Onboard Autonomous Rover Science

Rebecca Castano, Tara Estlin, Dan Gaines, Caroline Chouinard, Ben Bornstein,
Robert C. Anderson, Michael Burl, David Thompson, Andres Castano and Michele Judd
Jet Propulsion Laboratory
4800 Oak Grove Drive
Pasadena, CA 91109
818-393-5344
firstname.lastname@jpl.nasa.gov

Abstract—The Onboard Autonomous Science Investigation System (OASIS) was used in the first formal demonstration of closed loop opportunistic detection and reaction during a rover traverse on the FIDO rover at NASA's Jet Propulsion Laboratory. In addition to hardware demonstrations, the system has been demonstrated and exercised in simulation using the Rover Analysis, Modeling, and Simulation (ROAMS) planetary rover simulator [1]. We discuss several system enhancements including new planning and scheduling capabilities and image prioritization. We also describe the new end-of-traverse capability that includes taking a partial panorama of images, assessing these for targets of interest, and collecting narrow angle images of selected targets. Finally, we present several methods for estimating properties of rocks and provide a comparative assessment. Understanding the relationship of these methods is important to correctly interpret autonomous rock analyses performed during a traverse. [1] [2]

TABLE OF CONTENTS

1. INTRODUCTION

The Mars Exploration Rovers (MER) continue to make history as they endure the martian winter and send back valuable scientific data. In early November 2006, the *Spirit* MER reached an important milestone as it survived 1000 martian days (sols), traveling more than 6,876 meters (4.27 miles) [Figure 1]. Although the *Opportunity* MER has not been on Mars as long as *Spirit*, it has out-traveled its twin and had logged an even more astonishing 9,406 meters (5.85 miles). With hopes and expectations for the mission to continue well into the future, scientists are warming up to the possible benefits that autonomous operations might bring to the table.

There are a number of autonomous rover capabilities currently in development for future in-situ missions. The capability we focus on in this work, onboard autonomous rover science, continues to grow in importance as rover lifetime and travel distances increase. OASIS [2-6], an Onboard Autonomous Science Investigation System, is a JPL-managed project designed to maximize mission science on rover missions with long traverses.

OASIS is designed to operate onboard a rover identifying and reacting to serendipitous science opportunities. Science opportunities can include detection of dust devils and clouds [6, 7] and novel rocks that the rover has not seen before [3]. OASIS analyzes data the rover gathers, and then prioritizes the data based on criteria set by the science team. At the next opportunity for transmitting data back to Earth, the data is already prioritized – ensuring that the most valuable data is sent first.

As OASIS is working to prioritize the data, it is also searching for specific targets it has been told to find by the science team [3]. If one of these targets is found, it is identified as a new science opportunity and a "science alert" is sent to the planning and scheduling component of OASIS. After reviewing the rover's current operational status to ensure that it has enough resources to complete its traverse and act on the new science opportunity, OASIS changes the command sequence of the rover.

The rover is instructed to stop its current traverse, locate the rock that triggered the science alert, and take additional data (e.g., color image, closer grayscale image, spectrometer reading) on that rock. In addition, the system now enables the rover to either turn and collect data on the identified target rock or to drive to the target so that closer measurements can be collected. Once it has completed this additional measurement, the rover reverts back to its original plan and continues on its traverse.

1

[1] 1-4244-0525-4/07/$20.00 ©2007 IEEE.
[2] IEEEAC paper #1475, Version 4, Updated January 23, 2007

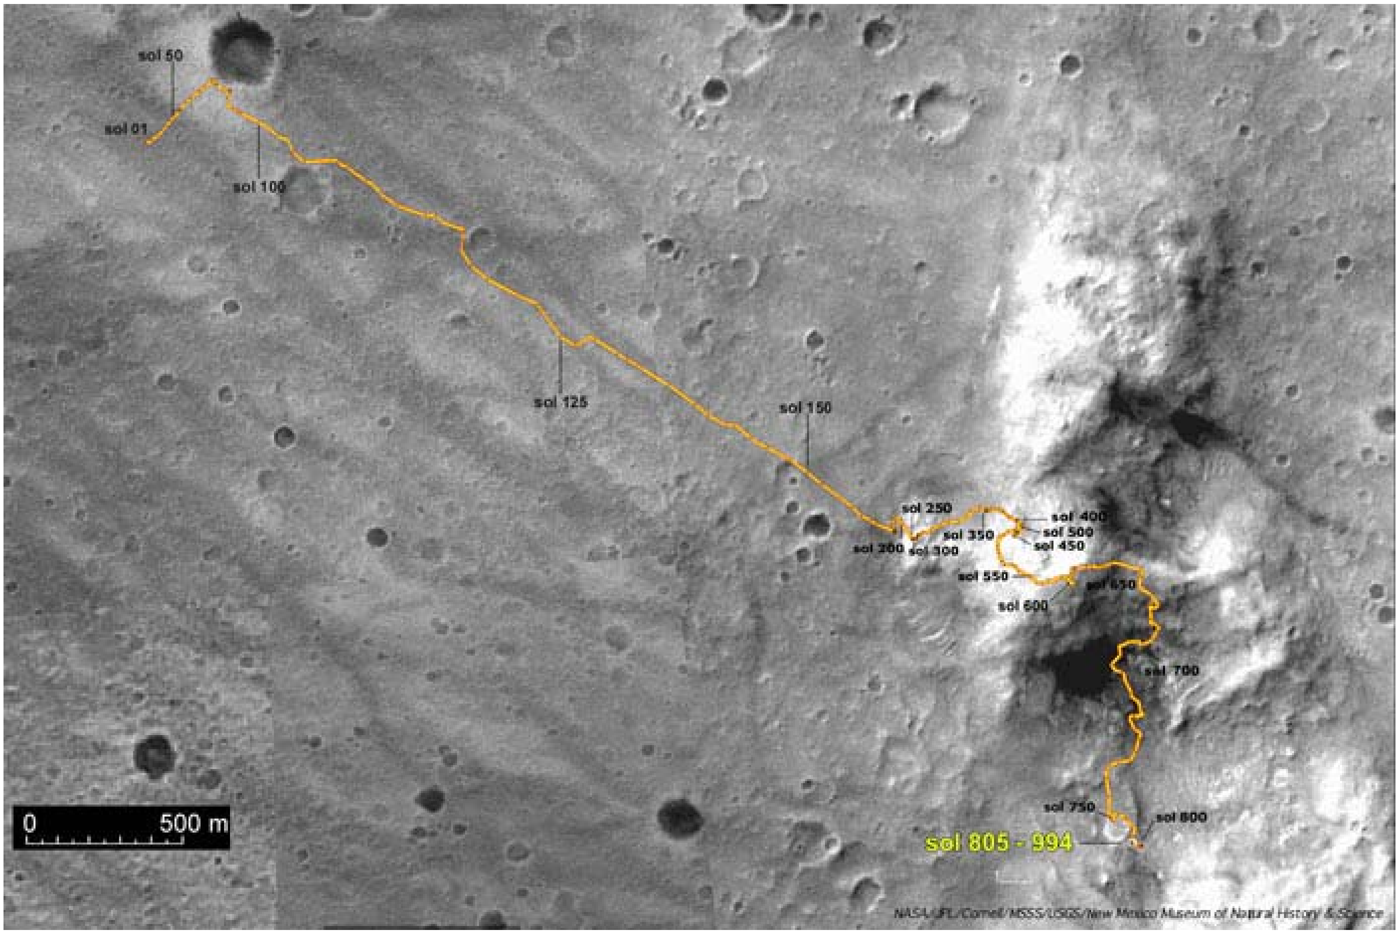

Figure 1 As of sol 986 (Oct. 11, 2006), *Spirit's* total odometry logged in at 6,876 meters (4.27 miles). *Spirit* passed its 1000 sol milestone in this location, as the milestone occurred during a three week solar conjunction (the rover did not move because during solar conjunction the sun interferes with transmissions between Mars and Earth). Image credit: NASA/JPL/Cornell/USGS/New Mexico Museum of Natural History and Science.

A new end-of-day capability was added to OASIS this year. After the rover has completed the image captures of its final location at the end of the sol, OASIS can run these images through its onboard rock detection program(s) and identify targets for a remote sensing instrument (such as a close-up panoramic image or a Laser Induced Breakdown Spectrometer, or LIBS). As resources permit, OASIS can instruct the rover to point to these targets, take data and return the data to Earth at its next communication cycle.

The OASIS system includes three primary components: feature extraction, data analysis and prioritization, and planning and scheduling. In Section 2, we briefly describe the system and these components. Section 3 addresses the development of Rockster, an algorithm developed to segment rocks in rover imagery. The importance of using multiple rock finders and setting rock finder parameters is outlined in Section 4, while Section 5 discusses using the results of these rock detectors on autonomous target selection and sampling. In Section 6 we discuss OASIS integration with Maestro. Section 7 describes how the ROAMS rover simulation environment supports the development, testing and experimentation of the OASIS system. Section 8 discusses RockIT, the OASIS Rock Identification Toolkit, which is a mature, cross-platform, graphical program originally designed to help geologists rapidly and accurately label rocks (or particles) in images. Finally, in Section 9 we describe related work and conclude with a summary in Section 10.

2. OASIS OVERVIEW

To assess and subsequently prioritize the scientific value of a set of collected images, the information within the images must first be extracted. A geologist in the field gets information about a site by identifying geologic features including the albedo[3], texture, shape, size, color, and

2

[3] Note that we are using the term albedo to refer to the brightness of a rock in the image. Technically this is the DN value; we are not calculating the true albedo of the rock which would require taking the incident solar flux,

arrangement of rocks, and features of the topography, such as layers in a cliff face. The geologist analyzes and assesses this data, and then initiates some action based on the analysis, such as taking a sample or taking some additional measurement of an interesting rock.

For an autonomous system to help a remote geologist investigate a traversed region, the system must be able to perform, albeit in a very simple way, these same types of functions. Clearly, one critical function is to identify rocks or other objects that may be of interest for the scientist.

There are three major components that comprise OASIS:

- **Extract Features from Images**: Detects objects of interest and provides a characterization of their properties. This module both locates rocks in the images and extracts rock properties (features) including shape, texture and albedo. Image analysis, when applied to a sequence of images taken in a stationary position for dust devil observation campaigns, can also be used to identify the motion of a dust devil moving across the field of view [7].

- **Analyze and Prioritize Data**: Uses the extracted features to assess the scientific value of the planetary scene and to generate new science objectives that will further contribute to this assessment. This module consists of three separate algorithms that analyze the collected data and prioritize the rocks. A new set of observation goals is generated to gather further data on rocks that either conform to the pre-set specifications of the science team, or are so novel in comparison to the other rocks, that another data measurement may be required.

- **Plan and Schedule New Command Sequence**: Enables dynamic modification of the current rover command sequence (or plan) to accommodate new science requests from the data analysis and prioritization module. A continuous planning approach is used to iteratively adjust the plan as new goals occur, while ensuring that resource and other operation constraints are met.

3. ROCKSTER

Rockster [8] is an algorithm developed to segment rocks in rover imagery. Rockster has been successfully integrated into a number of high-level demonstrations including the SOOPS year end demonstration (Science Operations on Planetary Surfaces [9] – a JPL interdisciplinary task to evaluate technology developments for their potential to make significant improvements in the overall science return of future missions) and live exercises of the OASIS/CLARAty [10,11] software carried out with the FIDO Rover in the JPL Mars Yard.

Figure 2 provides a high-level view of the steps involved in the algorithm. Like several other attempts at automatic rock segmentation (e.g., [5, 12]), the Rockster algorithm focuses on intensity edges in grayscale imagery. Rockster initially locates partial boundary contours of rocks using a procedure similar to the well-known Canny edge detector [13]. In particular, an intensity gradient is calculated over the image; ridges in the intensity gradient are linked together using non-maximum suppression, hysteresis thresholding, and edge-following yielding a set of raw contours.

This initial set of contours does not directly provide a usable segmentation of the rocks from the background due to various problems. These include: spurious contours from the sky-ground boundary (horizon line), texture within individual rocks, texture present in the background, incorrect linking choices at the junctions between contours, and unclosed contours around an object due to gaps in the gradient information (for example, areas along the rock boundary where the rock intensity and background intensity are too close to reliably separate). Rockster attempts to resolve these problems by splitting the initial contours into low-curvature fragments. Potential T-junctions that were missed by the edge detector are identified and used to further split fragments into even smaller pieces. A gap-filling mechanism is then applied to add new contour fragments between existing fragment endpoints. The final step is to regroup the edge fragments into coherent contours, which is accomplished through background flooding. Conceptually, water is poured into the image from the sides but the water is not allowed to cross over any edge fragments; thus, regions that are totally enclosed by edge fragments remain "dry" while other areas become "wet". Extracting contours around the dry areas yields the final rock segmentation.

---

the orientation of the rock facet, and the calibration of the camera into consideration.

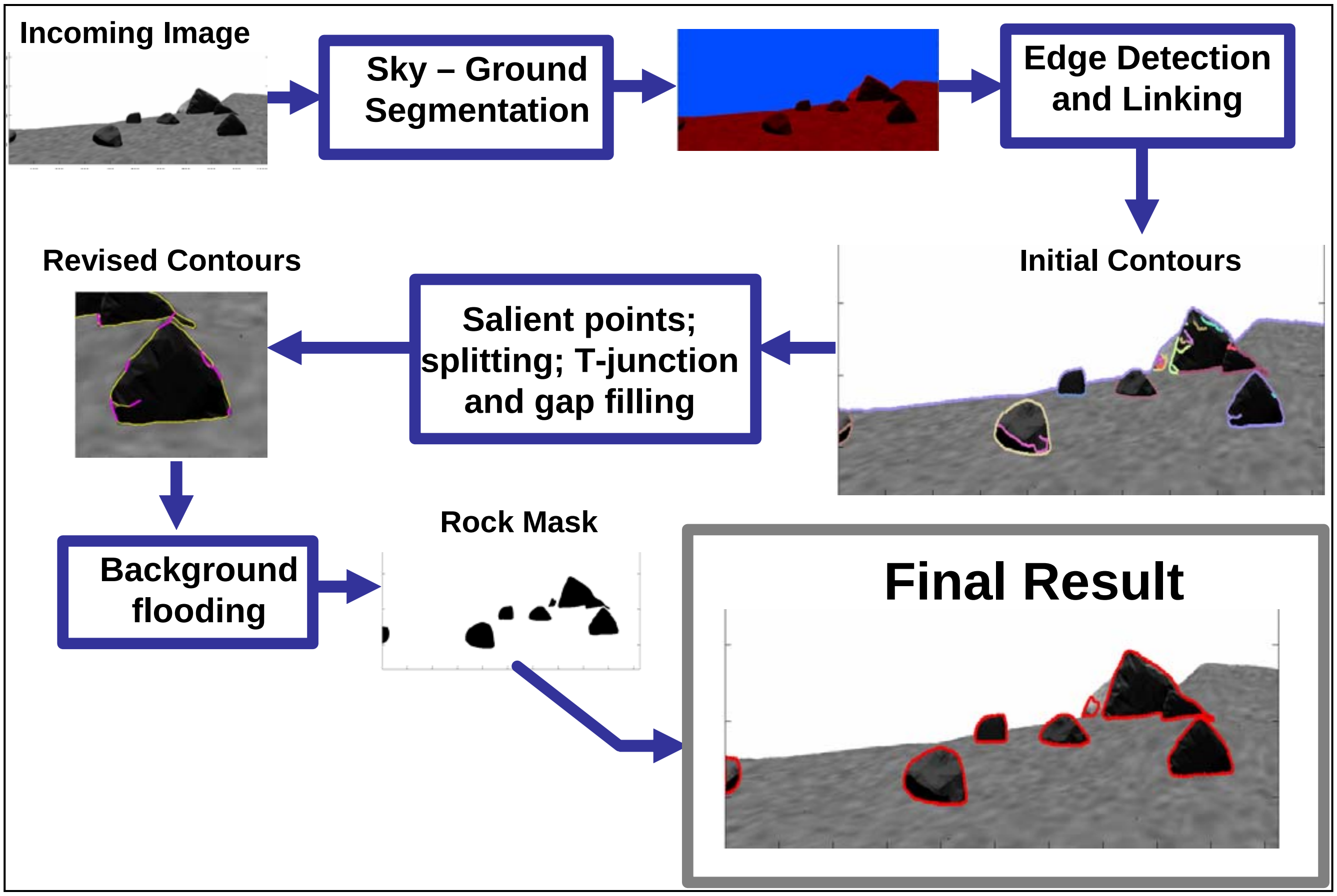

Figure 2 High-level view of the steps involved in the Rockster algorithm.

4. Multiple Rock Finders and Setting

Rock Finder Parameters

The OASIS system is designed with flexibility to allow it to be deployed on different rover platforms. In particular, we have used OASIS on the FIDO prototype rover and on the ROAMS (Rover Analysis, Modeling, and Simulation)[4] high-fidelity simulated rover. In both of these deployments, OASIS analysis algorithms were used on images obtained from hazard avoidance cameras and navigation cameras. Each type of camera, and each rover deployment, has differences in imaging and terrain characteristics that require minor parameter tuning.

There are two main types of parameters that we adjust, depending on the rover and cameras: crop boundaries and stereo vision settings. The crop boundaries are used to cut out regions of the image in which we do not want to find rocks. One example of this is to address artifacts in the image rectification process. As part of image processing, the images are rectified to remove radial distortion and align the left and right images in preparation for stereo range finding. This rectification process introduces artifacts in the form of black boundaries on the edges of images. Cropping is used to remove these regions from the image. For images obtained from hazard avoidance cameras, we also use crop boundaries to exclude the upper region of the image in which stereo quality is typically reduced.

The second type of parameter settings has to do with tuning stereo vision processing to deal with differences in hazard avoidance cameras and navigation cameras as well as the difference in terrain texture in the ROAMS simulation as compared with the JPL Mars Yard. The stereo vision settings include pyramid level, blob size and window sizes.

In addition to tuning parameters, we also support the application of different rock finding algorithms in different settings. In certain situations we want a more conservative classification of rocks in an image. That is, we want to minimize the number of false rocks that are identified. For example, when selecting rocks for autonomous targeting, it is important that we have high confidence in selecting areas of the image that actually contain rock. We have found the RockFinder algorithm [12] to have a low occurrence of false positives and, thus, we apply it for this type of

4

situation. For other applications, such as data summarization and novelty detection, we are interested in processing a large number of rock candidates to reduce the chance of missing something of particular interest. In this situation, we are willing to allow more false rocks in order to collect a wider range of input for novelty detection. We have found the Rockster algorithm to be well-suited for this type of application.

5. AUTONOMOUS TARGET SELECTION AND SAMPLING

Another application of the OASIS system is to provide autonomous targeting for science measurements that cannot be easily selected in advance. A number of rover remote sensing instruments have a very narrow field-of-view and thus require selection of specific focused targets for sampling. Such instruments include the MER Mini-Thermal Emission Spectrometer (mini-TES), the Mars Science Laboratory (MSL) ChemCam spectrometer, which performs Laser-Induced Breakdown spectroscopy (and is planned for launch in 2009), and infrared point spectrometers. Targeting these instruments by mission personnel on Earth currently requires a lengthy process. The typical scenario for selecting targets is to manually identify the targets using data that has already been downloaded on a previous sol. Thus, after reaching an end-of-day location, the rover must sit and wait until images can be analyzed and new measurement commands are uplinked (which at best will happen on the next sol).

By analyzing image data onboard, OASIS can autonomously select targets for these instruments and execute a set of measurement activities that do not exceed current rover resources. These techniques could be used, for example, on the MSL mission to select targets for the ChemCam instrument to sample (Figure 3).

To select potential targets during field testing, an image panorama of the surrounding area is taken. Using the RockFinder algorithm, rocks in the scene are autonomously identified. Next, a subset of these rocks is selected and prioritized based on scientist-selected criteria using the OASIS target signature algorithm. For instance, the target signature could be based on rock shape, albedo or size. Once a subset of target rocks is identified, points on these rocks are selected as specific instrument targets. Targets can also be prioritized based on their proximity or closeness to a pre-specified set of target features. For example, if light albedo rocks are preferred, the lightest colored rocks would be assigned the highest priority for sampling.

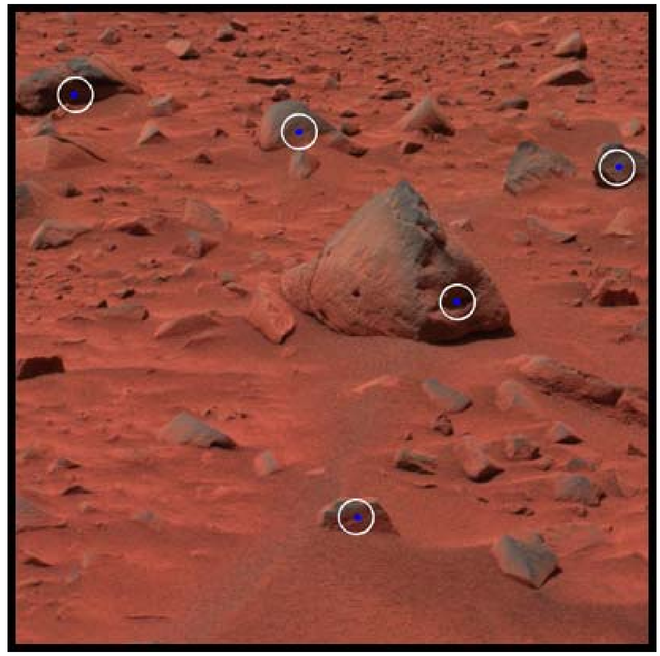

Figure 3: OASIS selects five potential targets for the ChemCam instrument to sample from an image taken on the MER mission. Autonomously selecting targets vs. blind sampling greatly increases the chances of accurately targeting a rock.

Since only a limited amount of time or resources may be available to take these measurements, the OASIS planning and scheduling subsystem is used to schedule only measurements that can be safely executed based on the rover health and current state. Currently, an iterative optimization approach is used to schedule additional science measurements within the rover's current command sequence. In this approach, science measurements that have been selected by onboard data analysis are added to the plan and the planner attempts to accommodate each measurement. If not all additional measurements can be added due to resource or other constraints, the planner iteratively deletes the lower priority measurements first. Typical constraints include a limited time window within which to add end-of-day measurements, a limited amount of rover energy for that sol (which may have fluctuated from predicted levels due to activities earlier in the day), and a limited amount of onboard memory to store science data.

To test OASIS on this application, we have done considerable testing with the FIDO rover. For testing, the rover was commanded to drive to an end-of-day location, where it took a full (or partial) panorama with the rover navigation cameras. Images were then analyzed by OASIS to locate rocks and identify targets. For this test, OASIS used a target signature which gave preference to targeting larger rocks, as suggested by the MSL ChemCam PI. Measurement requests for a subset of identified rocks were then sent to the planning subsystem.

Once the new measurement requests are received, the OASIS planner attempts to schedule as many new measurements as possible given rover constraints. New measurements are executed with the FIDO panoramic cameras, which provide an example of a limited field-of-view (FOV) instrument capable of making high resolution measurements of very specific target areas. Figure 4 shows an example of a large rock identified in a navigation camera image and the resulting panoramic camera image. This scenario was run successfully with FIDO using various rock configurations and onboard resource levels.

6. OASIS AND MAESTRO

As a preliminary demonstration, some the capabilities of OASIS have been integrated with Maestro, the science visualization and planning tool in use on the Mars Exploration Rovers mission [14]. The integration of Maestro allows users to design plans for execution on the rover and to visualize collected data. The initial integration enables users to command aspects of the OASIS system. First, a command to turn autonomy on and off was added to the command set in Maestro. Second, the capability to select the criteria to indicate high priority rock properties for downlink prioritization was added. Finally, while the rover is capable of collecting more data than can be downlinked there must be a mechanism for collecting this data. In this demonstration, the capability to command the collection of additional images for analysis by the OASIS system was incorporated into Maestro. Users can specify the pattern and quantity of additional images for analysis. Thus, when a traverse is planned users specify whether or not to perform autonomous science during the traverse. If they choose to activate autonomous science, they can design a sequence of mast camera acquisitions to acquire data for autonomous science as well as specify the frequency at which the rover should run this sequence. When data from the traverse is returned, users can view not only the standard Maestro downlink products, but also the rockfinder results on downlinked images. For more discussion on experiments using OASIS in an operational setting with Maestro and the ROAMS system conducted as part of the SOOPS task see [9].

7. OASIS AND ROAMS

To support development and experimentation, we have deployed the OASIS system in the ROAMS rover simulation environment [1]. ROAMS provides high-fidelity modeling of rover dynamics and terrain environment. The integration with ROAMS has been performed at the level of

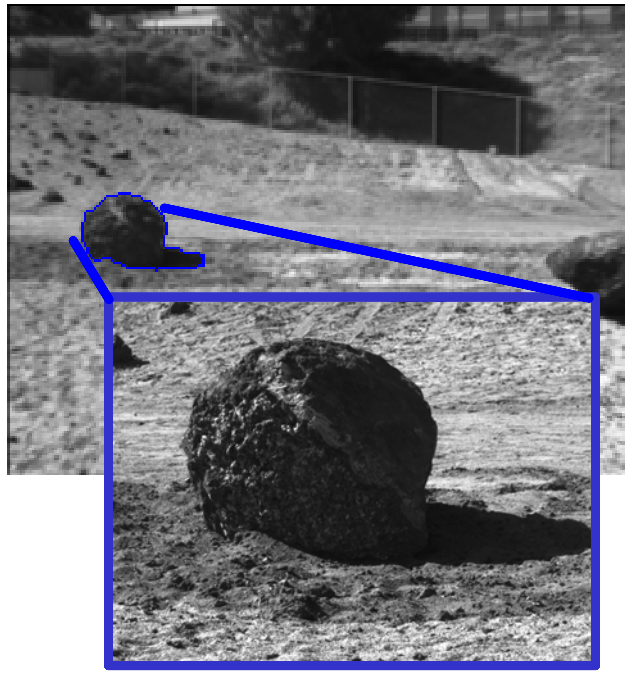

Figure 4: An example of a large rock autonomously selected and targeted by a limited FOV instrument on the FIDO rover (in this case, a high resolution panoramic camera).

CLARAty [10,11] actuators and sensors and, as a result, from a software standpoint deploying OASIS in ROAMS is identical to deploying OASIS on a physical rover. For example, OASIS uses the same interface code to command the rover wheels and mast and to acquire images from a camera whether on a simulated or physical rover. ROAMS has been a powerful tool in assisting with the development and debugging of the OASIS system. As described in Section 4, ROAMS has also supported experimentation to assess the impact autonomous science has on conducting surface operations for future missions.

8. ROCKIT

The OASIS Rock Identification Toolkit (RockIT) is a mature, cross-platform, graphical program originally designed to help geologists rapidly and accurately label rocks (or particles) in images. As images are labeled, RockIT reports both individual rock statistics and overall scene statistics. Golombek et al. [15] used RockIT to compare rock size distributions at several locations along the Spirit traverse. This past year, we have made improvements to RockIT so that it can serve as a visual front-end to the OASIS system. We should stress that RockIT is by no means required to use OASIS, but when available it provides a window into the inputs and outputs

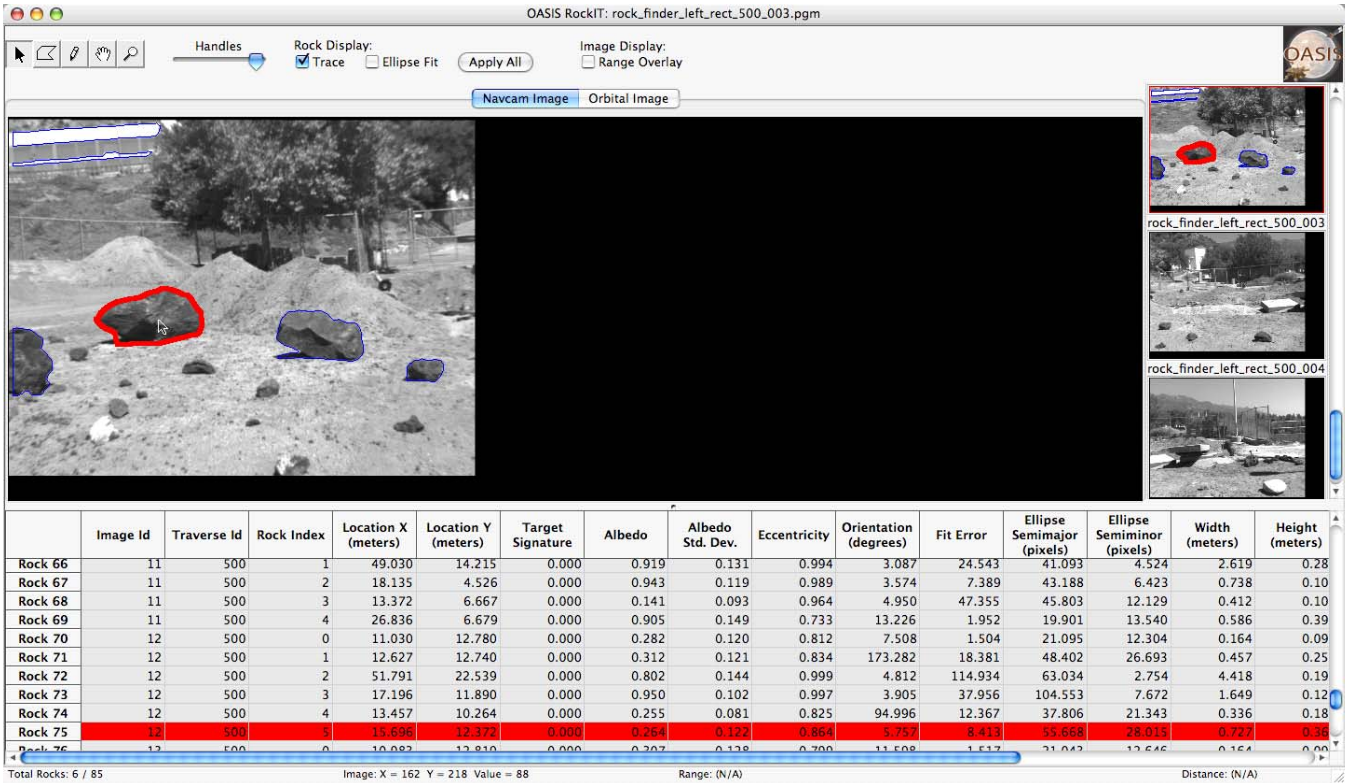

Figure 5. The OASIS Rock Identification Toolkit (RockIT) graphical user interface is comprised of four primary views: a main image display, a data summarization table, a scrollable list of thumbnail images (collected along a rover traverse), and an overhead map display (not shown; see Figure 6). An OASIS rock detection result is highlighted in the three views.

of the OASIS system. In this section, we provide a brief description of RockIT and how it supports OASIS capabilities.

The RockIT graphical user interface is comprised of four primary views: a main image display, a scrollable list of thumbnail images, a data summarization table, and an overhead map display. The main image display shows images, range data, and rock detections (Figure 5). When range data (e.g. derived from a stereo image pair) is available, RockIT will provide, as the mouse passes over each pixel, both the 3D (x, y, z) coordinates and the distance to that pixel in meters. For some parts of each image, range data may not be available, or range data beyond a certain distance may be too inaccurate. RockIT can mark these bad range regions with a translucent color (e.g., yellow) overlay. Rock detections are drawn on top of the composite image / range display. The combination of image data, range data, and detected rock detections forms a single dataset.

During a rover traverse, a sequence of datasets (images, range data, and rock detections) is collected. RockIT summarizes this sequence in its thumbnail view. The images of the traverse are rendered as image thumbnails and displayed in chronological order. Selecting a particular thumbnail will display the full image, as well as range data overlay and rock detections for that image in the main display area.

The data summarization view provides valuable features and statistics about each rock in the entire traverse. The statistics are derived by the underlying OASIS system, and while too numerous to describe in their entirety, include rock albedo, size, and shape features. For size, when range data is not present, 2D pixel area is reported. When range data is available, rock length and width are calculated. For shape statistics, RockIT performs fast, numerically stable, least squares ellipse fits [16] to each trace and reports the major and minor axes, eccentricity, and orientation of the fit ellipse. OASIS also reports the (x, y, z) location of each rock, which is useful for constructing an overhead map of the rover's world.

The 'overhead map' view provides a meaningful way to summarize a rover traverse in a single image (see Figure 6). The (x, y) location, albedo and size of each rock are used to create the map. Rocks on the maps are represented as colored ellipses. Each ellipse is centered at the (x, y) location of the rock (z is omitted for this simple 2D map). The ellipse major and minor axes are functions of the rock's size, and the ellipse is colored (filled) according to the rock's average albedo. In keeping with intuitive correspondence among views, selecting a rock in the

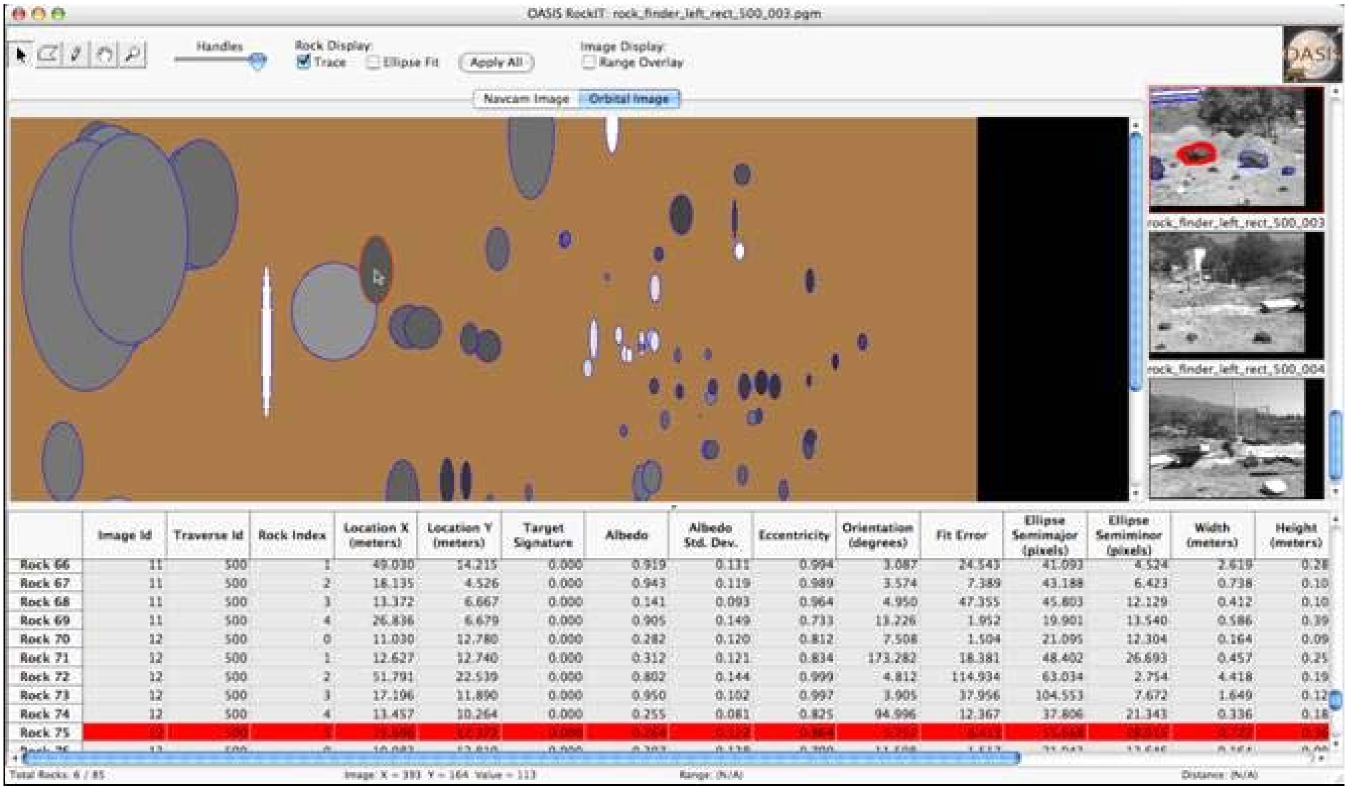

Figure 6. The same rock highlighted in Figure 5 is shown on the overhead map view. Rocks on the maps are represented as colored ellipses. Each ellipse is centered at the (x, y) location of the rock (z is omitted for this simple 2D map). The ellipse major and minor axes are functions of the rock's size, and the ellipse is colored (filled) according to the rock's average albedo.

overhead map or data summarization table will "jump" a user to the image containing that rock and highlight it.

The most recent RockIT development ties the four views together to simulate data downlink prioritization. RockIT can run various OASIS downlink prioritization schemes (including key target signature prioritization). For images that fall below the downlink prioritization threshold, RockIT covers the corresponding image and thumbnails. All the summary data provided by the table and overhead map, as well the user interaction and data linkage, still function as normal. The only change is that if a rock is selected in an image that did not meet the downlink criteria, a blank (black) image is displayed with a message that reads "Not Downlinked" (see Figure 7). We've found this simulated downlink feature to be particularly helpful in conveying the utility of OASIS to potential customers. Even if a particular image is not available, the data summarization table and overhead map provide a wealth of information about what was encountered during a traverse.

9. RELATED WORK

OASIS successfully synthesizes autonomous feature detection, geologic feature analysis, resource management, and execution on a physical rover testbed. It is unique in this high degree of autonomy and in the integration between planning, feature detection, and target selection. However, other previous projects have also investigated component technologies for autonomous science.

Researchers have developed several feature detection methods for terrestrial and Mars images. As early as 1999, preliminary OASIS feature extraction methods for textural image analysis and identification of key spectral signatures were being demonstrated [17], while Marsokod rover tests used shadow and edge features to identify rocks and strata in rover imagery [18]. This work also pioneered horizon detection for sky removal. The following year the Nomad Antarctic Meteorite Expedition (RAMS) demonstrated an autonomous search for meteorites on an Antarctic plateau [19 20] Here a robot platform used simple image segmentation to recognize rocks against glacial ice. Researchers have developed several rock detection and

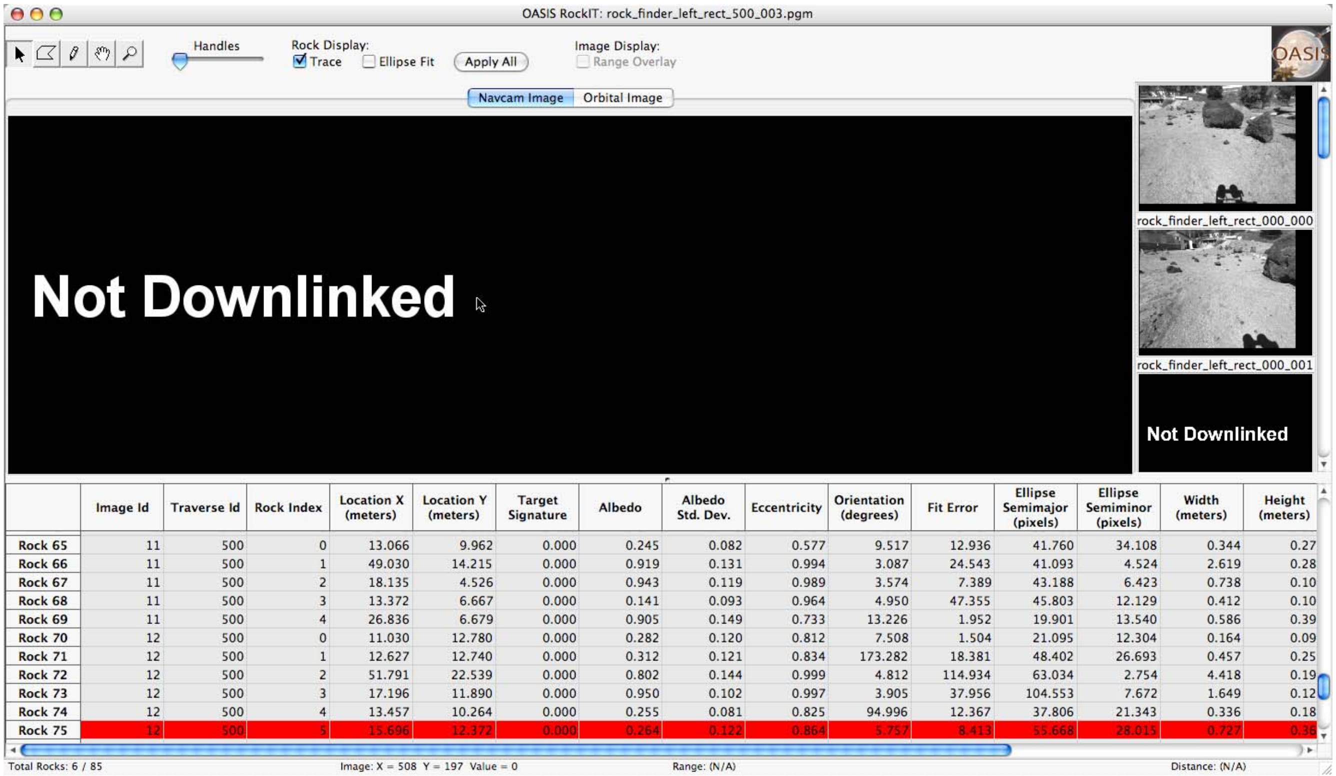

	Image Id	Traverse Id	Rock Index	Location X (meters)	Location Y (meters)	Target Signature	Albedo	Albedo Std. Dev.	Eccentricity	Orientation (degrees)	Fit Error	Ellipse Semimajor (pixels)	Ellipse Seminor (pixels)	Width (meters)	Height (meters)
Rock 65	11	500	0	13.066	9.962	0.000	0.245	0.082	0.577	9.517	12.936	41.760	34.108	0.344	0.27
Rock 66	11	500	1	49.030	14.215	0.000	0.919	0.131	0.994	3.087	24.543	41.093	4.524	2.619	0.28
Rock 67	11	500	2	18.135	4.526	0.000	0.943	0.119	0.989	3.574	7.389	43.188	6.423	0.738	0.10
Rock 68	11	500	3	13.372	6.667	0.000	0.141	0.093	0.964	4.950	47.355	45.803	12.129	0.412	0.10
Rock 69	11	500	4	26.836	6.679	0.000	0.905	0.149	0.733	13.226	1.952	19.901	13.540	0.586	0.39
Rock 70	12	500	0	11.030	12.780	0.000	0.282	0.120	0.812	7.508	1.504	21.095	12.304	0.164	0.09
Rock 71	12	500	1	12.627	12.740	0.000	0.312	0.121	0.834	173.282	18.381	48.402	26.693	0.457	0.25
Rock 72	12	500	2	51.791	22.539	0.000	0.802	0.144	0.999	4.812	114.934	63.034	2.754	4.418	0.19
Rock 73	12	500	3	17.196	11.890	0.000	0.950	0.102	0.997	3.905	37.956	104.553	7.672	1.649	0.12
Rock 74	12	500	4	13.457	10.264	0.000	0.255	0.081	0.825	94.996	12.367	37.806	21.343	0.336	0.18
Rock 75	12	500	5	15.696	12.372	0.000	0.264	0.122	0.864	5.757	8.413	55.668	28.015	0.727	0.36

Figure 7. RockIT can also simulate OASIS data downlink prioritization. The same rock highlighted in Figures 6 and 7 is now in an image whose downlink priority was too low. Even though this image is not available, the data summarization table and overhead map provide a wealth of information about this particular rock.

classification algorithms since these initial efforts. Some more recent work has employed stereo in 3D segmentation of rocks for instrument placement applications [21]. Machine learning-based approaches to detecting rocks in rover imagery have also been demonstrated [22].

Recent work provides evidence that automatic data analysis of detected features can extract meaningful geologic properties. The LITA Atacama rover expedition demonstrated that rock detection and classification strategies could characterize local surface geology in desert environments [23]. Laboratory studies show that under controlled conditions automatic feature extraction techniques already compare favorably to human geologists in quantifying some of rocks' visual attributes like angularity [24]. Autonomous classification of spectra for onboard analysis has also been demonstrated [25]. This growing body of research suggests that autonomous feature extraction can yield meaningful geologic analyses.

A number of planning and executive systems have been used for robotic applications. One approach directed towards rover command generation used a Contingent Planner/Scheduler (CPS) that was developed to schedule rover science operations using a Contingent Rover Language (CRL) [26]. CRL allows both temporal flexibility and contingency branches in rover command sequences. Contingent sequences are produced by the CPS planner and then are interpreted by an executive, which executes the final plan by choosing sequence branches based on current rover conditions. As compared to OASIS, only the executive is onboard the rover; planning is a ground-based operation and does not involve re-planning. Since only a limited number of contingencies can be anticipated and incorporated into the plan, CRL does not provide as much flexibility as OASIS when adjusting the sequence in response to unexpected events.

The LAAS-CNRS lab robotic control architecture [27] also uses onboard planning and execution to create initial plans and to provide re-planning capabilities. However, as compared to CASPER, the IxTeT planner uses a partial order CSP-based planning approach, which can require larger amounts of time for re-planning since a valid plan must be found at every search step. Further this system has not addressed the handling of opportunistic science and interacting with an onboard data analysis system.

Other similar approaches include Atlantis [28] and 3T [29], which used a deliberative planner and an executive on top of a set of reactive controllers. These approaches only use a batch planning approach in a limited fashion and do not provide online re-planning or support for opportunistic science.

The Autonomous Sciencecraft Experiment (ASE) [30] has demonstrated the capability of planning and data analysis systems to coordinate the behavior of the EO-1 Earth orbiting satellite. The Remote Agent Experiment (RAX) [31] demonstrated the ability of an AI planning and execution system to generate and execute plans onboard the NASA Deep Space One (DS1) spacecraft. RAX, however, used a batch approach to planning and could not dynamically re-plan. Further, since RAX and ASE were applied to spacecraft, neither handle the large uncertainty inherent in surface navigation and science.

10. SUMMARY

Onboard autonomy and science data analysis will have a significant impact on future landed missions. A number of elements of the OASIS system are being considered for use on current and future rover missions. The OASIS system is continuing to expand its capabilities for opportunistic science by increasing functionality in the feature extraction, data analysis and planning and scheduling components. Rock detection remains a non-trivial task, and will require continued investigation in the future. Future directions include using rock density maps and other tools to identify other scientifically interesting features such as geologic contact boundaries found during a traverse.

11. ACKNOWLEDGMENTS

The research described in this paper was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration. Specific funding for this work was provided by the Mars Technology Program, the DSMS Technology Program, and the Intelligent Systems Program.

Finally, our work this year could not proceed without the support and advice of scientists outside of the team, in particular: Albert Haldemann (JPL) and Matt Golombek (JPL).

REFERENCES

[1] A. Jain, J. Guineau, C. Lim, W. Lincoln, M. Pomerantz, G. Sohl, R. Steele, "ROAMS: Planetary Surface Rover Simulation Environment." In Proceedings International Symposium on Artificial Intelligence, Robotics and Automation in Space, i-SAIRAS 2003, Nara, Japan, May 19-23, 2003.

[2] R. Castano, T. Estlin, R. C. Anderson, D. M. Gaines, A. Castano, B. Bornstein, C. Chouinard, and M. Judd, "OASIS: Onboard autonomous science investigation system for opportunistic rover science," accepted for publication in *Journal of Field Robotics*.

[3] Rebecca Castano, Robert Anderson, Tara Estlin, Dennis Decoste, Forest Fisher, Daniel Gaines, Dominic Mazzoni, and Michele Judd, "Rover Traverse Science for Increased Mission Science Return," *Proceedings of the 2003 IEEE Aerospace Conference.* Big Sky, Montana, March 2003.

[4] Rebecca Castano, Michele Judd, Tara Estlin, Robert Anderson, Lucas Scharenbroich, Lin Song, Dan Gaines, Forest Fisher, Dominic Mazzoni, and Andres Castano, "Autonomous Onboard Traverse Science System," *Proceedings of the 2004 IEEE Aerospace Conference*, Big Sky, Montana, March 2004.

[5] R. Castano, M. Judd, T. Estlin, R. C. Anderson, D. Gaines, A. Castaño, B. Bornstein, T. Stough, and K. Wagstaff, "Current Results from a Rover Science Data Analysis System," *Proceedings of the 2005 IEEE Aerospace Conference*. Big Sky, Montana, March 2005.

[6] Rebecca Castano, Tara Estlin, Daniel Gaines, Andres Castano, Caroline Chouinard, Ben Bornstein, Robert C. Anderson, Steve Chien, Alex Fukunaga, and Michele Judd, "Opportunistic Rover Science: Finding and Reacting to Rocks, Clouds and Dust Devils," *Proceedings of the 2006 IEEE Aerospace Conference*, Big Sky, Montana, March 2006.

[7] A. Castano, Fukunaga, A., Biesiadecki, J., Neakrase, L., Whelley, P., Greeley, R., Lemmon, M., Castano, R., and Chien, S. "Autonomous detection of dust devils and clouds at Mars." To appear *Machine Vision and Applications,* 2007.

[8] M. Burl, "Rockster: Rock Segmentation through Edge Regrouping," Personal memo B-0120, Sept. 2006.

[9] A. F. C. Haldemann, M. McHenry, R. Petras, B. Bornstein, R. Castano, J. Cameron, T. Estlin, T. G. Farr, D. Gaines, A. Jain, M. Lee, C. Leff, C. Lim, I. Nesnas, M. Pomerantz, M. Powell, I. Shu, J. Wood, R. Volpe, "Simulation to Evaluate Autonomous Behaviors for Mobile Planetary Surface Science Missions," *Proceedings of the 2007 IEEE Aerospace Conference.* Big Sky, Montana, March 2007.

[10] I. A. Nesnas, A. Wright, M. Bajracharya, R. Simmons, T. Estlin, W. S. Kim, "CLARAty: An Architecture for Reusable Robotic Software," *SPIE Aerosense*

Conference, Orlando, Florida, April 2003.

[11] http://claraty.jpl.nasa.gov

[12] A. Castano, R. C. Anderson, R. Castano, T. Estlin, and M. Judd, "Intensity Based Rock Detection for Acquiring Onboard Rover Science," *Lunar and Planetary Science Conference*, Houston, TX, March 2004, abstract #2015.

[13] J. Canny, "A Computational Approach to Edge Detection", *IEEE Trans. Pattern Analysis and Machine Intelligence*, Vol. 8, pp. 679-698, 1986.

[14] J. Norris, M. Powell, M. Vona, P. Backes and J. Wick, "Mars Exploration Rover Operations with the Science Activity Planner." In *Proceedings of IEEE International Conference on Robotics and Automation (ICRA)*, 2005.

[15] M. P. Golombek, et al., "Assessment of Mars Exploration Rover Landing Site Predictions," *Nature*, vol. 436, #7, 44-48, 2005.

[16] A. W. Fitzgibbon, M. Pilu, and R. B. Fisher, "Direct Least Square Fitting of Ellipses," *IEEE Trans. on Pattern Analysis and Machine Intelligence*, Vol. 21, No. 5, pp. 476-480, 1999.

[17] M. Gilmore, R. Castano, T. Mann, R. C. Anderson, E. Mjolsness, R. Manduchi, and R. S. Saunders, "Strategies for autonomous rovers at Mars", in *Journal of Geophysical Research*, Vol. 105, No. E12, Dec. 2000 pp. 29223-29237.

[18] V. C. Gulick, R. L. Morris, M. A. Ruzon and T. L. Roush. "Autonomous image analysis during the 1999 Marsokhod rover field test", *J. Geophysical Research.* 104:E4, 2001. p. 7745-7764.

[19] L. Pedersen, "Robotic Rock Classification and Autonomous Exploration", *PhD Thesis, Robotics Institute, Carnegie Mellon University, CMU-RI-TR-01-14.* November 2000.

[20] M. D. Wagner, D. Apostolopoulos, K. Shillcutt, B. Shamah, R. G. Simmons and W. Whittaker. "The Science Autonomy System of the Nomad Robot", *IEEE International Conference on Robotics and Automation (ICRA)*, 2001 p. 1742-1749.

[21] L. Pedersen, "Science target assessment for Mars rover instrument deployment," *IEEE International Conference on Intelligent Robotics and Systems,* 2002. p.817-822.

[22] D. R. Thompson, S. Niekum, T. Smith and D. Wettergreen, "Automatic Detection and Classification of Geological Features of Interest", *Proceedings of the 2005 IEEE Aerospace Conference.* Big Sky, Montana, March 2005.

[23] D. R. Thompson, T. Smith and D. Wettergreen, "Data Mining During Rover Traverse: From Images to Geologic Signatures", *8th International Symposium on Artificial Intelligence, Robotics and Automation in Space (iSAIRAS)*, September 2005.

[24] H. Dunlop, "Automatic Rock Detection and Classification in Natural Scenes," *Masters Thesis, Carnegie Mellon University,* CMU-RI-TR-06-40. August 2006.

[25] B. Bornstein, R. Castaño, M. S. Gilmore, M. D. Merrill, J. P. Greenwood, "Creation and Testing of an Artificial Neural Network Based Carbonate Detector for Mars Rovers," *Proceedings of the 2005 IEEE Aerospace Conference,* Big Sky, Montana, March 2005.

[26] J. Bresina, K. Golden, D. Smith, and R. Washington, "Increased Flexibility and Robustness of Mars Rovers," Proceedings of the International Symposium, on AI, Robotics and Automation for Space, Noordwijk, The Netherlands, June 1999.

[27] R. Alami, R. Chautila, S. Fleury, M. Ghallab, and F. Ingrand, "An Architecture for Autonomy," *International Journal of Robotics Research*, 17(4) April, 1998.

[28] E. Gat, "Integrating planning and reacting in a heterogeneous asynchronous architecture for mobile robots," in *SIGART Bulletin 2*, 1991, 70-74.

[29] R. Bonasso, R. Firby, E. Gat, D. Kortenkamp, D. Miller, and M. Slack, "Experiences with an Architecture for Intelligent, Reactive Agents," *Journal of Experimental and Theoretical Artificial Intelligence Research,* 9(1), 1997.

[30] S. Chien, R. Sherwood, et al.,. "Using Autonomy Flight Software to Improve Science Return on Earth Observing One," *Journal of Aerospace Computing, Information, and Communication,* 2005.

[31] A. Jonsson, P. Morris, N. Muscettola, K. Rajan, and B. Smith, "Planning in Interplanetary Space: Theory and Practice, " *Proceedings of the Fifth International Conference on Artificial Intelligence Planning Systems,* Breckenridge, CO, April 2000.

BIOGRAPHY

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

Dr. Tara Estlin, *OASIS Planning and Scheduling team lead, senior member technical staff. Dr. Estlin is a member of the Artificial Intelligence Group at the Jet Propulsion Laboratory where she performs research and development of planning and scheduling systems for rover automation and multi-rover coordination. Dr. Estlin is currently the PI of the "An Onboard Scientist for Multi-Rover Scientific Exploration" and the "Integrated Resource and Path Planning" projects. These efforts are developing capabilities for onboard rover-command generation, resource planning and scheduling, and data analysis for single and multiple rovers. She is also a member of the "Robotic Autonomy Architecture" project and is a past member of the Long-Range Science Rover (LRSR) project, which developed the Rocky 7 and Rocky 8 rovers. Dr. Estlin has been at JPL since 1998. She received a B.S. in computer science in 1992 fromTulane University, an M.S. in computer science in 1994 and a Ph.D. in computer science in 1997, both from the University of Texas at Austin.*

Dr. Daniel Gaines *is a senior member of the Artificial Intelligence Group at JPL. His research interests are in integrated planning and execution and in machine learning to improve planning. Dr. Gaines received a Ph.D. in Computer Science from the University of Illinois at Urbana-Champaign. Before coming to JPL, Dr. Gaines was an Assistant Professor in Computer Science at Vanderbilt University. His work at JPL is primarily focuses on planning and execution techniques for planetary exploration rovers.*

Caroline Chouinard *is a member of the Artificial Intelligence Group at the Jet Propulsion Laboratory in Pasadena, California where she performs software development for planning and scheduling systems. She has contributed to the many stages of the rover's decision-making life-cycle, including initial plan generation, scheduling, execution, plan optimization and re-planning for events. She received a B.B.A in International Business and Finance in 1996 and a B.A. in Computer Science in 2000 both from the University of Texas at Austin. Her current interests include cognitive science, artificial intelligence, and graphics.*

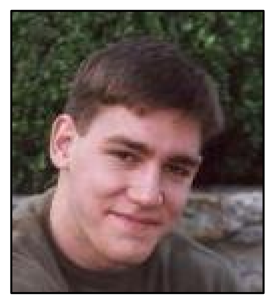

Ben Bornstein *is a research programmer in the Machine Learning Systems group at the Jet Propulsion Laboratory in Pasadena, CA. He enjoys bringing machine learning techniques and considerable hacking (programming) skills to bear to solve problems in geology, bioinformatics, and systems biology. He has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

Dr. Robert C. Anderson, *OASIS Science team lead, senior member technical staff. Dr. Anderson attended Old Dominion University in Norfolk, Virginia, where he received his Bachelor of Science degree in geology in 1979. In 1985, he received a Master of Science from Old Dominion University in geology with an emphasis on structural geology and mapping tectonic features surrounding the Tharsis region of Mars. In 1995, he received a Doctor of Philosophy from the University of Pittsburgh in geology with an emphasis on visible and near infra-red remote sensing. His Ph.D. research was centering on mapping Quaternary surfaces and soils around the Whipple Mountains of southwestern Arizona. Dr. Anderson worked on the successful Mars Pathfinder Project as science support for the Mineralogy and Geochemistry Science Operations Group. Currently Dr. Anderson is the Investigation Scientist for the Rock Abrasion Tool (RAT) and science support for Mission Operations on the Mars 03 mission. He works closely with the FIDO rover team and is presently the Science Team lead on the Onboard Autonomous Science Investigation System (OASIS) project. Dr. Anderson's research is centered on unraveling the geologic history of Mars.*

Dr. Michael C. Burl *is a Principal Member of the Technical Staff at theJet Propulsion Laboratory in the Machine Learning and Instrument Autonomy Group. He earned his Ph.D. in Electrical Engineering from Caltech in 1997 and is well-known for his work in applied machine learning, computer vision, data mining, and signal processing. He has served on the Program Committees for a number of top-tier computer vision and data mining conferences.*

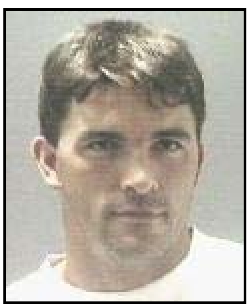

David R. Thompson *is a PhD candidate at the Robotics Institute at Carnegie Mellon University. He received a MSc. in Informatics from the University of Edinburgh. His current work involves computer vision and machine learing technologies for autonomous science. Recent projects include autonomous site characterization and geologic mapping during rover operations in the Atacama Desert of Chile.*

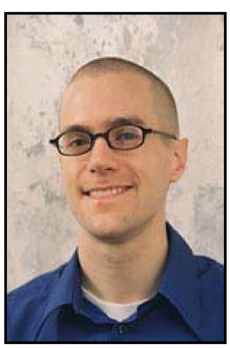

Dr. Andres Castano *received his B.S. and M.E.E. degrees from the University of Los Andes, Bogota, Colombia, and his M.S. and Ph.D. degrees from the University of Illinois at Urbana-Champaign, all in Electrical Engineering. At Urbana-Champaign he worked at the Beckman Institute in visual servo-control of manipulators and omni-directional cameras. In 1998 he joined the Information Sciences Institute at the University of Southern California where he was the lead roboticist of the Conro project on reconfigurable robots. In 2000 he joined the Machine Vision Group at the Jet Propulsion Laboratory where he currently works in feature detection and recognition. His research interests are robotics, computer vision and graph theory.*

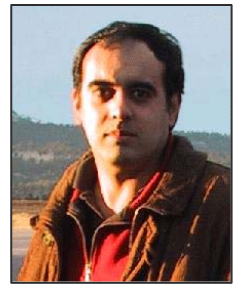

Michele Judd, P.E., *OASIS Research Manager and senior technical staff member of JPL's Science Division. Ms. Judd received her B.S. in Petroleum Engineering from Stanford University and her M.L.A. in Organizational Development from Southern Methodist University. A nationally certified Quality Manager and licensed professional engineer, Ms. Judd has over 16 years of experience in project management, both within JPL and for Mobil Oil.*