

11-2000

# Technology and Field Demonstration of Robotic Search for Antarctic Meteorites

Dimitrios S. Apostolopoulos  
*Carnegie Mellon University*

Michael D. Wagner  
*Carnegie Mellon University*

Benjamin N. Shamah  
*Carnegie Mellon University*

Liam Pederson  
*Carnegie Mellon University*

Kimberly Shillcutt  
*Carnegie Mellon University*

*See next page for additional authors*

Follow this and additional works at: <http://repository.cmu.edu/robotics>

 Part of the [Robotics Commons](#)

---

Published In

.

This Article is brought to you for free and open access by the School of Computer Science at Research Showcase @ CMU. It has been accepted for inclusion in Robotics Institute by an authorized administrator of Research Showcase @ CMU. For more information, please contact [research-showcase@andrew.cmu.edu](mailto:research-showcase@andrew.cmu.edu).

---

**Authors**

Dimitrios S. Apostolopoulos, Michael D. Wagner, Benjamin N. Shamah, Liam Pederson, Kimberly Shillcutt, and William L. Whittaker

---

Dimitrios S. Apostolopoulos  
da1v@ri.cmu.edu

Michael D. Wagner  
mwagner@ri.cmu.edu

Benjamin N. Shamah  
bshamah@ri.cmu.edu

Liam Pedersen  
pedersen+@ri.cmu.edu

Kimberly Shillcutt  
kimberly@ri.cmu.edu

William L. Whittaker  
red@ri.cmu.edu

# Technology and Field Demonstration of Robotic Search for Antarctic Meteorites

## Abstract:

*Meteorites are the only significant source of material from other planets and asteroids, and therefore are of immense scientific value. Antarctica's frozen and pristine environment has proven to be the best place on Earth to harvest meteorite specimens. The lack of melting and surface erosion keep meteorite falls visible on the ice surface in pristine condition for thousands of years. In this article we describe the robotic technologies and field demonstration that enabled the first discovery of Antarctic meteorites by a robot. Using a novel autonomous control architecture, specialized science sensing, combined manipulation and visual servoing, and Bayesian classification, the Nomad robot found and classified five indigenous meteorites during an expedition to the remote site of Elephant Moraine in January 2000. This article first overviews Nomad's mechatronic systems, and details the control architecture that governs the robot's autonomy and classifier that enables the autonomous interpretation of scientific data. It then focuses on the technical results achieved during field demonstrations at Elephant Moraine. Finally, the article discusses the benefits and limitations of robotic autonomy in science missions. Science autonomy is shown as a capable and expandable architecture for exploration and in situ classification. Inefficiencies in the existing implementation are explained with a focus on important lessons that outline future work.*

**Keywords:** robotic meteorite search, science autonomy, Bayesian meteorite classifier

## 1. Introduction

Robotic planetary exploration is subject to severe communications and data management restrictions that prevent the return of more than a fraction of the scientific data that can be acquired. To maximize mission return, future generations of robots must achieve a new level of autonomy that goes beyond data collection and achieves autonomous summarization, interpretation and discovery in the absence of frequent human oversight. Automation of the search for and classification of Antarctic meteorites offers a unique case for early demonstration of robotics in a scenario analogous to geological exploratory missions to other planets and to the

Earth's extremes. Moreover, the discovery of new meteorite samples is of great value because meteorites are the only significant source of extraterrestrial material available to scientists.

Antarctica's frozen and pristine environment has proven to be the best place on Earth to harvest meteorites. The lack of melting and surface erosion keep meteorite falls visible on the ice surface. Antarctica's cold, pristine environment and low weathering rates contribute to the preservation of meteorites over thousands of years. Scientists of the National Science Foundation's Antarctic Search for Meteorites program (ANSMET) and other international organizations have collected more than 20,000 Antarctic meteorite samples the vast majority of which originate from the asteroid belt. Only ten specimens have been identified as martian or lunar in origin. ANSMET's searches 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. 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 they are 2 cm to 15 cm across, however they can measure as much as 1 m in diameter and weigh tens of kilograms. Generally their shape is flat elliptical or near spherical with extruded features and a glossy brownish to black color that is attributed to a fusion crust. Their weight is a function of elemental composition. Common elements are iron and nickel mixed with silicates.

A robotic search for Antarctic meteorites faces challenges in addition to those confronted by humans. Vast amounts of knowledge and experience are required to autonomously explore a region and identify meteorite samples from many varieties of terrestrial rocks without input from a field scientist. To accomplish this goal the problem of meteorite search was decomposed into search, detection and

classification aspects. Over a three year program component technologies were identified and tested, and three expeditions to Antarctica were performed. This program's first expedition to Antarctica occurred in January of 1998. Component technologies such as a panoramic camera, a ground penetrating radar and important sensors for meteorite classification such as a metal detector, magnetometer and optical reflection spectrometer were individually evaluated in field conditions. The most important result of the first season was the validation of optical reflection spectroscopy as a method of discriminating meteorites from terrestrial rocks [Pedersen98]. During the second expedition in November 1998, the Nomad robot sought out and classified meteorites in the Patriot Hills region of Antarctica. A human assistant deployed a fiber optic cable of a spectrometer to take readings of rocks *in situ*. This system classified 50 rocks and 15 planted meteorites. The robot also autonomously navigated 10 km of polar terrain using stereo vision, a laser rangefinder and pose sensors [Apostolopoulos99].

Finally, during a four-week expedition to Elephant Moraine in January of 2000, Nomad accomplished the first ever robotic meteorite discoveries and achieved an unprecedented demonstration of autonomous robotics in a science mission to an extreme terrestrial environment. The robot classified five indigenous meteorite samples based on autonomous interpretation of spectral data acquired from indigenous samples. Nomad's success proved the rudiments of "scientific autonomy" as a viable model for robotic exploration of unknown environments. During the Elephant Moraine demonstrations, Nomad explored approximately 2500 m<sup>2</sup> of blue ice following predetermined search patterns and classified 42 spectra of targets which were originally identified as potentially interesting targets by its high-resolution imaging sensor. Nomad's mechanisms proved quite reliable in Antarctica's extreme conditions.

Nomad's expedition yielded several important technical lessons:

- Nomad's science autonomy architecture is a feasible model for robotically performed scientific exploration and *in situ* classification missions.
- Rock and meteorite classification has shown marked improvement when trained on data taken in field conditions.
- Open-loop spectrometer placement cannot guarantee a quality of data equal to that taken in a laboratory under controlled conditions, although a closed-loop placement system should be able to overcome this limitation.
- The state-of-the-art speed for robotic mechanisms and autonomy must be increased to create an efficient meteorite search robot.
- Wide angle imagery could be used to quickly disregard rocks that are clearly not colored like meteorites.

An overview of Nomad's systems is presented in Section 2 with an emphasis on technologies critical to the search and *in situ* classification capabilities. Section 3 details the control architecture that governs the robot's exploration autonomy

and Bayesian-based classifier that enables the autonomous interpretation of geologic data. Section 4 summarizes the Elephant Moraine expedition and discusses both the overall results and specific performance of the target acquisition, visual servoing, search and classification components. Section 5 describes the key technical lessons of the expedition with emphasis on issues of scientific autonomy. Finally, section 6 captures the most significant results of this program.

## 2. Nomad, a Meteorite-Seeking Robot

Nomad is an autonomous field robot with advanced control, perception and cognition of a planetary rover and the first robot to exhibit intelligence to explore an extreme polar environment and autonomously classify indigenous rocks (Fig. 1). Nomad combines advanced mobility with hierarchical distributed control to safely navigate in unknown environments.



**Figure 1: Nomad investigating a meteorite sample using its manipulator arm at Elephant Moraine, Antarctica.**

The robot is equipped with a sophisticated sensor payload that enables navigation and science. A laser rangefinder is used to detect obstacles and roll / pitch / yaw sensors make the robot cognizant of dangerous terrain. A high-resolution color camera mounted on a pan / tilt unit detects and studies rock targets on the ice and snow. A manipulator arm places a visible-to-near-infrared reflection spectrometer within a centimeter of a rock face. For precision placement a camera mounted on the manipulator's wrist allows the sensor deployment software to visually servo to a rock. Nomad runs safeguarded navigation and science autonomy as the two major autonomy software systems that enable safe path following and the planning, coordination and execution of search functions, respectively.

### 2.1 Robotic Mobility

Exploiting an all-wheel independent drive configuration, actuated deployable chassis, electronically coordinated steering and articulated-frame averaging suspension, Nomad possesses superior terrainability [Shamah98].

Nomad features individual propulsion drive units that reside inside each wheel. The advantages of in-wheel propulsion include individual traction control, greater control flexibility and sealed drive units. The chassis can expand or compact by driving two pairs of four-bar linkages with two electric motors, one on each side of the robot. This transforming action leads to a significant increase in vehicle footprint and enables steering by differential actuation of the two deployment motors. Nomad's transforming chassis increases stability by deploying the wheels beyond packaging dimensions once the robot begins a traverse. As shown in Table 1, Nomad's transforming chassis also allows explicit steering by using the same two actuators used to deploy the wheels to cause changes in wheel heading. The transforming chassis is based on the motion of four-bar linkages connected to each wheel. The wheels are actuated in pairs such that the two right wheels move synchronously (as do the two left wheels) to achieve smooth heading change [Rollins98].

Nomad's two steering and four in-wheel driving motors are controlled by software running with the VxWorks real-time operating system. Electrical current and voltage monitors record drive motor power consumption. These data are used for analysis of Nomad's mobility over Antarctic terrain. Encoders determine the position of steering and in-wheel motors for the purposes of driving and dead reckoning. Pose and other telemetry are transmitted via the Network Data Delivery System (NDDS) to other computing modules [RTI00].

## 2.2 Sensing, Navigation and Visualization

Differential GPS is the primary method of localizing Nomad. A Trimble 7400MSi remote DGPS unit and base station are used. This setup provides centimeter position accuracy in north by east coordinates. The robot's position is updated at a rate of five times per second. Position data are transmitted to Nomad's real-time computer via an RS-232 serial line where they are broadcast to all software modules using NDDS.

Two stereo camera pairs were initially used by Nomad to create depth maps for obstacle avoidance. However, findings from the November 1998 expedition showed stereo vision to be ineffective on blue ice and snow in Antarctica due to an extreme lack of texture [Moorehead99]. In place of the cameras a SICK laser rangefinder is used for obstacle detection (Fig. 2). Terrain maps are generated with laser range data at 5 meters ahead of the robot.

A panoramic camera has been used to perform landmark navigation experiments [Deans98]. This system points a Sony color CCD camera up at a spherical mirror to acquire 360 degree views around the robot (Fig. 2). During the latest expedition this sensor was used as a remote viewing sensor to record images of the scenery around the robot. This allowed researchers in the base camp to see what kind of terrain Nomad was traversing as it moved between meteorite search sites.

<b>Physical</b>	
Mass:	725 kg
Power:	~ 3500 watts maximum
Size:	2.4 x 2.4 x 2.4 m deployed
Speed:	0.5 m/s maximum 0.15 m/s nominal
<b>Science</b>	
Visual:	3-CCD, 640 x 480, 24-bit color CCD, 512 x 494, wrist cam
Spectral:	300-1100 nm spectrometer
Deployment:	3 DOF manipulator
<b>Navigation Sensors</b>	
Position estimation:	gyrocompass GPS wheel encoders
Obstacle avoidance:	single line laser
<b>Locomotion</b>	
Propulsion:	4 wheel independent drive
Steering:	left / right coordinated (1 actuator per side)
<p>Nomad chassis fully stowed</p> <p>Nomad chassis fully deployed</p> <p>Coordinated steering Left &amp; right sides servoed to proper steering angles</p> <p>Point turn Left &amp; right sides servoed to equal but opposite angles</p>	
<b>Communication</b>	
Data rate:	1 Mbps
Equipment:	2.4 GHz Wireless Ethernet 900 MHz backup radio

**Table 1: Nomad specifications**

## 2.3 Manipulation

Nomad's manipulator arm serves the critical purpose of accurately deploying and positioning the spectrometer, which is used for classification of designated samples. The manipulator is composed of two rotational joints and a


prismatic joint. The necessary workspace for the manipulator takes into account the maneuverability of the robot and the workspace of the high-resolution camera for viewing targets (Fig. 2). Another consideration is the fact that the arm must not interfere with the laser range finder and the high-resolution camera during driving modes.

The manipulator's SCARA (Selective Compliant Articulated Robot for Assembly) configuration exhibits simplicity of design, simplicity of control, compactness and minimum torque requirements [Spong89]. The key advantage to this configuration is that if the base of the manipulator is level then gravity does not affect the motion of the first two linkages. For level deployment the motors of the first two joints need only account for the inertia of the system.

The sensor head located at the tip of the manipulator's wrist accommodates the spectrometer, a camera, light sources, and one rotational degree of freedom (Fig. 3). The ideal placement of the spectrometer is at 45 degrees to the rock sample surface.

## 2.4 Science Sensors

Potential meteorites are found using a high-resolution 3-chip CCD camera mounted on a pan / tilt head on Nomad's sensor mast (Fig. 2). Zoom and focus are controlled over an RS-232 line with the use of custom-designed lens interface hardware. This allows Nomad's autonomy software to capture fully zoomed images of any rocks it finds on the ice or snow. During a search the lens is zoomed out to find new rocks. Its maximum field of view is 28 degrees.




**Figure 2: Nomad components: (1) manipulator arm (shown stowed), (2) high-resolution science camera, (3) laser rangefinder, (4) panoramic camera, (5) radio modem antenna, (6) high-gain wireless ethernet antenna.**

Nomad deploys its manipulator arm to place a reflection spectrometer within 1cm of the sample. Software guides the arm above an estimate of the sample's position. To overcome errors in this estimate, Nomad uses a 92° wide field of view camera on the wrist of its arm to perform visual servoing.

Once the sample is found in a color image from the arm camera, Nomad can refine its knowledge of the rock's position through a five-step visual servoing process. The arm is moved between the first and second images to create a stereo baseline from which an estimate of the rock's location can be made. To increase accuracy, the wrist is moved closer to the ground for two more images that further refine the location estimate. The wrist is then lowered to within 10 cm of the rock. Due to small deviations of the wrist's horizontal position during vertical translations, another image is taken at this wrist position for a final error correction before taking a spectrometer reading [Wagner00].

Once the manipulator wrist has been servoed to the current sample, a visible light reflection spectrometer obtains spectral data of rock samples. One end of a shielded fiber optic cable is placed close to the rock, while the other end is attached to an Ocean Optics S2000 spectrometer (Fig. 3). Light emitted by lamps on the wrist reflects off the rock and back into a fiber optic cable. Nomad's science autonomy software uses this spectrum to distinguish between rocks and meteorites.



**Figure 3: Close-up of Nomad's science sensors mounted on the manipulator's wrist.**

## 2.5 Computing and Control

Three computers control Nomad: one for real-time control and two dedicated to autonomous operations. The real-time system uses the VxWorks operating system and runs on a Pep VM62 68060 mounted on a VME backplane. Steering arc commands from the navigation software are translated into appropriate steering and drive motor commands using PID control loops. Velocity feedback is used to control wheel motors and position feedback to control its electronically coordinated steering actuators. The real-time system also performs analog to digital conversion so data from several sensors can be broadcast to other systems as NDDS telemetry. This includes encoder values, DGPS position data as well as current and voltage monitors for each of the motors.

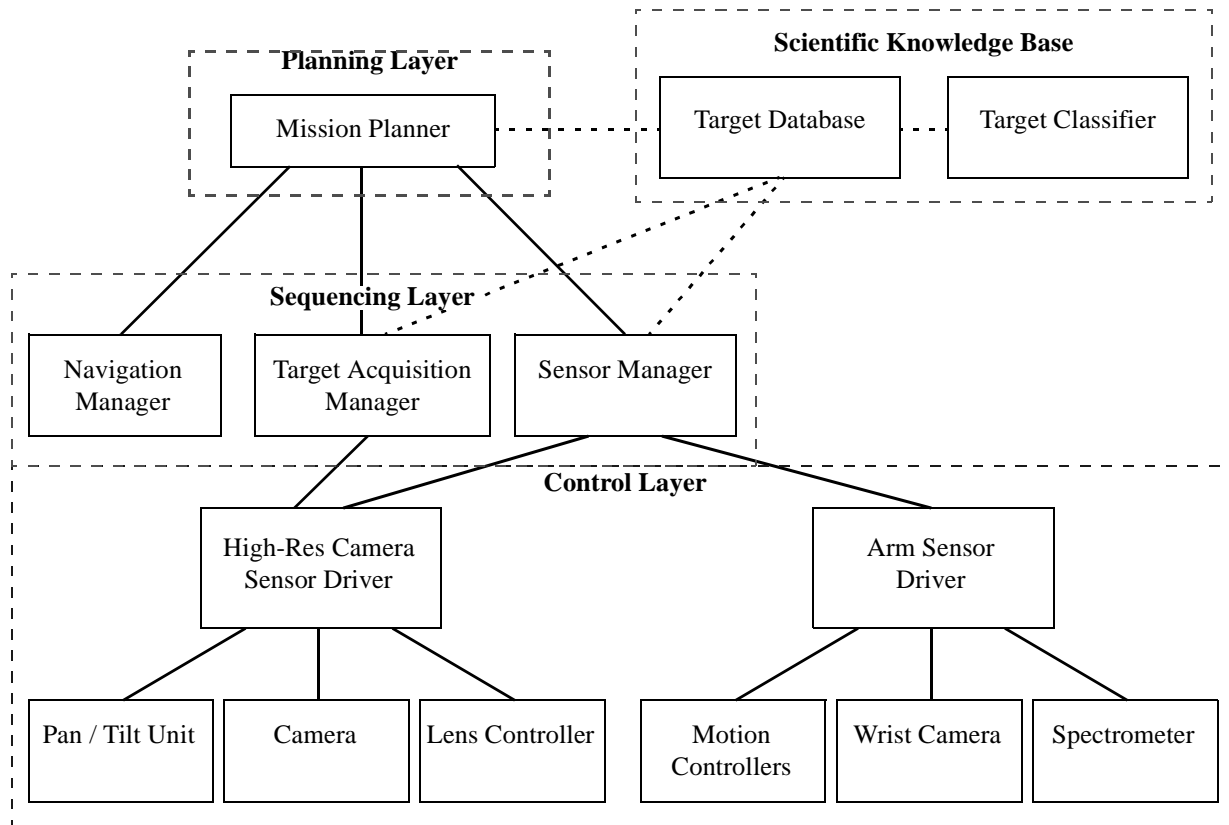


Figure 4: Nomad's SAS control architecture

The autonomous navigation software uses the Red Hat Linux version 5.2 operating system. The hardware is a 133 MHz Pentium processor on a PCI bus inside an industrial PC case. It communicates with the laser rangefinder and sends steering arcs to the real-time system using NDDS. Several of the science autonomy modules run on this computer as well.

Nomad's third computer is a 200 MHz Pentium Pro on a PCI bus, also mounted in an industrial PC case which runs the remainder of the science autonomy software. This computer contains a PCI Imagenation frame grabber, serial ports, an Ocean Optics analog/digital converter and a National Instruments DAQ card to interface with a wide variety of sensors and actuators including three cameras, a pan / tilt unit, a motion controller board, a spectrometer and solar panels. This is all performed under the Red Hat Linux version 5.2 operating system.

### 3. Autonomous Search and Classification

Robots seeking scientific discoveries through intelligent search and data classification must have the ability to plan and execute a search, identify potentially interesting targets and reason about their finds in a deliberate manner. This class of autonomous science robots must also master a seamless coordination of autonomous navigation functions, science sensor deployment and data acquisition. The development and implementation of robotic seeking for Antarctic meteorites onboard Nomad is based on these fundamental

requirements and has been manifested in a hierarchical control architecture called the Science Autonomy System.

#### 3.1 Science Autonomy System (SAS)

The Science Autonomy System (SAS) can be compared to a three-tiered architecture (Fig. 4). Three-tiered architectures typically distinguish control, sequencing and planning in a hierarchical structure [Gat98]. The lowest layer, the *control layer*, contains the robot's sensors and actuators. These sensors allow the science autonomy system to observe the world. This layer contains control loops between software, sensors and actuators to create *primitive behaviors* for operation such as calibrating and deploying sensors. Raw data sensed by the control layer are utilized in two operational modes: *acquisition* and *identification*. In acquisition mode, science autonomy attempts to find new targets of interest in the world. In identification mode, study and classification of a target takes place. To carry out both objectives, the physical sensors may have to function in different ways or coordinate with other sensors. The *sequencing layer* is responsible for this coordination by interpreting the results of the control layer. The highest level, the *planning layer*, considers the results of the sequencing layer and creates a plan that will optimize mission variables such as energy costs and scientific information gains [Shillcutt98]. Commands are sent to the sequencing layer to deploy sensors and construct navigation plans. The planning layer is therefore responsible for merging navigation and

science autonomy within the SAS. Outside each layer lies the *scientific knowledge base*, composed of a database and a classifier. The knowledge base classifies targets based on sensor data and provides the planning layer with estimates of the information that would be gained by further study of each target. It also contains information about all science targets registered to the system, and communicates with many other modules in the SAS. Throughout its architecture, the SAS exhibits a hierarchical control scheme, with each layer in the architecture responsible for awaiting completion of lower layer actions.

### 3.1.1 Control Layer

#### Primitive Behaviors and Sensor Drivers

Primitive behaviors are tasks that can be composed into more complex goal-directed behaviors by the sequencing layer. In the SAS, these include sensor deployment, sensor data acquisition and sensor calibration. The primitive behaviors of Nomad's science autonomy system are defined by the *sensor driver* interface. This abstract interface allows the sequencing layer to communicate with the control layer in a well-defined manner, defined by the following methods:

- **startup** initializes all subcomponents and prepares the sensor for deployment.
- **calibration** runs any calibration routines that are specific to the sensor hardware.
- **deployment** deploys the sensor to a target coordinate given the current robot pose.
- **dataAcquisition** captures and saves sensor data. Deployment should be done before this method is called.
- **unDeployment** returns the sensor to its startup state.
- **diagnostic** tests the sensor and returns the most critical error.
- **shutdown** prepares the sensor to cease being used.
- **getCost** returns the cost (e.g. time and energy requirements) of deploying the sensor to a given science target coordinate.
- **getWorkspace** returns the workspace area of the sensor.
- **getStatus** returns the status information of the sensor without taking any actions (unlike the diagnostic method).

These methods translate into the primitive behaviors available to the sequencing layer. Multiple sensors can inherit the sensor driver class and implement these methods in hardware-specific ways. However, the specifics of sensor hardware are abstracted from the sequencing layer. All interfacing with the mobility platform itself is currently done through the autonomous navigation system, a self-encapsulated architecture.

#### Sensor Drivers on Nomad

The sensor drivers Nomad uses include a high-resolution camera sensor driver and a manipulator sensor driver. The sensor driver's deployment method converts estimated DGPS coordinates of a target to pan and tilt angles. The data acquisition method stores the image on the CCD along with

meta-data such as the pixel coordinates of any rocks in the image. Nomad's manipulator contains multiple sensors and actuators, all of which are encompassed by the manipulator sensor driver. Deployment involves coordination between three axes of a motion control board and a color CCD camera mounted on the wrist of the arm to visually servo the wrist down to the potential meteorite target. Again, estimated DGPS target coordinates are passed into the deployment module and errors in these estimates are overcome by visual servoing. Data acquisition involves storing spectrometer data and images from the wrist-mounted camera. Calibration of the spectrometer is performed after every fourth spectrum taken. Calibration requires placing the wrist so the spectrometer is directly above a white reflectance standard calibration target.

### 3.1.2 Sequencing Layer

This layer is comprised of three modules: the *target acquisition manager*, the *sensor manager* and the *navigation manager*. These modules respectively coordinate the recognition of new science targets, the analysis of existing science targets and the safeguarded execution of mobility commands. Modules in the sequencing layer are aware of what sensor drivers exist in the system and what types of data these sensors provide, but they are not aware of implementation specifics of any given sensor. They both interpret data and execute command sequences to accomplish goals of the planning layer.

#### Target Acquisition Manager

*Target acquisition* describes the process of finding previously unlocated science targets for later investigation. Sensed data from the control layer enters the target acquisition manager and estimated locations of new targets are sent to the planning layer. Methods for acquiring new targets include segmentation of imagery or analysis of metal detector readings. These methods are dependant on the sensor data, meaning that an *a priori* knowledge of the sensor data structure is required to find new targets. Therefore the target acquisition manager uses *target acquisition drivers* that have the data structure knowledge required to find new targets. Target acquisition drivers may communicate with multiple sensor drivers to acquire sensor data with which to find new targets. Once this examination is complete, they pass the target acquisition manager an object that is independent of the sensor data format.

In the current implementation of Nomad's science autonomy system only one target acquisition driver exists, the high-resolution camera acquisition driver. Images from this camera are processed to segment rocks from background ice (Fig. 5). A linear combination of blue and green color ratios is calculated for each pixel in an image window. Shadows in the image generally create several areas of noticeably different blue-green color ratios. A high standard deviation of these ratios initiates an intensity-based shadow compensation routine. Pixels with a low blue-green ratio are designated as rock. The circumference of these regions are found, instead



of the entire connected region, resulting in a time savings of up to 175%. The approximate centroids of these regions are then calculated. Regions that are too small, too large or too similar to lines (often a representative of image artifacts) are ignored.

Using the pose of the robot and an assumption that the ground is a flat plane between the robot and the new targets, the DGPS location of each target is calculated and placed in the target database.



**Figure 5: Raw (top) and segmented (bottom) images processed by the high-resolution camera acquisition driver.**

### Sensor Manager

The sensor manager carries out commands from the planning layer that request sensor data of existing science targets. The planning layer commands the sensor manager to deploy a particular sensor to a target. The sensor manager retrieves the target's record from the database, which contains its estimated DGPS coordinates. The sensor manager module is responsible for executing a sequence of commands that perform any sensor calibration required, deploy the sensor to the target and acquire sensor data for analysis. Once the sensor is finished, the new data are put into the target database for access by the rest of the system, and the sensor manager indicates success or describes errors. Similarly to the target acquisition manager, the sensor manager uses

*sensor manager drivers* to encapsulate knowledge of sensor data structure.

There are two sensor manager drivers in the current implementation of Nomad's science autonomy system: a high-resolution camera sensor manager driver and a manipulator arm sensor manager driver. The manipulator sensor manager driver is the more complex sensor manager driver in the system. It communicates with both the wrist camera and the arm sensor drivers to coordinate the proper deployment of the arm. When the robot has maneuvered into the manipulator's workspace, the manipulator sensor driver takes a second image of the target with the high-resolution camera sensor driver. This image provides an updated estimate of the target's location to ensure proper manipulator placement.

### Navigation Manager

The navigation manager is responsible for converting high-level mobility plans from the planning layer to specific steering arcs to pass to the autonomous navigation system. For instance, the planning layer may request that a specialized search pattern be executed, e.g. spiral shaped pattern, and the navigation manager will begin sending desired steering commands to the autonomous navigation system every second. A separate obstacle detection module simultaneously sends its desired steering commands to the navigation system, which arbitrates between the two inputs and sends a final steering command to the robot. [Moorehead99]. The *navigation manager* is responsible for following a mission planner chosen search path using the pure pursuit path tracking algorithm [Coulter92]. When executing patterns, it dynamically updates the robot's next waypoint every second, keeping the waypoint approximately 5m ahead of the current robot position throughout the path. This lookahead distance enables the robot to quickly return to the path after avoiding obstacles or examining targets, without creating oscillations around the path. For the curved paths of the spiral pattern, a shorter lookahead distance is used. Any type of search pattern can be interrupted, saved, and resumed later, even if the robot is moved in the meantime [Skillcutt99]. The navigation manager also enacts pre-planned maneuvers provided by the mission planner, sending steering commands one by one to the navigation system. Such maneuvers may be necessary to put desired targets into the workspace of the robot's sensors.

#### 3.1.3 Planning Layer

This layer is responsible for deliberative tasks that create plans to be carried out by the sequencing layer. The mission planner is the sole module in this layer. It first initiates a search pattern via the navigation manager, and starts the target acquisition manager, continually listening for input from other modules. When new targets are found by the target acquisition manager, the planner stops the robot and adds the targets to its list. Considering each target in the list, in combination with each available sensor, the mission planner requests an estimate of the information that could be

gained by additional sensor data. A separate information gain value is defined by the classifier for each combination of targets and sensors. If the information gain for a particular target/sensor pair is below a threshold, then that pair is ignored and no further cost calculations are performed for it.

If at least one target and sensor pair passes this first cutoff test, additional costs are calculated. The distance to the target from the robot, the time cost of using the sensor and the need for maneuvering the robot into the sensors workspace are all considered. These costs are compared for each target / sensor pair, and the best one is selected for further investigation. If no targets pass the information gain test, then the mission planner decides to continue the pattern and search for more targets.

After a decision to investigate, the mission planner may request a maneuver to move the robot and its sensors into range of the target. The mission planner then activates the sensor manager, requesting the desired data on the selected target. New sensor data are placed in the database for the classifier to process. The classifier then puts its findings back into the database, triggering the notification message for which the mission planner is waiting. If the target has a high probability of being a meteorite, the remote observers will be alerted. Otherwise, the mission planner re-analyzes its target list to determine if more target / sensor pairs should be selected at the current time. The more confident the classifier is about its results, the less information additional data will provide, leading to no further examination of the target. When no other sensors or targets pass the information gain cutoff test, the mission planner will resume the search pattern.

If the mission planner decides to command sensor deployment but the robot is too far from the target, a maneuver is planned. Maneuver planning is performed using a model of how the robot responds to five commands: turn left, turn right, go forward, go backward or stop. The model also contains state information about the previously commanded move. To speed the planning process, the model was limited to just those five motions and discretized into 1 second intervals, which has been experimentally shown to be sufficient. The A\* search algorithm is used to find the lowest cost plan [Hart68]. The search goal not only includes the position but also the heading of the robot. The robot's yaw is important because sensors have workspaces limited both by distance and angle to the target.

Although A\* ensures an optimal plan, the mission planner augments the A\* algorithm using heuristics to decrease average execution time. It checks if the robot will need to back up first, to compensate for the limited turning abilities of the robot which are not included in the optimistic cost heuristic. The A\* algorithm also finds only the first possible plan, which is at the edge of the sensors workspace. To place the robot closer to the center of the workspace, the maneuver planner adds an additional five steps to the plan, picking the set of steps which result in a final position of the targets closest to the center of the manipulators workspace.

### 3.2 Rock / Meteorite Classifier

The target classifier lies at the core of Nomad's autonomous capability to classify meteorites and rocks, encapsulating the system's scientific knowledge. Its role is to use sensor data to classify a science target as belonging to predefined classes. Based on its current assumed knowledge of a target, it also calculates the information gain that would result from readings from each remaining sensor in the system.

The classifier is responsible for deciding, based on sensor data, the likelihood that a sample is either a meteorite or terrestrial rock. The nature of this task, and because it has to be integrated with the operation of a mobile robot, imposes several conditions. First, the classifier must learn how to classify from examples. Human experts have trouble explaining how they can identify many rocks, and will often refer to an example. In the words of a veteran Antarctic meteorite searcher, "First you find a few meteorites, then you know what to look for." However, it is difficult to acquire large sets of training data under realistic field conditions. This has required two earlier expeditions to Antarctica, as well as visits to the Arctic and the Atacama desert in Chile. Prior application domain specific knowledge must therefore be embedded in the classifier to minimize the amount of example data needed for it to learn how to discriminate different rock types. Finally, field conditions change unexpectedly, so the classifier must adapt as the robot operates.

The robot has multiple sensors which are not deployed all at once. The classifier must accept incomplete data and compound evidence as more sensor data become available. Moreover, there is a cost associated with the use of each sensor. Sensors such as a spectrometer are particularly expensive to use as the robot must maneuver to bring rock samples into its workspace. Therefore it is desirable to make an initial classification with low deployment cost, long range sensors such as a color camera, reserving expensive sensors for verification of promising rock samples. A corollary of this concept is that the classifier should accept prior evidence from other sources, including expert knowledge on what to expect in a particular location. Rock classes are often ambiguous, and the distinctions between certain types fuzzy at best [Dietrich79]. The classifier must handle this ambiguity and indicate several likely hypotheses if a definite classification cannot be achieved.

These requirements argue convincingly for a Bayes network based approach. Bayes networks are graphical models of the statistical relationships between various random variables, such as a rock type and sensor readings. As variables are observed (sensor measurements obtained), the posterior distribution of the remainder can be efficiently calculated using Bayes rule. This is a principled way of compounding evidence, and prior information is naturally represented by prior probabilities [Cheeseman88]. Moreover, the intuitive graphical structure of a Bayes network makes it easy to encode physical constraints or domain specific knowledge into the network topology, thus reducing the intrinsic

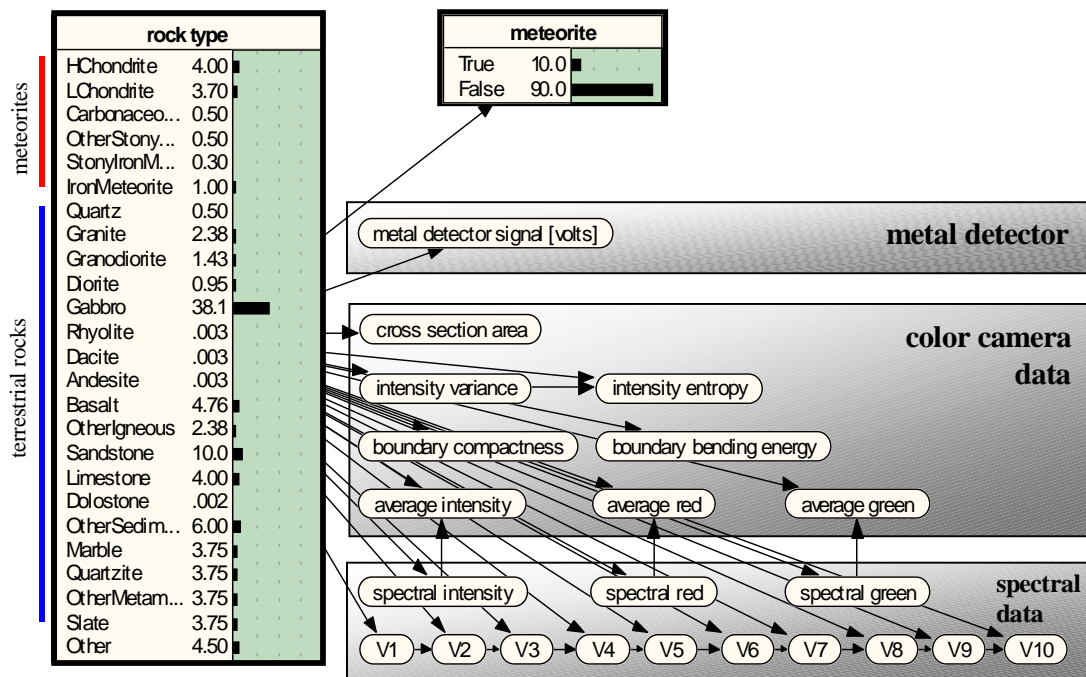


Figure 6: Bayesian network for rock and meteorite classification.

dimensionality. [Charniak91] gives a simple introduction to Bayes networks, while [Pearl88] is the definitive reference.

The classifier works asynchronously whenever new sensor data are placed in the target database. Once it receives notice of this event, it classifies the science target based on both the new data received and data previously recorded in the database. Therefore each new sensor reading enhances previous classifications rather than replacing them. New targets have default class membership probabilities and maximum information gain estimates both of which are updated upon new sensor readings.

### 3.2.1 Bayes Network Topology

Figure 6 depicts the Bayes network developed for rock and meteorite classification using feature vectors computed from color images, spectra and metal detector data. These features have been deliberately chosen to minimize statistical interdependence, modeled by the links between feature nodes and the intermediate node *rock type*, which enumerates all terrestrial rocks and meteorite types likely to be encountered.

The result sent to the SAS system is the network-computed probability of a rock being a meteorite. Due to uncertainty in the estimation of prior probabilities this cannot be regarded as the “true” probability, but rather as a measure of classifier confidence. In practice this is not an issue, as a threshold confidence value for selecting rocks is chosen based on rate of classification (ROC) curves and acceptable meteorite recognition and false positive rates.

### 3.2.2 Sensors and Feature Vectors

The raw sensor data is processed to extract appropriate feature vectors to be used by the Bayes network classifier. In Antarctica Nomad acquired reflectance spectra and color images of sample rocks. Relevant features in a spectrum are the location and size of peaks and troughs and the average reflectance over different wavebands: red, green and blue (Fig. 7 (a)). Spectral troughs and peaks are detected by correlating the spectrum with a set of 10 templates over a 50 nm region of support. Restricting the degree of overlap between templates minimizes statistical dependencies among the resulting spectral features. Normalizing the correlation coefficients makes them conditionally independent of the average spectral intensity and robust to changes to scale. This is important because field spectral measurements have large scaling errors. The result is a 13 element real valued feature vector from the original 1000 element spectrum.

Color images are harder to interpret (one of the rocks in Figure 7 (b) is a meteorite). First the rock needs to be segmented from the background of snow and ice in the image, using a partially observable Markov model [Pedersen98]. Features of interest are the rock cross sectional area (a proxy for size, requiring that the scaling of the images be known), average color, and simple texture and shape metrics. These are useful because meteorites tend to be small and dark compared to terrestrial rocks. All real valued features are quantized prior to being entered into the Bayes network, which cannot handle continuous quantities.

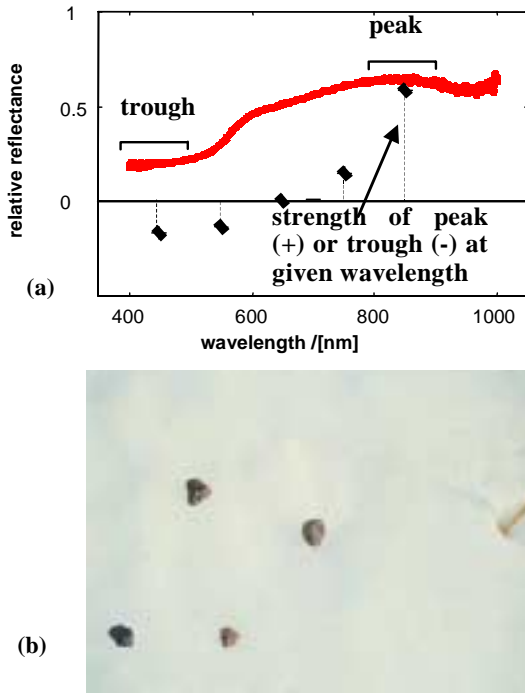


Figure 7: Example spectrum with associated features (a) and color image of meteorite and terrestrial rocks (b).

### 3.2.3 Network Training

The conditional probability matrices (CPMs) describe the probability distributions of network sensor feature nodes given rock type and other parent nodes. They are learned by incorporating examples using the algorithm in [Spiegelhalter93]. This is a stable and true Bayesian learning algorithm. Furthermore, each training sample is weighted to reflect its frequency of occurrence for the rock type that generated it. This is especially important as multiple sensor readings are taken from single rock samples, especially those exhibiting significant surface variation.




Figure 8: Rate of classification curves based on laboratory visual, spectral and combined data.

Using meteorite and terrestrial rock data acquired in the laboratory, partitioned into 75% training, 25% testing cross

validation sets, the Rate of Classification (ROC) curves in Figure 8 are generated. Note the superior classification with spectra versus classification with color images only. In fact, given a spectrum, a color image does not improve classification. However, it is easier to acquire a color image than a spectrum, so a color camera is still useful for preliminary screening.

### 3.3 Meteorite Search and Classification Process

The operation of the SAS during meteorite search is as follows: a command to search an area designated by DGPS coordinates is given. Nomad's mission planner commands a pattern to exhaustively search the area. The robot's navigation and obstacle avoidance modules supervise search path progress and maneuver Nomad around potentially dangerous obstacles.

While the search pattern progresses, Nomad's target acquisition manager processes images from the high-resolution camera to detect rocks. A threshold-based color segmentation is applied to discover dark rocks on generally bright ice or snow. New targets are placed into the database and the mission planner's target list. The mission planner stops the robot and selects the best target to examine using a set of deployment cost parameters, such as distance.

The mission planner commands Nomad to maneuver to a target for investigation with the spectrometer. Once the target is within workspace of Nomad's manipulator arm, the robot stops and Nomad's sensor manager deploys the arm. The manipulator is initially deployed to a target position determined from the high-resolution camera image. Then using the downward looking camera in the wrist of the manipulator visual servoing is used to position the tip of the spectrometer to within 1 cm of the rock.

The sensor manager triggers the acquisition of spectroscopic readings once the spectrometer probe is positioned just above the rock target. A detailed analysis of spectral features is autonomously performed by the classifier. This improves the initial classification and enables the distinguishing of terrestrial rock types. The classifier reports its results to the target database.

Upon completion of the classification process, the sensor manager returns the arm to its stowed position and returns control to the mission planner. Should the target be given a high probability of being a meteorite, Nomad waits for human validation. If there is another target within the manipulator's workspace, the robot stays in its location and repeats the arm deployment and classification cycle. Otherwise, the mission planner maneuvers Nomad to the location of the next target in the list or resumes its search pattern to seek new targets. The autonomous search is complete when Nomad covers the designated area and classifies all identified targets.

## 4. Robotic Meteorite Search Demonstration

Nomad's robotic technologies were validated during a field expedition to the remote site of Elephant Moraine in January 2000.

### 4.1 Elephant Moraine Expedition

Elephant Moraine, located at 76 deg 16' S, 157 deg 12' E is a barren region of Antarctica that is known to contain numerous meteorites. Because of this, ANSMET has conducted searches of Elephant Moraine in the past. The site is located approximately 160 miles northwest of McMurdo Station, the main logistical center of Antarctica. The field team landed in McMurdo in late December, 1999. Supplies for the field expedition to Elephant Moraine were collected and staged by support personnel. The McMurdo area also provided an environment in which to test Nomad's capabilities. Functions never previously tested in Antarctica, such as target acquisition and the manipulator, were validated. Nomad and the field team consisting of six robotics researchers and one ANSMET meteorite expert were airlifted to Elephant Moraine on January 7 and January 10, respectively.



**Figure 9:** Nomad executing target acquisition experiments near the Elephant Moraine camp.

A camp consisting of living tents and an operations tent was set up. By January 14, the robot was again assembled and ready for operation. Experiments gathering data and quantifying the performance of manipulator arm deployment, image segmentation, target acquisition and classification were performed until January 19. At this point, meteorite search demonstrations began in earnest. Ten demonstrations were performed in various locations between January 20 and 29. The camp was taken down and the team departed on February 2, 2000.

### 4.2 Technical Results

#### 4.2.1 Overall Search and Classification Results

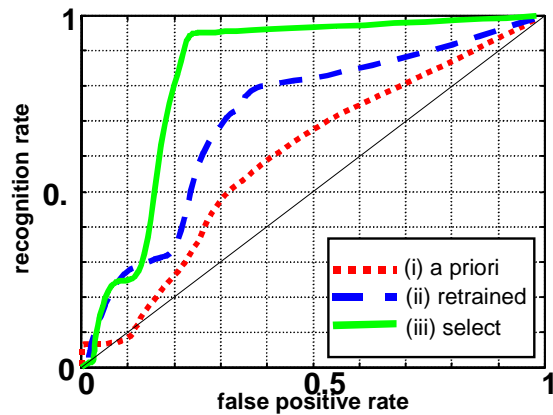
During the ten demonstrations performed at Elephant Moraine, a total area of approximately 2500 m<sup>2</sup> were searched in around 16 hours spread over 10 days. Three

meteorites were found and correctly classified during autonomous searches, out of a total of 42 targets examined. An additional two meteorites were correctly classified during tests performed without patterned searches. Sensor deployment was performed with a high degree of autonomy, although touch sensors designed to lower the manipulator's wrist directly above a rock failed, forcing users to teleoperate the wrist at the very end of each deployment. All other mechanisms functioned reliably even in extremely cold conditions. Image segmentation required for target acquisition and manipulator visual servoing proved capable in many conditions, although its parameters had to be hand-tuned for different lighting conditions that ranged from bright direct sunlight to diffuse overcast conditions. Rock / meteorite classification was effective overall but systematically poor for a particular class of rocks upon which the robot had not been trained but were common at Elephant Moraine.

Item	Results
Patterned Search	2500 m <sup>2</sup> searched in 16 hours 37% of time spent searching 47% of time spent deploying sensors 6% of time spent maneuvering robot to target
Classification	3 meteorites correctly classified 22 rocks correctly classified 17 misclassifications <b>42 total spectra classified</b>
Solar Data	Two 10 W panels recorded solar irradiation continuously
Weather Data	Temperature and wind speed recorded continuously
Panoramic Cam	42 images taken
Public Participation	Daily download of 1-2 MB recorded data to publicly accessible web site

**Table 2:** Nomad's performance at Elephant Moraine.

A problem with the camera zoom control prevented the acquisition of close up color images of samples, so classification had to be done with spectra only. Overall performance is indicated by the ROC performance curves in Figure 10. These were generated from a test set of rocks and meteorites in the moraine, some of which were analyzed multiple times. Figure 10 (i) is using the a priori classifier built from the lab data used to generate Figure 8, acquired prior to arrival in Antarctica. Performance is less than initially expected (Fig. 8). There is a notable improvement in Figure 10 (ii), the ROC curve for the same classifier further trained with field data (from 8 rocks and 2 meteorites not in the test set).



**Figure 10: Rate of classification curves based on data acquired by Nomad in Antarctica.**

Even with retraining, classification is systematically bad for a particular class of rocks (hydro-thermally altered dolerites and basalts) that were found in Elephant Moraine. These rocks are stained red with iron oxide (rust) whose spectrum has a prominent peak at 900 nm, precisely where many meteorite spectra also have a peak. This is not surprising, given that most meteorites contain metallic iron, and therefore can have surface rust. Unfortunately, these rocks were absent from the initial training set and not initially expected in this area. Performance is much better if these rocks are removed from the test set Figure 10 (iii) and the retrained classifier is used.

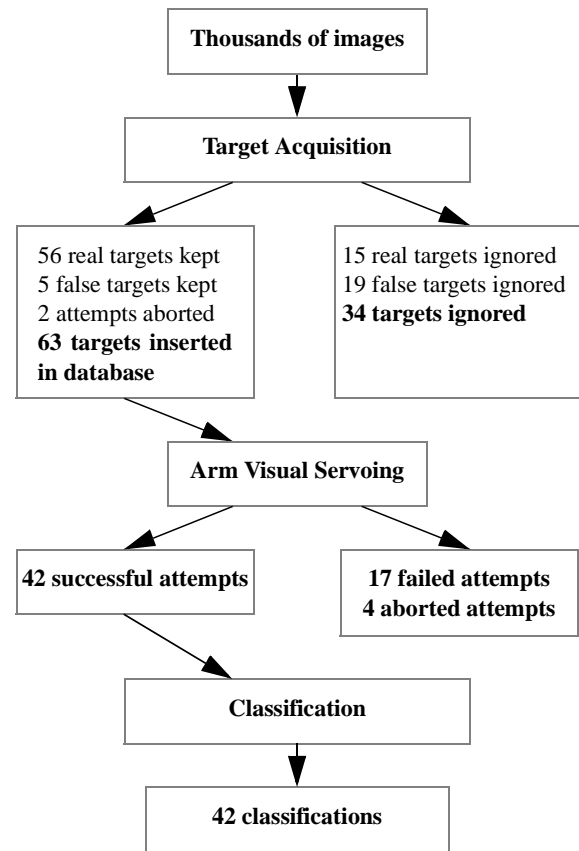
#### 4.2.2 Target Acquisition

Nomad reliably acquires a target in two steps. Initially, the robot scans the ground 2.5 meters in front of the robot looking for targets. During initial testing at Elephant Moraine, researchers discovered that this single step was not reliable enough to ignore false positives.

Several cases occurred when shadows appeared as dozens of individual targets. The team found that when Nomad's camera attempted to zoom into these false targets, they would be lost entirely or continue to appear as dozens of individual targets. Real targets were much more likely to be recognized as either a single or a small number of individual targets. Therefore Nomad's target acquisition manager software was modified. If a target initially detected was lost or appeared as a large number of small targets, it would be ignored.

Target acquisition success is therefore measured by its ability to keep real targets and ignore false targets that may have been initially detected. During the course of ten demonstrations, the target acquisition manager analyzed thousands of images. 97 images included real or false targets that were segmented. Target acquisition was successful on 75 of these cases (a 79% success rate). Of the 75 successes, 56 real targets were analyzed and 19 false targets were correctly ignored. Only one of these successful real targets required human intervention, as the rock examined was a small basalt

whose diameter was around 1 cm. Failure modes include ignoring or incorrectly imaging real rocks or attempting to analyze false targets. Images of objects other than the target could be captured, such as a shadow or the robot's manipulator arm itself. This could occur if the rock was in front of the robot. Real targets were thrown out or incorrectly imaged 15 times. False targets were analyzed five times. In two cases, the team overrode the SAS by aborting target acquisition. These decisions were made if the robot repeatedly came across the same target many times. Note that high-resolution images were not used for classification purposes due to a problem with the camera zoom control.



**Figure 11: Target acquisition and classification results.**

#### 4.2.3 Visual Servoing

Success in visual servoing of Nomad's manipulator was defined by proper placement of the manipulator arm above the sample and acquisition of a spectrum. The manipulator was deployed successfully on 42 out of 63 attempts not manually aborted (72% success rate). Unfortunately the touch probe mounted on the wrist, which was designed to halt the wrist when it reached the proper height above a sample, failed in the cold environment. Therefore teleoperation from the base camp was required to safely guide the wrist down for the final 10 cm to the sample. Teleoperation was performed by viewing wrist camera images from the base camp and verifying downward

movement steps with an observer in the field. Not considering this teleoperation problem, 13 of the successful attempts required some other type of human intervention. Poor segmentation was by far the most common reason requiring additional intervention. In these cases, the human operator visually selected the rock's position from a wrist camera image, which allows visual servoing to continue automatically. In four cases, the team overrode the SAS not to deploy the arm, mostly when repeated targets were encountered.

These arm deployment results can be further analyzed with respect to demonstration location. When autonomous searches were performed outside the moraine (demonstrations before January 28), 18 out of 23 attempts were successful - a 78% success rate. When the robot was taken into the moraine for the last two demonstrations January 28 and 29, only 25 out of 36 attempts were successful resulting in a lower, 69% success rate. Snow fell on the evening of January 27, partially or completely covering all rocks with a slight layer of snow. As seen in Figure 12, this snow caused exposed portions of a single rock to be disconnected. To the robot's segmentation routines these disconnected regions would appear as two distinct targets. During each step of Nomad's visual servoing process, it is important that the same target be selected. If two targets are in close proximity, errors in the targets' initial DGPS estimates cause uncertainty in target selection, resulting in a failed arm deployment attempt.




Figure 12: High-resolution image of snow-covered rock.

#### 4.2.4 Image Segmentation

One of the biggest problems for both human and robotic meteorite hunters are the varying lighting conditions in Antarctica. The conditions change from complete white out situations where nothing is visible to clear bright days when shadows and sun glare prevail. The best conditions are when there is medium cloud cover such that shadows and glare are minimized.

Image segmentation performed by the target acquisition and arm visual servoing systems thresholded images based on the blue-green ratios of pixels. Shadow compensation is used to make the system more robust in varying lighting conditions. However, the values of multiple parameters must be tuned to make this system effective in different lighting conditions. Specifically, the performance is highly sensitive to parameters that affect when and how shadow compensation is performed. This factor decreased the efficiency of Nomad's searches because segmentation performance had to be briefly analyzed before each demonstration. Using solar panel responsiveness to detect sky conditions, along with a model of the solar ephemeris, can provide information for autonomously selecting the best parameters [Shillcutt00].

Three main lighting conditions were seen: direct sunlight, shadows and diffuse light, or overcast, conditions. Several images from the high-resolution camera were combined to produce Figure 13. It plots the color ratio versus intensity of pixels in each of the three above mentioned lighting conditions. Nomad's current segmentation method thresholds these images by drawing a vertical line in each of these images. Pixels to the left of the line would be classified as rock. Figure 13 confirms field experience that diffuse lighting conditions aid in finding rocks. The background ice pixels cluster away from those of rocks. As expected, the worst case scenario was when shadows fell into the image. In this case, a single vertical line cannot be drawn that effectively discriminates shadowed background ice from rocks. The shadowed rocks group near the bottom of the plot due to their dark color. But their pixels' color ratios are not distinct from those of shadowed ice, making segmentation difficult.


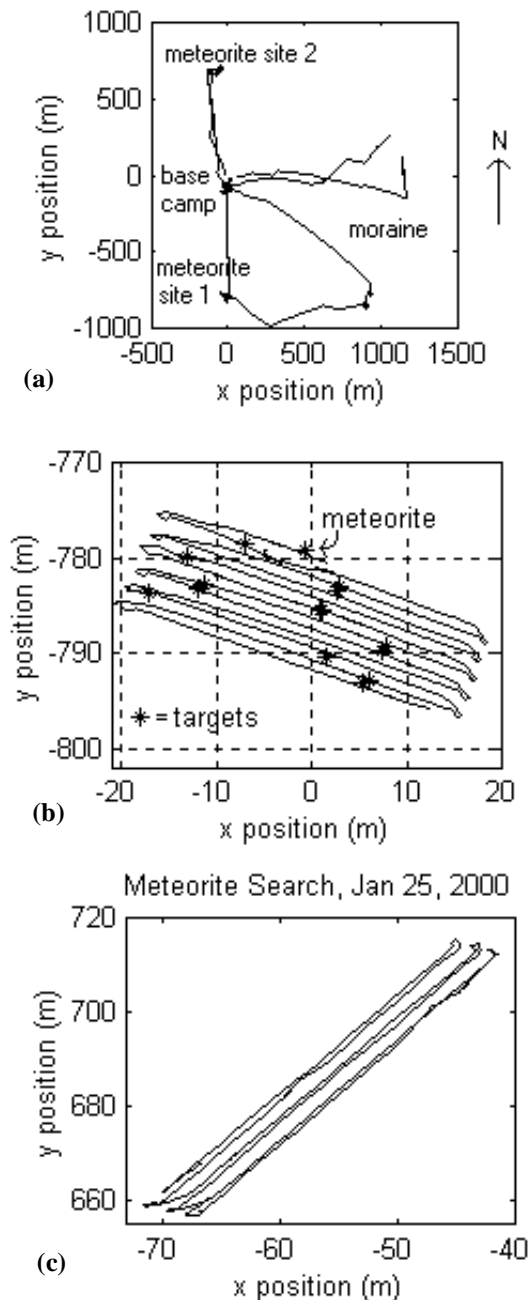


Figure 13: Pixel color ratio vs. intensity of three distinct cases of lighting. From left to right: direct sunlight, shadows and diffuse light in overcast conditions.

#### 4.2.5 Coverage Planning and Patterned Search

Navigational activity during the Antarctic expedition was recorded using the robot's differential GPS sensor. A ground plot of all the locations Nomad visited is shown in Figure 14. Nomad traveled from the base camp to the southern moraine location on January 19, and then to meteorite site 1 on January 20. Nomad returned to base on January 23, and then went on to meteorite site 2 on January 24. Nomad drove to the northern moraine location, via the base camp, on January 26 - 27, and returned to base on January 30.



**Figure 14:** DGPS plots of Nomad's: (a) overall expedition search paths, (b) patterned search leading to the discovery of the first meteorite, and (c) patterned search in an wide ice field

A close-up view of the first meteorite search demonstration area is shown in Figure 14 (b), where three searches took place from January 20 - 22. In the third search, 15 targets were investigated, one of which was successfully identified as a meteorite.

The accuracy of the coverage patterns followed during the expedition was marred occasionally by nearly steady-state distance errors occurring during the straight portions of the rows. The robot would tend to follow a line slightly to the side of the desired row, causing one pair of rows to be pinched together, while the adjacent pair was spread apart (Fig. 14 (c)). The cause of this path deviation was compass heading error, which made the robot think it was heading toward its waypoint when it was in fact heading slightly away from it. This error seemed to vary depending on the robot's location, and in addition, was possibly aggravated by a slow recovery of the compass after the robot turned around at the end of the rows. One solution to this problem is using more feedback to correct for steady-state errors. Also, calculating the heading from DGPS readings can be used in combination with the compass to produce a more robust heading value.

## 5. Technical Lessons

### 5.1 Productive Field Autonomy for Science Missions

Nomad's search efficiency was limited by two major factors. First, the operating speed was slow. The robot travelled at 15 cm/s during its searches. Even if the robot never paused to analyze a rock it could only travel 0.54 km in an hour. The manipulator arm also moves slowly, both because of a low velocity and because only one joint was actuated at a time. The mechanical speed limit of this system has clearly not been explored. The second major factor limiting Nomad's search efficiency was target acquisition. The high-resolution camera's maximum field of view was less than 28 degrees, and the camera was pointed at the ground approximately 2.5 m in front of the robot. This resulted in a search row width of only 1.5 m. This results in a maximum search area of around 800 m<sup>2</sup>/hr if the robot never pauses to deploy its arm.

In areas near moraines, the robot spent approximately 50% of its time deploying its manipulator to take spectra of the large number of rocks. Many of these rocks were clearly not meteorites due to their size and color. Unfortunately, images from the high-resolution camera were judged unsuitable for classification. The bright sunlight reflecting off the ice closed the lens' auto-iris giving most rocks a black appearance. The classifier uses black coloring as an indication of a meteorite, so this problem caused rocks to be incorrectly classified as meteorites. Were it not for this problem, many rocks should have been quickly judged to be terrestrial in origin. If information gains would have been truly implemented in the system, it is reasonable to expect that a meteorite hunting robot could instantly decide that a large percentage of terrestrial rocks are not worth costly spectrometer



investigation. This would clearly speed up the process, resulting in a larger total area searched.

## 5.2 Data Quality Assessment and Interpretation

During meteorite searches, the SAS had no knowledge of sensor data quality. Often the spectrometer was not properly positioned over rock samples, either missing them completely or approaching them at a bad angle to the rock's surface. In theory, the classifier has the capability to recognize outliers, but cannot discriminate between those due to unusual rocks and those due to bad readings. Furthermore it did poorly at recognizing ice because the spectral features were chosen to optimize rock recognition. There was no external measure of confidence in the data.

While the visual servoing algorithm of the arm sensor driver was effective there was not a reliable method of sensing the safe and ideal final height of the sensor head. Therefore, it was teleoperated for the final 10 cm in the vertical direction. The only way to verify placement was visual inspection by a researcher. The researcher's subjective comments were relayed to base camp only after the robot had finished its analysis.

An automated system to verify correct sensor placement is essential for both proper classification and to inform the mission planner that more data is needed. It would provide a measure of confidence that the data being classified was valid. The sensor head must therefore have some form of closed-loop control routine that receives feedback on the position of the spectrometer relative to a rock sample. This represents a difficult problem because of space constraints on the sensor head, the variable geometry of rocks and the sensitivity of the spectrometer.

Sensor noise and systematic effects due to autonomous robot placement of sensors on samples in the unstructured and uncontrolled polar environment are significant. They are hard to know a priori and need to be learned from data acquired by the robot, and in field conditions, as demonstrated by the significant improvement in classification achieved after field retraining.

With the caveat that training be continued using data acquired by the robot in the field, the Bayes network approach to robotic rock classification is a viable approach to this task. Nomad did autonomously identify several meteorites. However, in areas with iron-oxide stained, hydrothermally altered rocks the reflection spectrometer must be supplemented by other sensors, such as metal detectors, magnetometers or more exotic spectrometers (thermal emission or Raman), obviously at greater cost.

## 5.3 Design of the Science Autonomy System

The modularity of the SAS resulted in robustness and manual adaptability to unforeseen problems. If one module in the system experienced difficulties, it could be individually restarted without disabling the entire system. The system was easily altered when high-resolution imagery proved unsuitable for classification. Additional alterations were

made during the mission such as changes to the target acquisition system.

Some sequencing details of the architecture did cause operations to slow unnecessarily. For instance, the sensor manager places sensor data in the database only after all sensor action is completed (i.e. the sensor is returned to stowed position). The classifier begins processing when it receives a notification of new data from the database. Once the mission planner sees that the classifier is finished with the current target, it either returns to the search or begins study of the next target in its queue. In the case of Nomad's manipulator arm, sensor data are ready as soon as a spectrometer reading is finished. However, it would be unsafe for the robot to move until the arm is finished stowing. Because the operation of the classifier and mission planner essentially depended on only one signal from the sensor manager -- that data were ready -- execution time was wasted during arm stowing that could have been used for classification. If the sensor manager had two signals, one to notify of new sensor data and one to indicate that the sensor motion is complete, this inefficiency could be remedied.

Several times during its demonstrations, Nomad ran over rocks it had already seen. While this was mostly due to a small search row width, the robot often did have knowledge of the rock's location as it ran it over. These rocks could be designated as obstacles in the mission planner's global map and therefore avoided.


Future work could also be done in the realm of target acquisition. Currently only one target acquisition driver is used, but if multiple drivers were operational there is no means to fuse their data together. For instance, if an image showed a rock on the ice and a metal detector being swept in front of the robot found a signal, the system could only recognize two new targets even if they were really the same object seen by different sensors.

A good way to address this problem could be through the use of evidence grids [Moravec89]. Not only would this allow fusion of data from different sensors, but multiple readings from a single sensor could be combined to give a more accurate target location. For each image that the existing target acquisition driver captures, all suitable segmented objects are designated as unique targets. The frequency at which the driver captures these images is designed to minimize the chances that a particular rock is seen twice, but even when working this unnecessarily constrains the amount of information the system receives about a rock. Evidence grids are a proven solution to this problem, and have been applied successfully in obstacle detection. Repeated evidence from single sensors, or evidence from other sensors is compounded using Bayesian statistics. This is effective at eliminating false positives. For example, lighting effects, shadows and reflections often caused spurious detections which might disappear in images taken from a different viewpoint. Because targets are entered into the database every time they are detected, often the same rock would be entered multiple times if it was imaged more than once. An

evidence grid approach would be immune to this as targets are automatically tracked, and only entered into the database when sufficiently certain and close enough for investigation.

## 6. Conclusions

Through autonomous search and unsupervised classification of geological samples, this program produced the first discovery of meteorites by a robot and demonstrated the rudiments of robotic scientific autonomy.



**Figure 15: Nomad's meteorite finds: The top row shows three ordinary meteorite types found and classified on January 22/25/26, respectively. The images were acquired by the wrist camera during placement of the spectrometer probe (visible in all three images). The bottom row shows the fourth ordinary meteorite found on January 27 and a special type found on January 29 and classified as "interesting rock." Nomad's classifier had been trained on this rare type only once before using a sample provided by the Smithsonian Institution.**

Central to Nomad's success is the effective integration of science autonomy with real-time data interpretation. Robotic planning based on science mission parameters has proven essential to achieve an effective search, despite the slow pace of Nomad. Other critical technical conclusions have been found that outline future research directions. These include:

The primary technical conclusions are:

- Autonomous geological classification requires that the robot can independently evaluate the quality of sensor data. This implies the need for methods to assess the accuracy of science sensor functionality, placement by mechatronic means and recalibration. This is possibly the greatest challenge for autonomous science robots.
- Autonomous search strategies must take into account the various environmental and terrain properties. The primary reason for the success of the target acquisition system was the exploitation of rich color and reflectance information in distinguishing rocks from the ice / snow background. Target acquisition was greatly challenged in the moraine where rock density increased and terrain became a mix of soil and ice.
- Planning and sensor selection based on information gain and deployment costs are key elements for productive exploration. Scientific exploration tasks often entail the study of tremendous areas containing large numbers of unremarkable objects. Exploration is further constrained by a robot's power, communication or computing limita-

tions. An autonomous robot that considers information gain is drastically more efficient at scientific exploration because it can automatically disregard uninteresting objects while utilizing its resources for notable targets only.

Future scientific missions could clearly benefit from Nomad's technologies. Robust, high-speed, autonomous robots imbued with science autonomy and sensors like Nomad's could effectively perform geological classification missions in remote polar regions and other planetary analogous environments. The envisioned robotic autonomy could enable pioneering discovery missions such as search for life at the Earth's extremes and beyond.

## Acknowledgments

This program has been supported by the National Aeronautics and Space Administration (NASA) under grants NAG5-7707 and NAG9-1090. The authors would like to thank Melvin Montemerlo, David Lavery and Chris Culbert of NASA for their long-standing support of this work. We would also like to thank the National Science Foundation (NSF) and Antarctic Search for Meteorites project (ANSMET) for providing scientific guidance, field expertise and excellent logistical support. We are especially appreciative of the contributions of Scott Borg, Ralph Harvey, John Schutt and William Cassidy.

This work would not have been possible without the dedicated efforts of the Nomad Meteorobot team. Thank you all for a wonderful endeavor.

## References

- [Apostolopoulos99] D. Apostolopoulos, M. Wagner, W. Whittaker, "Technology and Field Demonstration Results in the Robotic Search for Antarctic Meteorites," *Proceedings of the International Conference on Field and Service Robotics*, Pittsburgh, PA, 1999, pp.185-190.
- [Charniak91] E. Charniak, "Bayesian Networks Without Tears" in *AI Magazine*, Winter 1991, 12(4), pp. 50-63.
- [Cheeseman88] P. Cheeseman, J. Kelly, M. Self, J. Stutz, W. Taylor, D. Freeman, "AutoClass: A Bayesian Classification System," *Proceedings of the Fifth International Conference on Machine Learning*, Ann Arbor, MI, 1988, pp. 54-64.
- [Coulter92] R. Coulter, "Implementation of the Pure Pursuit Path Tracking Algorithm," *Technical Report CMU-RI-TR-92-01*, Robotics Institute, Carnegie Mellon University, January, 1992.
- [Deans98] M. Deans, Web Page: [http://www.frc.ri.cmu.edu/projects/meteorobot/Expedition97/Experiment\\_Panoramic/Experiment\\_Panoramic.html](http://www.frc.ri.cmu.edu/projects/meteorobot/Expedition97/Experiment_Panoramic/Experiment_Panoramic.html), Robotics Institute, Carnegie Mellon University, 1998.
- [Dietrich79] R. Dietrich, B. Skinner, *Rocks and Minerals*, J. Wiley & Sons, New York, 1979.
- [Gat98] E. Gat, "On Three-Layer Architectures", *AI and Mobile Robots*, D. Kortenkamp, P. Bonasso and R. Murphy eds., MIT/AAAI Press, Cambridge, MA, 1998.
- [Hart68] P. Hart, N. Nilsson and B. Raphael, "A Formal Basis for the Heuristic Determination of Minimum Cost Paths," *IEEE Transactions on Systems Science and Cybernetics*, SSC-4(2), 1968, pp.100-107.

[**Moorehead99**] S. Moorehead, R. Simmons, D. Apostolopoulos, W. Whittaker, "Autonomous Navigation Field Results of a Planetary Analog Robot in Antarctica", *International Symposium on Artificial Intelligence, Robotics and Automation in Space*, Noordwijk, Holland, 1999.

[**Moravec89**] H. Moravec, "Certainty Grids for Sensor Fusion in Mobile Robots," in *Sensor Devices and Systems for Robotics*, Alicia Casals, Editor, Springer-Verlag, Berlin, 1989, pp.243-276.

[**RTI00**] RTI, 2000, Web Page: <http://www.rti.com>.

[**Pedersen98**] L. Pedersen, D. Apostolopoulos, W. Whittaker, G. Benedix, 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, 1998.

[**Pearl88**] J. Pearl, *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*, Morgan Kaufmann, San Mateo, CA. 2nd edition 1991.

[**Rollins98**] E. Rollins, J. Luntz, B. Shamah, W. Whittaker, "Nomad: A Demonstration of Transforming Chassis," *Proceedings of the 1998 IFAC Workshop on Intelligent Components for Vehicles*, Seville, Spain, March 1998.

[**Shamah98**] B. Shamah, D. Apostolopoulos, E. Rollins, and W. Whittaker, "Field Validation of Nomad's Robotic Locomotion," *Proceedings of the 1998 SPIE International Conference on Mobile Robots and Intelligent Transportation Systems*, Boston, MA, 1998, pp. 214-222.

[**Shillcutt98**] K. Shillcutt, W. Whittaker, "Modular Optimization for Robotic Explorers," *AAAI Fall Symposium on Integrated Planning for Autonomous Agent Architectures*, Orlando, FL, 1998.

[**Shillcutt99**] K. Shillcutt, D. Apostolopoulos, W. Whittaker, "Patterned Search Planning and Testing for the Robotic Antarctic Meteorite Search," *1999 International Topical Meeting on Robotics and Remote Systems for the Nuclear Industry*, American Nuclear Society, Pittsburgh, PA, April 25-29, 1999.

[**Shillcutt00**] K. Shillcutt, "Environment Based Navigational Planning for Robotic Explorers," Ph.D. Thesis, Robotics Institute, Carnegie Mellon University, 2000.

[**Spiegelhalter93**] D. Spiegelhalter, A. Dawid, S. Lauritzen, R. Cowell, "Bayesian Analysis in Expert Systems," *Statistical Science*, 8(3), 1993, pp. 219-283.

[**Spong89**] M. Spong, M. Vidyasagar, *Robot Dynamics and Control*, John Wiley & Sons, New York, NY, 1989.

[**Wagner00**] M. Wagner, "Experimenter's Notebook: Robotic Search for Antarctic Meteorites 2000 Expedition," *Technical Report CMU-RI-TR-00-13*, Robotics Institute, Carnegie Mellon University, June, 2000.