

Technology and Field Demonstration Results in the Robotic Search for Antarctic Meteorites

Dimitrios Apostolopoulos, Michael Wagner, William “Red” Whittaker
E-mail: {davl, mwagner, red}@frc.ri.cmu.edu

Field Robotics Center, The Robotics Institute
Carnegie Mellon University
Pittsburgh, PA 15213 USA
phone: (412) 268-7224; fax: (412) 268-1488

ABSTRACT: *Robotic search for meteorites in Antarctica is an ideal test case for demonstration and field validation of planetary science rovers. Antarctica’s lengthy diurnal cycle, its harshness and its remote location present conditions and challenges similar to those encountered in missions to the poles of the Moon and Mars. This project has researched and developed the technologies and capabilities of an autonomous robot for the search of Antarctic meteorites. Nomad, a robot that explored 220 km of the Atacama desert in 1997, was winterized and outfitted with sensors and onboard intelligence for detection and in situ classification of rocks and meteorites. Nomad’s autonomous perception and navigation capabilities are appropriate for excursions in polar environments.*

This article first introduces the science and search for Antarctic meteorites and overviews Nomad’s robotic technologies. It then details Science Autonomy, Nomad’s control architecture and its functionality. The remainder of the article discusses Nomad’s performance in a meteorite search as it traversed ice terrains and endured harsh conditions in Patriot Hills, western Antarctica, during a six week expedition in late fall of 1998.

1. Overview

Meteorites provide the only significant source of geological material from other planets and asteroids. They are also an invaluable source of information in our quest to learn about the formation and evolution of the solar system and the origins of life. Meteorites fall randomly to the Earth, yet unique terrestrial regions contain substantial concentrations. The premier example is Antarctica, in particular the eastern regions of the Transantarctic Mountains. Large numbers of Antarctic meteorites drift with the ice flow and concentrate on *stranding surfaces*. These are areas in which the ice flow is blocked by mountains and ice deflation due to wind ablation and sublimation, which exposes meteorites and rocks. Antarctica’s cold, pristine environment—as well as its low weathering rates—contributes to the preservation of meteorites.

Human searches for Antarctic meteorites are based on a sound understanding of probable sites. Using satellite and aerial imaging, scientists of NSF’s Antarctic Search for Meteorites program (ANSMET) target bedrock that blocks ice drifting from the central to the coastal regions and ice

stranding surfaces. In the last 30 years, ANSMET and others have collected more than 20,000 Antarctic meteorite samples. A very small fraction of these have been identified as Martian or lunar in origin and are of exceptional scientific value. Searches under ANSMET are systematic and are largely confined to visual inspection of ice surfaces and moraines (blue ice fields containing large collections of rocks). Scientists identify meteorites, mark their locations with flags, map reference areas and carefully collect the meteorites. Critical to this process is absolute protection against contamination; the meteorites are preserved in their frozen state and are shipped to NASA’s Johnson Space Center for analysis and classification. The collection success rate of ANSMET teams is very high; 96 out of 100 rocks identified in the field as meteorites are proved to be so. Antarctic meteorites vary in size and composition. Typically 2-cm to 15-cm across, they can measure as much as 1 m in diameter and can weigh tens of kilograms. Generally, their shape is flat elliptical or near spherical with extruded features and a glossy brownish-to-black color attributed to a fusion crust. Their weight depends on their elemental composition. Common elements found in meteorites are iron and magnesium mixed with silicates. Numerous nonmetallic meteorites have also been found.

Searching for meteorites in Antarctica is an unprecedented and daunting task for robots. Without preliminary evaluation of sites, vast areas must be searched. To make the search a realistic challenge to a robot, stranding surfaces should be the first to be searched because they tend to have a much higher meteorite density. Additionally a systematic search approach can be taken that involves intelligent planning of sensor deployment and navigation to maximize the search area while meeting time and power requirements. To address the issues of autonomous search for Antarctic meteorites, this project winterized Nomad. This robot was originally developed for long distance exploration of barren terrains. This project augmented its sensing and intelligent capabilities to produce a robust combination of autonomous exploration of remote environments with efficient search and classification of meteorite samples.

In late fall of 1998, Nomad sought out and classified meteorites in the Patriot Hills region of Antarctica. Earlier, in January 1998, a scouting expedition to Antarctica field validated individual robotic components and meteorite detection sensors for the purpose of quantifying their performance before incorporating them aboard Nomad. The goals for Nomad's expedition were the robotic search for and classification of rocks and meteorites as well as autonomous navigation of polar terrain. The primary demonstration evaluated the robot's classifier algorithms to distinguish and characterize rocks and meteorites using high resolution imagery and reflection spectroscopy. The second parallel demonstration conducted a field evaluation of autonomous capabilities including navigation of ice and snow using stereo vision and laser perception as well as area coverage planning. Nomad correctly classified 10 of 15 planted meteorites and 38 of 42 native rocks, and accomplished 10 km of autonomous navigation. The Antarctic excursion provided invaluable technical and logistical lessons. The following sections profile Nomad's key navigation and meteorite search capabilities, and summarize the experiments and results from the robot's excursion at Patriot Hills.

2. Nomad: Robotic Antarctic Explorer

Nomad is a planetary prototype robot with mechanical, sensing and navigation capabilities appropriate for autonomous search of Antarctic meteorites (Figure 1). The four-wheel rover underwent extensive modifications and upgrades over its original design used in the Atacama desert to meet the performance requirements of a robust winterized rover capable of prolonged excursions in Antarctica [1][2]. Mechatronic modifications included the use of pneumatic studded tires for improved traction on ice and snow, utilization of materials, bearings, electronics, connectors, and lubricants designed for low-temperatures, and custom-made casings and seals for terrain and science sensors. Nomad's computing enclosure as well as all sensor casings was equipped with heaters and temperature monitors for thermal control. Nomad's computing was augmented to include a specialized processor for onboard analysis and classification of science data and high level planning. Dedicated processors for navigation autonomy and real-time control were upgraded and all computing was unified under the Linux operating system.

Nomad uses velocity feedback to control its propulsion and position feedback to control its electronically coordinated steering. The robot, which operates entirely with electricity, is powered by an internal combustion engine generator and is capable of eight hours of operations between refueling. Stereo vision and laser rangefinding are the sensing modalities for obstacle avoidance and navigation. A high resolution CCD camera and a spectrometer are the primary sensors for scientific data acquisition. A metal detector and magnetometer have also

been used to collect scientific data. Communication is performed via wireless ethernet while maintaining line of sight with the base camp.



Figure 1: Nomad in the midst of a snowstorm in Antarctica. During its expedition to Patriot Hills, the robot endured wind chills as low as -70°C and winds gusts up to 75 km/hr.

3. Science Autonomy: An Architecture for Robotic Meteorite Search in Antarctica

Science autonomy is Nomad's control architecture that enables the robot to execute a command to search a specified area, avoid any obstacles encountered and to perform an *in situ* classification of potential meteorites found. The structure of science autonomy can be described as a three-level hierarchy (Figure 2). The lowest level is that of the *physical sensor*. These sensors allow the science autonomy system to observe the world. However science autonomy has two operational modes: acquisition and identification. In acquisition mode, science autonomy attempts to find new targets in the area around the robot. In identification mode, study and classification of a target takes place. To carry out both objectives, the physical sensors may have to function in very different ways or may possibly even have to coordinate with other sensors. Therefore the *action coordination* level of the system organizes various physical sensors in different ways to accomplish the goals of the next level: *planning*. The planning level examines the results generated by the action-coordination level and creates a plan that will optimize mission variables such as power and time limitations. [3]

The *mission planner* coordinates all of these activities. It selects an optimal area coverage pattern to exhaustively search an ice field. Throughout the search, Nomad's autonomous navigation system operated by the *navigation manager* serves two purposes. First, it seeks to keep the robot away from obstacles. It does this by integrating terrain data from stereo cameras and a laser rangefinder as well as by monitoring inertial sensors for dangerous roll or pitch values. Second, it implements waypoint navigation,

which enables the robot to maneuver within a particular tolerance to a set of specified goal locations. Waypoint navigation receives a set of Differential Global Position System (DGPS) coordinates from a remote human operator and directs Nomad in straight lines between the coordinates. The module attempts to bring the robot within an error radius of each way point. Because polar weather such as blowing snow can limit the effectiveness of terrain sensors, an error recovery module monitors the status of the robot for problems not noticed by the terrain sensors. In its current form, it detects two kinds of problems. It can detect when all possible steering arcs are impassible because of obstacles. Also it can determine when the roll or pitch of the vehicle has exceeded nominal values. If a problem is detected, error recovery will suspend waypoint navigation. It will also initiate a back-up maneuver along the robot's previous route. [4]

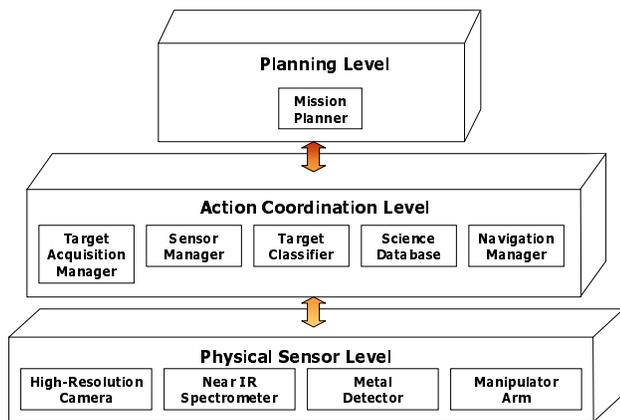


Figure 2: Science Autonomy architecture.

During area coverage pattern execution, a *target acquisition manager* monitors data from a set of acquisition sensors capable of finding new meteorite targets on the ice field. A high resolution three-chip color CCD camera mounted on a pan/tilt unit serves as an acquisition sensor. Approximately every second, pixels in an image are classified into rock or ice categories. When one or more groups of pixels appear as possible meteorites, the target-acquisition manager calculates estimates of the targets' DGPS coordinates based on the assumption that the ground is a perfectly flat plane. All of this information is entered into a central *science database*.

The mission planner places new targets into a priority queue sorted by distance from the robot. Considering one target at a time, the mission planner requests sensor-deployment costs and the cost of maneuvering the robot as it advances toward the target within the work space of the sensors. These costs are weighed against the benefits of continuing a search pattern and an estimate of the information that could be gained by additional sensor data. This estimate is a measure of entropy from the classifier's confidence in its classification of the target from prior

sensor data. The mission planner can then decide to continue the pattern or to investigate the potential meteorite. After the decision to investigate, the mission planner then selects a sensor with the highest information-gain-to-deployment-cost ratio. If necessary the robot may move within the sensor's field of operation, which contains the target. The sensor manager then coordinates deployment of the sensor to the target's estimated DGPS location. This eliminates the requirement that the acquisition sensor track the target during deployment. However the location estimate may contain errors due to deviations of the ground from a flat plane. After any necessary calibration, sensor-data acquisition takes place. The new data are placed in the database, and the classifier processes them to update its classifications and estimated information gains. Finally the mission planner closes the science autonomy control loop by analyzing the new estimated information gains and deployment costs for other sensors at Nomad's disposal. If the classifier is very confident in its results, further sensor data will most likely only verify the current classification. The mission planner will therefore decide to end its investigation of the current target. Otherwise additional sensors are used.

The *target classifier* lies at the core of Nomad's autonomous capability to classify meteorites and rocks. The Bayes network-based classifier uses the visible spectrum (400 - 1,000 nm) for identification of rocks. The classifier detects peaks and troughs in spectral signature by computing the normalized correlation coefficient of the spectrum with a fixed set of Gaussian templates, each with a predetermined width and central wavelength. A positive coefficient indicates a reflection peak, whereas a negative value indicates a trough. Spectrally flat areas have zero correlation. The classifier also computes normalized red, green and intensity coefficients. The intensity coefficient is fairly unreliable but still has considerable diagnostic value. A Bayes network, encoding the statistical distribution of spectral features (correlation coefficients and red, green and intensity values) for each rock type along with their assumed prior probabilities, computes a confidence value from the spectra for each rock type in the network. The network confidence for each rock type corresponds to the posterior probability of the rock being examined, given the current spectral data. [5]

4. Antarctic Demonstration

4.1 Patriot Hills Expedition

The November 1998 expedition led by a team from Carnegie Mellon University; the University of Pittsburgh; NASA Ames Research Center; Laboratoire d'Analyses et d'Architecture des Systems, France; and the Chilean Antarctic Institute (INACH) conducted field work at Patriot Hills, Antarctica. The team reestablished a base camp from a visit to the region earlier in the year at 80°

18° South/81° 16' West and carried out demonstrations and component experiments for six weeks.

The camp had line-of-sight communication with Nomad, which facilitated autonomous navigation and meteorite classification field tests (Figure 3). Experiments included systematic patterned searches, autonomous navigation of various terrains, and automatic classification of rocks and meteorites. In addition, independent experiments of component technologies took place involving landmark-based rover localization, terrain mapping using a millimeter-wave radar, and rover traction control.

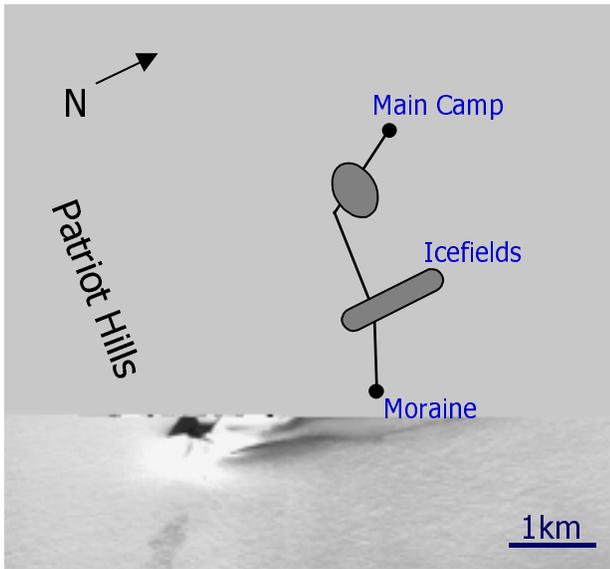


Figure 3: An overhead view of Patriot Hills with annotations of the areas where Nomad performed autonomous exploration and classification of rocks and meteorites.

4.2 Classification Experiments and Results

The Bayes network-based rock/meteorite classifier performed assessments based on the images and reflectance spectra obtained in the field. The spectrometer and its computer were among the modules that Nomad used to classify rock samples. The spectrometer’s sensing head was mounted at the end of a flexible fiber light pipe and had its own light source that could be calibrated on a neutral white surface. The sensing head was shrouded so that, when it was held on the surface of a rock, external light was excluded and the internal light source produced the resulting reflected light spectrum received by the computer.

Because no meteorites were discovered at Patriot Hills during the scouting expedition in early 1998, the team performed classification experiments using meteorite samples loaned by the Smithsonian Institution and the Office of the Curator at Johnson Space Center [6]. Fifteen of the specimens were placed in proximity to the robot for in situ analysis through autonomous interpretation of high-resolution images and reflection spectra. Then Nomad was

teleoperated along the Patriot Hills moraine where classification tests were performed on 42 native rocks. In all cases, a human “assistant” calibrated and deployed the spectrometer; all data was processed aboard the robot. [7]

Because the classifier output provides evidence of a rock being a meteorite and does not produce a binary classification, it is difficult to make precise inferences on the success rate of the field measurements. A simplistic measurement can be obtained by classifying a rock as either a meteorite or a terrestrial rock depending on whether the classifier determines it to be more or less likely to be of extraterrestrial origin given the spectral data. Nomad’s classifier succeeded in correctly classifying 65% of the meteorite samples and 90% of rocks (Figure 4).

The sampling set of meteorites included five “meteowrongs” (terrestrial rocks that bear a superficial resemblance to meteorites) picked up by ANSMET in previous expeditions. They were all rejected by the classifier as being too spectrally inconsistent to be meteorites. Quartzite was the terrestrial rock most commonly mistaken to be of meteoritic origin, whereas the iron meteorites were most likely to be mistaken for terrestrial rocks. As part of the independent component technology demonstrations, the team conducted foot searches for meteorites at Martin Hills and Pirrit Hills—two areas in close proximity to Patriot Hills. Two meteorites were discovered in these field trips and were correctly classified by Nomad’s Bayes-based system. [8]

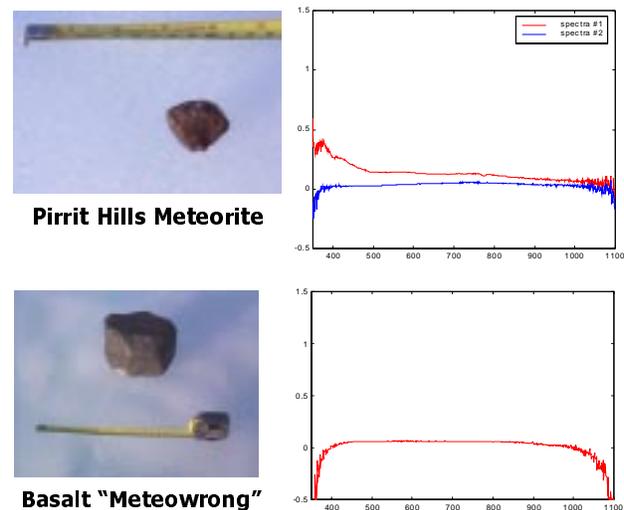


Figure 4: Characteristic spectra of a meteorite (top) and a “meteowrong” (bottom). The Pirrit Hills meteorite was initially considered to be an ordinary rock. Nomad’s classifier correctly recognized it as a meteorite with a high, 86.2% confidence. The classifier correctly classified the basalt as a terrestrial igneous rock despite its spectral similarity to one of the two spectra taken from the Pirrit Hills meteorite. [7]

Automatic sample classification based on image processing and spectral analysis of optical reflection

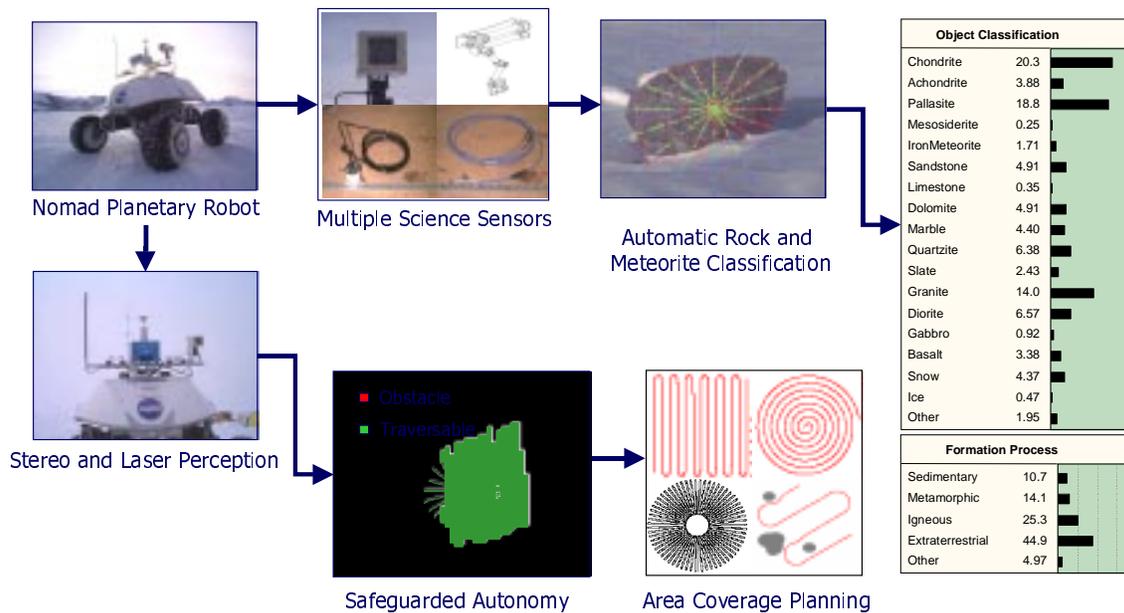


Figure 5: Meteorite classification and navigation technologies involving the use of science autonomy. The upper track shows the three rock/meteorite sensing models (high resolution vision, metal detection, spectroscopy), rock image acquisition, and output from the Bayes classifier. The lower track illustrates Nomad’s sensor head, which includes a laser range finder and two stereo pairs, a map developed by its autonomous navigation indicating traversable (green) and non traversable (red) areas and a selection of patterns that the mission planner can issue when Nomad explores an area.

combined with Bayes networks has proved very effective in the automatic discrimination of meteorites from terrestrial rocks. Moreover, classification using visible to near- infrared reflection spectroscopy yields sufficiently reliable results. The biggest asset of Bayes network-based rock classification is the ease with which data from other sensors such as magnetometers, metal detectors, and microscopic cameras, can be combined with it to improve the quality of the process and ultimately the success rate of sample classification.

4.3 Autonomous-Navigation Experiments and Results

The autonomous navigation system was tested under several weather and terrain conditions. Weather conditions included bright sun, various degrees of cloudiness and many cases of blowing snow. The terrain of the test sites included moraines and ice fields. Moraines are rocky regions that are the most likely areas for finding meteorites. However, they also include rocks over 40 cm in height that pose potential obstacles for the robot. Ice fields are generally very flat with soft depressions. The field tests involved waypoint navigation with stereo vision and laser.

Stereo vision was tested on snow, blue ice and moraines at Patriot Hills as well as under three different weather conditions (clear, overcast, snowy). Under all conditions, stereo could not find enough texture in the scene to produce sufficiently dense disparity maps for navigation. Polarizing filters did improve performance on blue ice, but the results were not sufficient for navigation. The terrain

type had very little effect on the results. The moraine was sparse, so most scenes consisted of blue ice, not rocks. However the weather did have a significant impact on the stereo results. Sunny days provided the best results with blowing-snow conditions yielding the second-best results. Overcast conditions proved the most difficult for the use of stereo, because clouds diffused the sunlight. This factor, combined with the Lambertian surface of a snow field, caused illumination to be almost uniform everywhere. No contrast existed, and it was very difficult even for humans to see depth. Under these conditions, stereo provided almost no terrain information.

The single-line-scan laser unit was tested under the same conditions as the stereo system. Terrain type had no effect on the laser. Even the specular surface of the blue ice fields had no effect on the return signal. Overcast conditions also had no effect on the active laser sensor. However the laser did have problems during periods of blowing snow because the laser beam reflected off of snowflakes. If the beam reflected back to the laser unit, a short-distance measurement would be recorded. If the beam reflected away from the unit, no return signal would be received. During mild conditions of blowing snow, filtering was able to overcome these effects. In heavy storms, such filtering was ineffective. Under those conditions, use of the laser was impossible.

Despite the absence of stereo data, Nomad achieved 10.3 km of autonomous navigation at Patriot Hills using its laser range finder as the primary terrain- perception sensor.

Because the vast majority of terrain negotiated was featureless, stereo matching from disparity maps was impractical. Lack of contrast due to excessive light diffusion during overcast days further hampered stereo. Laser range finding was effective in all cases with the exception of situations in which beam interference from drifting snow led to incorrect range measurements. [4][9]

5. Discussion

The technologies described in this article will greatly influence future Earth-based and planetary exploration involving geological classification. Currently the primary goal for Nomad is to autonomously discover a meteorite in Antarctica. The results of classification experiments conducted during and after the 1998 expedition provide encouraging evidence of Nomad's ability to correctly identify new meteorites in Antarctica. To meet this ambitious and important goal, several technologies are being advanced.

First, in response to problems with the use of laser and stereo vision in polar conditions, more robust obstacle avoidance schemes are being developed. Navigation primarily using inertial and state sensors must be advanced to allow not only recovery from dangerous situations but integration in the autonomous planning and following of paths. Methods to automatically detect new rock targets are also being implemented. These methods must reliably differentiate between rocks and ice under all applicable lighting conditions and terrain types. Furthermore a manipulator arm is being built to precisely deploy sensors to a rock target. The spectrometer used in the expedition must be placed within 1 cm of a target at approximately a 45 degree angle in two dimensions. Inaccurate readings can result if the sample is also underexposed to ambient light. Successful readings will involve effective visual servoing techniques. Finally this project will integrate and validate Nomad's geological classification techniques with high-level cost analysis to create a science robot that can autonomously replan its mission based on the latest scientific results. Rather than deploying all available sensors to every target, this will allow the robot to save power and time by deploying quick and low-power sensors initially. It will then make an initial classification and then decide if further investigation is worthwhile.

Nomad's effective past performance and improving technology make it attractive for use in the study of other areas where rock acquisition and classification by humans are difficult or dangerous, as well as very distant locations where teleoperation is impossible but the benefits to science are significant. Nomad represents the beginning of a new field of advanced robots: scientific explorers capable of making valuable discoveries without direct human guidance.

6. Acknowledgments

This work is supported by the Telerobotics Program of NASA's Office of Space Science, Advanced Technology and Mission Studies Division under Grant NAG5-7707. We would like to thank all members of the Nomad team, participants in the expedition, supporting organizations, and many volunteers that helped along the way. Special thanks go to William Cassidy and Pascal Lee for their role as lead scientists and their field support.

7. References

- [1] Wettergreen, D., Bapna, D., Maimone, M., and G. Thomas, "Developing Nomad for Robotic Exploration of the Atacama Desert," *Autonomous Systems*, Vol. 26, No. 2-3, Amsterdam, Holland, February 1999.
- [2] Bapna, D., Rollins, E., Murphy, J., Maimone, M., Whittaker, W., and D., Wettergreen, "The Atacama Desert Trek Outcomes," *Proceedings of the 1998 International Conference on Robotics and Automation*, Leuven, Belgium, May 1998.
- [3] Shillcutt, K., Apostolopoulos, D., and W., Whittaker, "Patterned Search Planning and Testing for the Robotic Antarctic Meteorite Search," *Proceedings of the 1999 ANS International Meeting on Robotic and Remote Systems for the Nuclear Industry*, Pittsburgh, PA, April 1999.
- [4] Moorehead, S., Simmons, R., Apostolopoulos, D., and W. Whittaker, "Autonomous Navigation Field Results of a Planetary Analog Robot in Antarctica," *1999 International Symposium on Artificial Intelligence, Robotics and Automation in Space*, Noordwijk, Holland, June 1999.
- [5] Pedersen, L., Apostolopoulos, D., Whittaker, W., Benedix, G., and T. Roush, "Sensing and Data Classification for Robotic Meteorite Search," *Proceedings of the 1998 SPIE International Conference on Mobile Robots and Intelligent Transportation Systems*, Boston, MA, November 1998.
- [6] Lee, P., W.A. Cassidy, D. Apostolopoulos, et al., "Search for Meteorites in the Patriot Hills Area, Ellsworth Mountains, West Antarctica," *Meteoritics & Planetary Science* 33, A92-93.
- [7] Pedersen, L., Apostolopoulos, D., Whittaker, W., Cassidy, W., Lee, P., and T. Roush, "Robotic Rock Classification Using Visible Light Reflectance Spectroscopy: Preliminary Results from the Antarctic Meteorite Search Program," *30th International Lunar and Planetary Science (LPSC) Conference*, Houston, TX, 1999.
- [8] Lee, P., W.A. Cassidy, D. Apostolopoulos, et al. (1999), "Search for Meteorites at Martin Hills and Pirrit Hills, Antarctica," *30th International Lunar and Planetary Science (LPSC) Conference*, Houston, TX, March 1999.
- [9] Vandapel, N., Moorehead, S., Whittaker, W., Chatila, R., and R. Murrieta-Cid, "Preliminary Results on the Use of Stereo, Color Cameras, and Laser Sensors in Antarctica," *Proceedings of the 1999 International Symposium on Experimental Robotics*, Sydney, Australia, April 1999.