several additional ANN architectures and parameters.

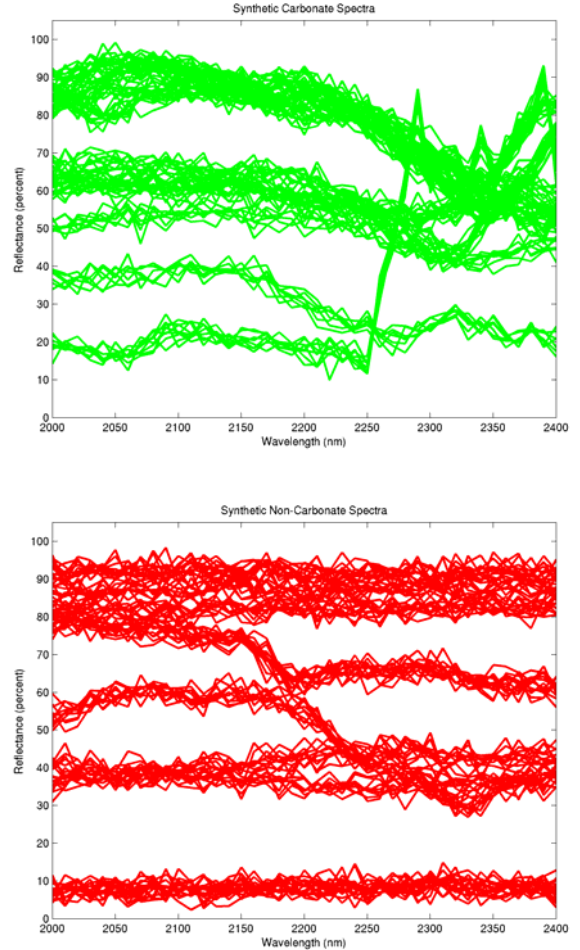


Figure 1 – A random sample of 100 carbonate (top) and 100 non-carbonate (bottom) spectra (out of 10,000) used to train the carbonate detector. We created these “synthetic” spectra by applying our generative model to ASTER spectra for the minerals listed in Table 1.

Table 1: Mineral spectra used for ANN training

Non-carbonates	Carbonates
Hematite	Cerussite
Magnetite	Strontianite
Clinocllore	Witherite
Biotite	Calcite
Albite	Rhodochrosite
Quartz	Siderite
Augite	Dolomite
Orthoclase	Azurite
	Malachite

To search the space of likely network architectures and parameters, we took a cross-validation approach. First, we divided the 10,000 spectra into a 9,900 sample training set that we used to train each network. The remaining 100 samples were used to test the accuracy of each network after training was complete.

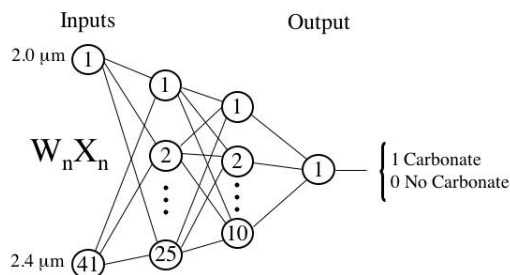


Figure 2 – The neural network architecture of the carbonate detector. The network has an input layer, two hidden layers, and an output layer. Input (X_n) are reflectance measurements taken from 2000 to 2400 nm in 10 nm increments. Output is a single real value in the range [0, 1], with a threshold set at 0.3 to produce a binary result: carbonate ($\geq .3$) or no carbonate ($< .3$). Weights (W_n) are determined via the standard Backpropagation algorithm.

After some trial and error, we discovered a robust carbonate detector with 41 inputs and three successive layers with 25, 10, and one, neuron(s) respectively (see Figure 2). The single neuron layer is the detector output and reports values between zero and one. All detectors were trained with the standard Backpropagation algorithm and sigmoidal activation units. For the training dataset, all carbonates were assigned an output value of one and all non-carbonates a value of zero. The nature of Backpropagation and the sigmoid transfer functions are such that the output unit rarely produces values at these two extrema. Instead values closer to zero indicate the absence of carbonate and values closer to one indicate the presence of carbonate. To choose the best detection threshold, we examined how well the detector was able to separate spectral data taken during our Silver Lake field test (see next section). We found the carbonate detector was clearly able to separate multiple spectra, taken from several distances (two, four, six, and eight meters) and directions (left, right, and center), of a limestone rock from data taken from a granite, gneiss, and metagabbro. We varied the detection threshold from zero to one in .001 increments, and calculated the percentage of limestone spectra correctly recognized as containing carbonate. This percentage is the true positive percentage. We also calculated the true negative percentage, i.e. as the same threshold varies from zero to one, the percentage of the other rocks correctly identified as not containing carbonate. A plot of these two percentages is shown in Figure 3. From this plot, there is a range of thresholds from 0.2 to 0.3 that minimize the number of false positives and maximize the number of true positives. We selected 0.3 for future tests to conservatively minimize false positives for datasets.

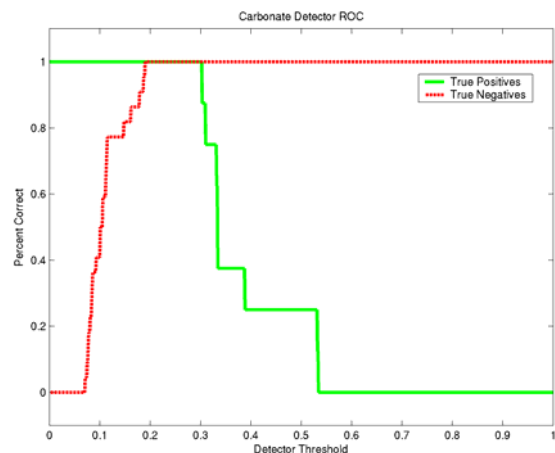


Figure 3 – Carbonate detector receiver operator characteristic curve for the Silver Lake Field data. The solid (green) line shows the percentage of true positives as a function of a threshold on the detectors output unit. The dashed (red) line shows the percentage of true negatives. The ideal threshold for the detector, accounting for all true positives and minimizing the number of false positives, is 0.3.

5. DETECTOR TESTING

After training, the ANN classifier was tested on data collected in the field at Silver Lake, CA. Silver Lake Playa is approximately 12 by 3 km and lies 3 km north of Baker, CA. Silver Lake was used as one of the test sites for the field prototype Mars rover, FIDO (Field Integration Data and Operations). The playa's long-term arid environment, lacustrine features and sediments make it a geological analog for sites on Mars which may contain evidence of water.

Spectrometer measurements were taken using a FieldSpecFR (Analytical Spectral Devices, Inc., ASD) fiberoptic spectrometer operating over the 350-2500 nm wavelength range. This wavelength range includes distinctive spectral features for carbonates, Fe-bearing minerals, and clay minerals. Measurements were taken over a range of distances and azimuths in order to accommodate rock heterogeneity and variations in lighting. The test data consisted of 30 field spectra of four different types of rocks. In our experiments on this limited test set, the neural net correctly identified the presence or absence of carbonate minerals in all cases (see Figure 4) [4].

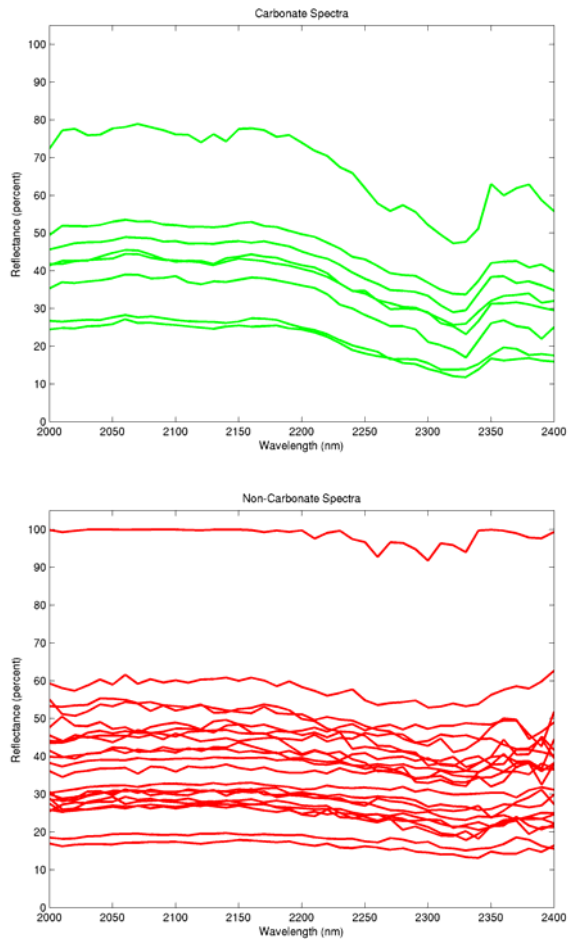


Figure 4 – Results of carbonate detector on the 2000–2400 nm region of 30 field spectra collected at the Silver Lake site. The absorption at ~2330 nm present in each of the carbonate spectra (top) allowed them to be successfully distinguished from non-carbonate (bottom) by the detector.

A second field test was performed on dolomitic marbles in northwest CT. Spectra were collected of fresh (weathered surface removed using a hammer) and weathered surfaces; 399 spectra were collected of 30 targets at three outcrops. In our preliminary analysis, we find the detector correctly identified carbonate in the spectra of 6 targets that were fresh and dominated by dolomite. The detector failed to detect carbonate in weathered surfaces of the dolomite-bearing rocks with one exception being a rock with a sugary texture prone to removal and thus exposing fresh crystals. Spectra that passed the detector had strong carbonate absorptions at both 2160 nm and 2340 nm. Weathering of the dolomite to clay minerals reduces these absorptions and making the spectra unrecognizable to the detector. Such coatings may limit the ability of a rover-based detector to recognize primary mineralogy remotely.

To examine this issue further, we sought to test the detector in more realistic Mars-like environments. Of particular

concern is the ability to discern primary rock mineralogy through the ferric dust detected on the surface of Mars by earth-based, orbital and landed instruments. Martian dust coatings strongly affect the spectral characteristics of surface rocks potentially masking the underlying substrate rock. To test the limits of the performance of the carbonate detector, we performed a series of experiments to simulate the deposition of Mars-analogue ferric dust onto a calcite substrate [5]. The detector correctly recognizes a calcite crystal beneath dust layers of up to ~100 μ m and 80% coverage. Dust layers of this magnitude may be expected on rocks within regions of aeolian activity and dust mobilization. The carbonate detector should accurately recognize coarse-grained calcite in dusty rocks in such environments.

6. FLIGHT READINESS

The detector consumes minimal CPU time and storage and thus is ideal for use in onboard rovers during long traverses.

A single run of the detector requires 1285 floating-point multiplication operations and also 1285 floating-point addition operations. While the sigmoidal activation function can be more costly to compute, it is often feasible to replace function calculation with two table lookups followed by a linear interpolation. In total each detector run requires no more than 3,000 operations. This is well within the capability of current flight microprocessors. We realize that often such processors do not contain math coprocessor capable of performing floating-point operations. To run on such processors, the floating-point multiplications and additions would need to be converted to fixed-point. The detector itself is written in 100% pure ANSI C and has been compiled on Linux, Windows, and MacOS X with no source code changes required. The source code and network weight file consume less than 100 KB of storage.

The carbonate detector can be combined as a part of a larger integrated on board science analysis system. Such a system called the Onboard Autonomous Science Investigation System (OASIS) is currently under development [1]. In this system, data is analyzed during a rover traverse for interesting scientific targets. Carbonate minerals are one form of interesting target and the carbonate detector would provide a beneficial form of feature extraction. The carbonate detector could be used in several ways. One mode of operation would be to use a fixed sampling strategy during a traverse where samples are collected either by targeting rocks or just pointing in a fixed direction. This data would be analyzed for key signatures. Upon identification of a carbonate, the data would be saved along with an image of the target to transmit at a downlink opportunity. Thus, the data would be prioritized for downlink, where data indicating detected targets is given highest priority. This allows the rover to collect at the maximum rate that the instrument and processing allow

rather than limiting data collection based on downlink bandwidth. Another mode of operation that the carbonate detector may be used is in reaction to analysis of data from another instrument. In this mode, engineering data collected for navigation (e.g. navigation and hazard camera images) would be analyzed for candidate targets of interest. Spectral measurements of identified candidate targets would be collected and analyzed to determine if the candidate has interesting composition as hypothesized from the engineering data. In each mode of operation, the analysis results can be used for prioritizing data for downlink or as a science alert for the rover to take action in response to the detection.

7. CONCLUSIONS

We have developed an ANN-based carbonate detector that correctly identifies carbonate minerals in laboratory, field and Mars-analogue settings. The success of the neural net classifier at identification of carbonates in field spectra demonstrated the potential ability of this technique to autonomously identify critical mineralogies on Mars.

8. FUTURE DIRECTIONS

There are several areas in which further work needs to be done to have a system of mineral detectors suitable for use for onboard a rover. The first area that we are exploring is the use of a Support Vector Machine (SVM) classifier. Neural networks were initially selected over SVM's largely due to the significantly less computational requirements for analysis. Recent advances in SVM methods have now led to a reduction in computational requirements [3]. Although, in general, still greater than neural nets, SVM's may be a competitive alternative and have demonstrated higher accuracy in many other application areas. One of the benefits of SVM's is that fewer training examples are needed in order to get a high quality classifier. For example, an SVM classifier may use 500 training examples versus 10,000 training examples for a neural net.

A second area that we are exploring is detectors for additional mineral classes. Finally, we recognize that while we have modeled the minerals as linearly mixing together in forming rocks, more often, rocks are intimate mixtures where the reflected spectra is not a linear combination of the constituent minerals weighted in proportion to the percentage of each mineral present. We are developing nonlinear models that more closely represent the observed mixing behavior.

9. ACKNOWLEDGEMENTS

This work was sponsored by NASA's Applied Information Systems Research Program. We thank Ted Roush, Ed Guinness, and Brian Ebel for field instrumentation and assistance in the collection of the spectra from Silver Lake. Richard Maclin contributed to the development of the code used for the creation of the neural network training data. In part, this work was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under contract with the National Aeronautics and Space Administration.

10. REFERENCES

- [1] Castaño, R., Judd, M., Estlin, T., Anderson, R. C., Scharenbroich, L., Song, L., Gaines, D., Fisher, F., Mazzoni, D., Castaño, A., "Autonomous Onboard Traverse Science System," *Proceedings of the 2004 IEEE Aerospace Conference*, Big Sky, Montana, March 2004.
- [2] Clark, R.N., Swayze, G.A., Gallagher, A.J., King, T.V.V., Calvin, W.M., "The U. S. Geological Survey, Digital Spectral Library: Version 1: 0.2 to 3.0 microns," *U.S. Geological Survey Open File Report 93-592*, 1993, <http://speclab.cr.usgs.gov>.
- [3] DeCoste, D., Mazzoni, D., "Fast Query-Optimized Kernel Machine Classification Via Incremental Approximate Nearest Support Vectors," *International Conference on Machine Learning (ICML)*, August 2003.
- [4] Gilmore, M. S., Castaño, R., Mann, T., Anderson, R. C., Mjolsness, E. D., Manduchi, R., Saunders, S., "Strategies for autonomous rovers at Mars," *J. Geophys. Res.* 105, 29,223–29,237, 2000.
- [5] Gilmore, M. S., Merrill, M. D., Castaño, R., Bornstein, B., Greenwood, J. P., "Effect of palagonite dust deposition on the automated detection of calcite in visible/near-infrared spectra," *ICARUS* 172, 641–646, 2004.
- [6] Hapke, B., *Theory of Reflectance and Emittance Spectroscopy*, Cambridge Univ. Press, New York, 1993.
- [7] Hooke, S. The JPL ASTER Spectral Library v1.1, 2000, <http://speclib.jpl.nasa.gov>.
- [8] McSween, H. Y., "What we have learned about Mars from SNC meteorites," *Meteoritics* 29, 757–779, 1994.

[9] McClelland, J.L., Rumelhart, D.E. Explorations in parallel distributed processing. MIT Press, Cambridge, MA, 1998.

11. BIOGRAPHIES



Ben Bornstein is a research programmer in the Machine Learning Systems group at the Jet Propulsion Laboratory in Pasadena, CA. He enjoys bringing machine learning techniques and considerable hacking (programming) skills to bear to solve problems in geology, bioinformatics, and systems biology. He has designed and implemented software systems for several Caltech biology labs, the Institute for Genomics and Bioinformatics (IGB) at UC Irvine, USC Children's Hospital, and JPL's Mars Exploration Rover (MER) project. He is also the inventor of and lead developer for LIBSBML, an open-source library for the Systems Biology Markup Language (SBML). Ben received a B.Sc. in Computer Science from the University of Minnesota Duluth in 1999 and is currently pursuing a M.Sc. in Computer Science at the University of Southern California.



Rebecca Castaño Supervisor, Machine Learning Systems (MLS) group at JPL, OASIS Technical team lead. She received her Ph.D. in Electrical Engineering from the University of Illinois with her dissertation in the area of computer vision. Dr. Castaño has been advancing the state of the art in onboard science analysis methods for the past five years and has been lead author on numerous publications in this field. From 1999 - 2001, Dr. Castaño served as the application lead for phenomenological computational field geology efforts in the MLS Group. She is currently the technology lead for science data processing for the Autonomous Sciencecraft Experiment on the New Millennium Program's ST6 project. Dr. Castaño is also the Team Lead of the Onboard Autonomous Science Investigation System (OASIS) project and the subtask lead for Data Analysis on the automated Multi-rover Integrated Science Understanding System (MISUS). Her research interests include machine learning, computer vision and pattern recognition.



Martha Gilmore is an Assistant Professor of Earth and Environmental Sciences at Wesleyan University. Prior to this she was a postdoctoral researcher in the Earth and Planetary Science Division at the Jet Propulsion Laboratory. Her primary research involves the use of images of the surfaces of Mars and Venus to interpret geologic process and history. Dr. Gilmore is also involved in the future of planetary exploration including strategies for the return of samples from Venus and Mars, the exploration of Venus with balloons, and the use of artificial intelligence to improve the geological capabilities of Mars rovers. She previously served as member of the Mars Surveyor 2001 lander team and has contributed to site selection for the Mars Exploration Rovers. Dr. Gilmore is a member of the Space Studies Board's Committee for Planetary and Lunar Exploration. She received her Ph.D. in Geological Sciences from Brown University.



Matthew Merrill M.A. student, Wesleyan University. He received his B.A. in Earth and Environmental Sciences at Wesleyan University in 2002.



James Greenwood is a Research Assistant Professor of Earth and Environmental Sciences at Wesleyan University. Dr. Greenwood is an expert on petrology of martian and chondritic meteorites, with extensive experience using electron and ion microprobes. Recent work includes study of the isotopic signatures of microbes and their detectability at Mars. He was a member of the Mars Pathfinder mineralogy team and performs research on the nature and chemistry of martian soils. He received his Ph.D. in Geological Sciences from Brown University.