

Current Results from a Rover Science Data Analysis System

Rebecca Castano, Michele Judd, Tara Estlin, Robert C. Anderson,
Daniel Gaines, Andres Castaño, Ben Bornstein, Tim Stough, and Kiri Wagstaff

Jet Propulsion Laboratory
California Institute of Technology
Pasadena, CA 91109
(818) 393-5344
Rebecca.Castano@jpl.nasa.gov

Abstract— The Onboard Autonomous Science Investigation System (OASIS) evaluates geologic data gathered by a planetary rover. This analysis is used to prioritize the data for transmission, so that the data with the highest science value is transmitted to Earth. In addition, the onboard analysis results are used to identify science opportunities. A planning and scheduling component of the system enables the rover to take advantage of the identified science opportunity. OASIS is a NASA-funded research project that is currently being tested on the FIDO rover at JPL for the use on future missions.

In this paper¹, we provide a brief overview of the OASIS system, and then describe our recent successes in integrating with and using rover hardware. OASIS currently works in a closed loop fashion with onboard control software (e.g., navigation and vision) and has the ability to autonomously perform the following sequence of steps: analyze gray scale images to find rocks, extract the properties of the rocks, identify rocks of interest, retask the rover to take additional imagery of the identified target and then allow the rover to continue on its original mission.

We also describe the early 2004 ground test validation of specific OASIS components on selected Mars Exploration Rover (MER) images. These components include the rock-finding algorithm, RockIT, and the rock size feature extraction code. Our team also developed the RockIT GUI, an interface that allows users to easily visualize and modify the rock-finder results. This interface has allowed us to conduct preliminary testing and validation of the rock-finder's performance.

TABLE OF CONTENTS

1. INTRODUCTION	1
2. OASIS SYSTEM OVERVIEW	4
3. HARDWARE INTEGRATION	5
4. ROCKIT RESULTS.....	7
5. CONCLUSIONS AND FUTURE WORK	7
6. ACKNOWLEDGEMENTS.....	8
REFERENCES.....	8
BIOGRAPHY	9

1. INTRODUCTION

One of the first objectives discussed in *A Renewed Spirit of Discovery: The President's Vision for U.S. Space Exploration*², is to “Implement a sustained and affordable human and robotic program to explore the solar system and beyond.” A major component of the in-situ, robotic exploration program will include the effective utilization of rovers on planetary surfaces.

In-situ exploration of a new rocky planet (or satellite) generally begins with the study of its geology, as rocks are excellent recorders of planetary history. While the 1976 Viking mission landers offered striking insights into the history of Mars, their geologic investigation capabilities were limited to the length of their robotic arms (see Figure 1), and the view of their cameras.

In 1997, the Mars Pathfinder mission provided the first rover, Sojourner, to ever travel upon another planetary surface. The farthest radial distance that Sojourner traveled (see Figure 2) was a little over 12 meters on Sol 74, and the total distance traveled during its lifetime was about 100 meters³. More importantly for the following discussion, the longest traverse Sojourner made in one sol (a Martian day)

² Released in January 2004. An electronic version is available at: http://www.whitehouse.gov/space/renewed_spirit.pdf

³ See <http://mars4.jpl.nasa.gov/MPF/rovercom/rovfaq.html#faq1> for a table listing Sojourner's technical accomplishments.

¹ 0-7803-8870-4/05/\$20.00© 2005 IEEE, IEEEAC paper #1212, Version 3, Updated December 23, 2004

was a little over 7 meters on Sol 32. The last signal received from the rover was on Sol 83, signaling the end of

a mission that lasted 12 times longer than its design lifetime of seven days (<http://solarsystem.nasa.gov/missions>).

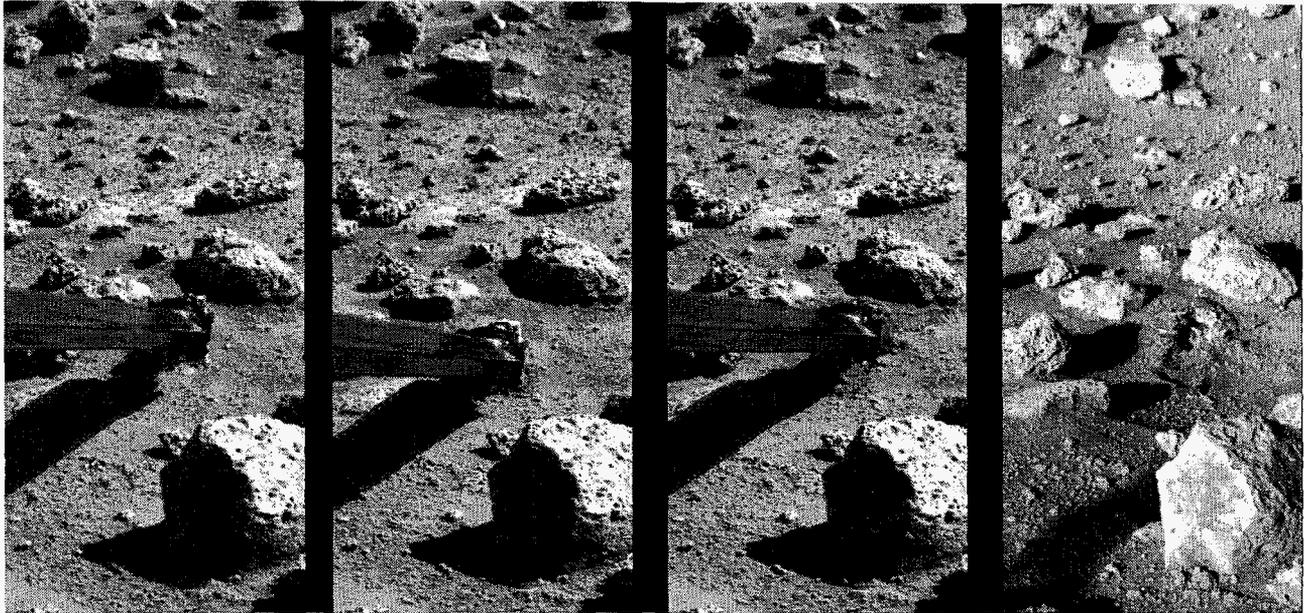


Figure 1: The robotic arm of the Viking 2 Lander extends to collect a sample of soil for analysis. Credit: NASA. JPL Planetary Photojournal ID number: PIA00145.

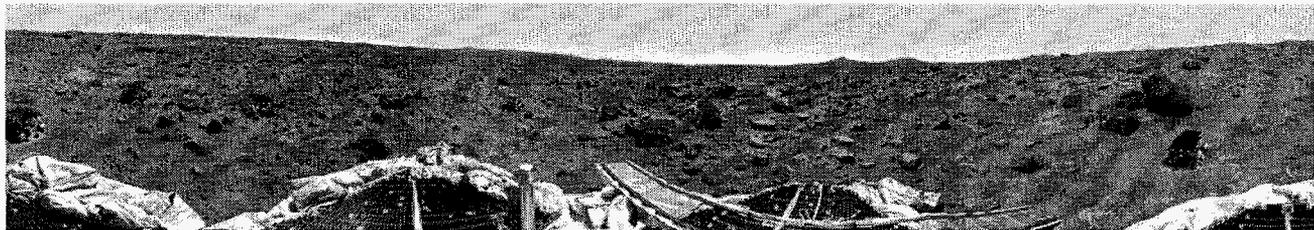


Figure 2: Various images of the Sojourner rover shot by the Pathfinder cameras have been placed into this panorama. Since the camera's position was consistent, it is possible to see these images of the rover in the context of the entire landscape. This provides a visual, snapshot record of the rover's travels. Panorama made available by Dr. Carol Stoker, NASA AMES.

The 2004 Mars Exploration Rovers (MER), Spirit and Opportunity, are currently acting as geological assistants to the remote human scientists back on Earth. They are directed to travel to a science site and are then instructed to visit rocks of interest and perform various scientific measurements on them. This is a slow and tedious process due to round-trip communication time, transmission bandwidth, and transmission opportunity constraints.

On April 17, 2004, Opportunity set a new one-sol driving record of 140.9 meters (462 feet) on Sol 82 of its extraordinary mission. This record is unlikely to survive the end of the MER mission. As of September 3, 2004, Spirit

had traveled over 3,600 meters (see Figure 3) and had already survived 148 sols past its 90 day design lifetime.

This hardware durability and increased mobility has forever changed the expectations and future requirements for rover operations. Long, extended missions make it expensive and difficult to maintain a fully staffed operations team. Besides cost, other concerns that lengthy in-situ missions raise include the physical and mental toll on the scientists, engineers and their families.

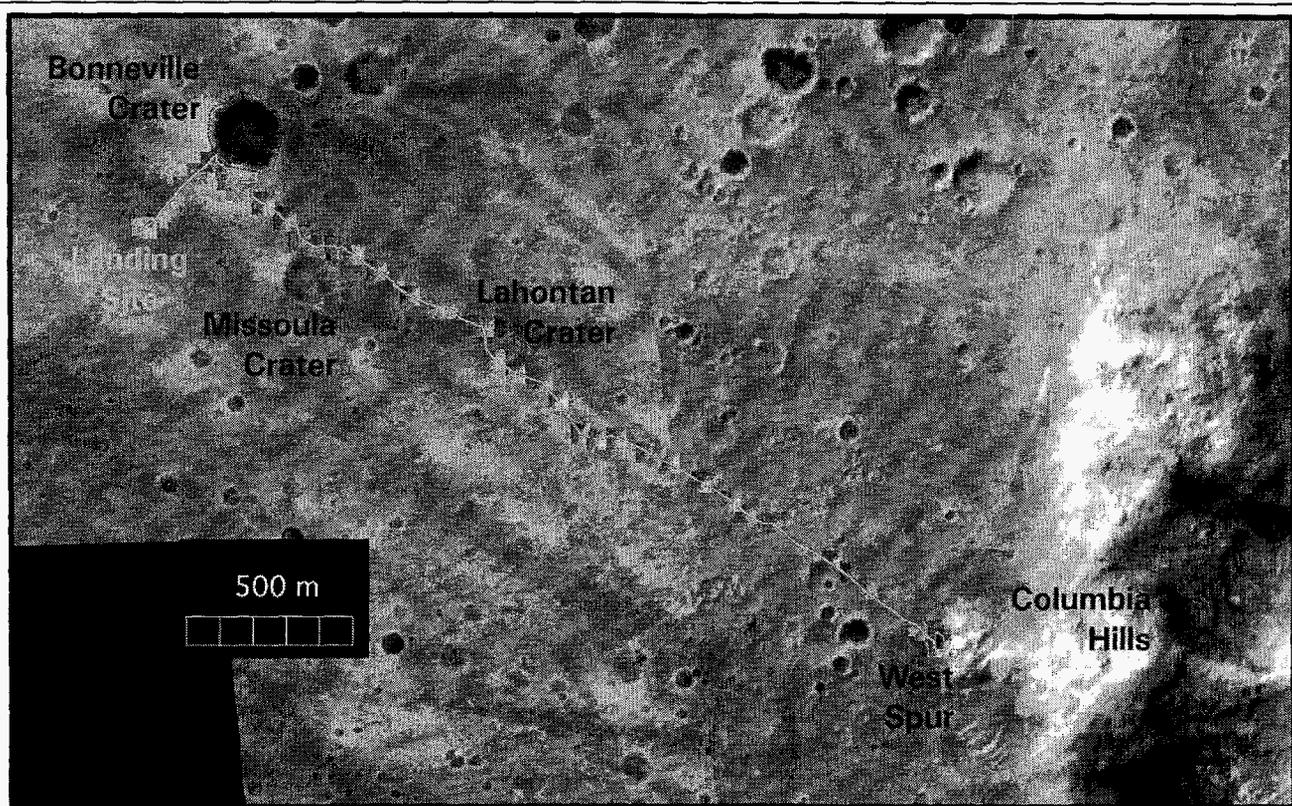


Figure 3 This map shows the 3,622.7 meter traverse of NASA's Mars Exploration Rover Spirit through the rover's 238th martian day, or sol (Sept. 3, 2004). The background image consists of frames from the Mars Orbiter Camera on NASA's Mars Global Surveyor orbiter. Inset images along the route are from Spirit's navigation camera. From its landing site, Spirit drove up to the rim of "Bonneville" crater on the far left, then to the north rim of "Missoula" crater, followed by a long drive across the plains to the "West Spur." The scale bar at lower left is 500 meters (1,640 feet). North is up. Credit: NASA/JPL. JPL image number PIA06872.

Scientists are starting to acknowledge the benefit that autonomous operations might bring to the table and are now requesting "smarter" rovers with a greater capability for autonomy for future missions⁴.

There are a number of autonomous rover capabilities currently in development for future in-situ missions. One key capability, autonomous onboard science, continues to grow in importance as rover travel distances continue to dramatically increase. OASIS [1, 2], 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 during a long traverse. It analyzes geologic 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. 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. Once it has completed this additional measurement, the rover reverts back to its original plan and continues on its traverse.

This science alert scenario has now been effectively demonstrated in a closed loop fashion with control software

⁴ A recent example is Dr. Steve Squyres (Athena Science Payload Principal Investigator for the MER mission) at the AAS conference in Pasadena, California (11/16/2004) during his Carl Sagan Memorial Award Lecture.

(e.g., navigation and vision) that is already onboard. This successful demonstration of OASIS, however, did not come without a price. A large investment of time and troubleshooting had to be supplied to integrate OASIS with the hardware and the hazard avoidance, position estimation and control software components. This paper will outline lessons learned from this experience.

Integration of OASIS onto a rover is a key step toward validating the technology for a future mission; however, it is just one step forward in a very long validation journey. Another validation effort under way is the rigorous performance testing of RockIT, the OASIS rock-finding algorithm.

During the validation of RockIT, it became clear that some type of graphical user interface (GUI) was needed for scientists to interact with RockIT and begin their evaluation of the software. OASIS developed a GUI that permitted scientists to quickly visualize and edit the rock-finder results. This new GUI allowed the scientists and the OASIS team to conduct preliminary testing and validation of the rock-finder's performance.

RockIT is being evaluated on how accurately it finds rocks in grayscale images using the three different performance criteria detailed in this paper. RockIT is tested against images that are taken from a number of various images sets – Mars Yard (which is a test environment that is similar to Mars) images taken by the FIDO rover at JPL, field test images taken by FIDO in the Arizona desert, and microscopic imaging data that have been transmitted to Earth by the MER rovers. Once the rocks (or grains in the case of the microscopic imager) have been found, they are sent to a rock size feature extraction code that can quickly output rock size distribution data for geologists.

In the past year, significant progress has been made on advancing the robustness of the OASIS system; to put this progress in perspective, however, an overview of the OASIS system and its desired autonomous capabilities is required.

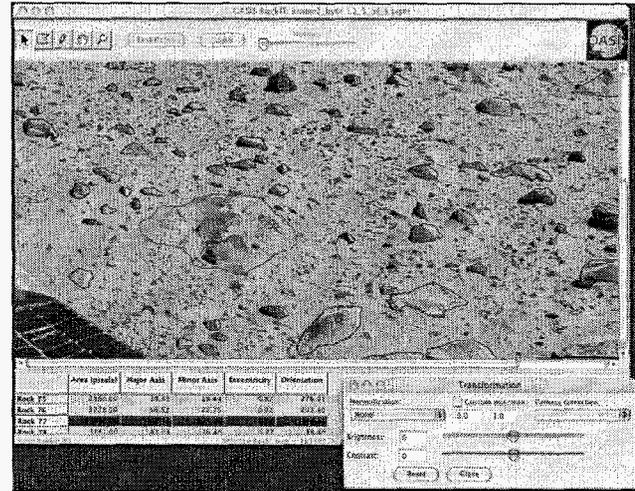


Figure 5 Screenshot from RockIT GUI used to provide manual, hand labeling (ground truth) to compare against the rock detector results.

2. OASIS SYSTEM OVERVIEW

To assess and subsequently prioritize the scientific value of a set of collected grayscale images, we must first extract the information found within the images. A geologist in the field gets information about a site by identifying geologic features including the albedo (fraction of incident light that is diffusely reflected from a surface – essentially the brightness of the rock in the image), texture, shape, size, color, and arrangement of rocks, and features of the topography such as layers in a cliff face. The geologist analyzes and assesses this data, and then takes some action based on the analysis, such as taking a sample or taking some additional measurement of an interesting rock.

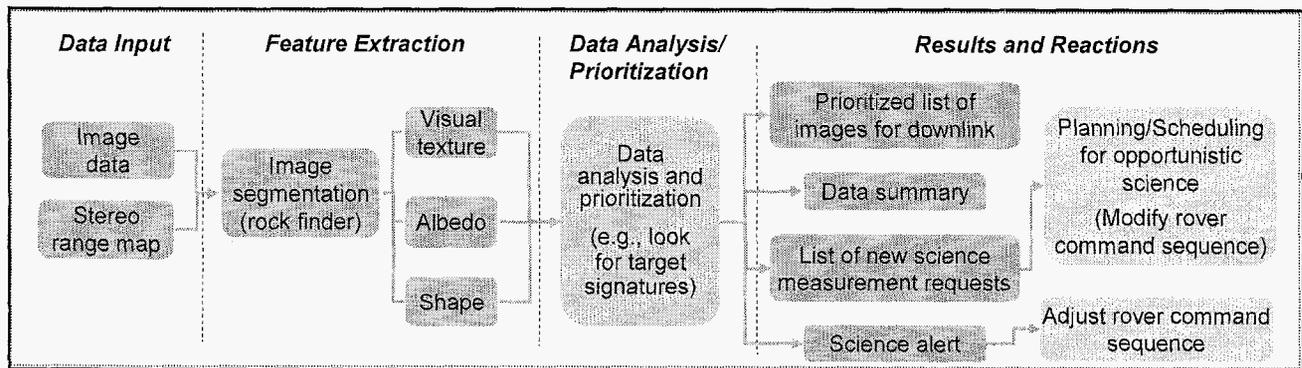


Figure 4 Overview of OASIS System.

In order for scientists to allow an autonomous system to help investigate the traversed region, the system must be able to perform, albeit in a very simple way, these same types of functions. This system thus acts as a geologist's assistant who helps point out rocks of interest to the geologist.

There are three major components that comprise OASIS:

- **Extract Features from Images:** Enables extraction of features of interest from collected images of the surrounding terrain. This module both locates rocks in the images and extracts rock properties (features) including shape, texture and albedo.
- **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 prioritization algorithms that analyze the collected data and prioritize the rocks. The results from these three algorithms are then fed into a unified prioritization algorithm that provides two downlink products: a prioritized list of images for downlink and a table that summarizes all of the rocks found on the traverse. A new set of observation goals is also generated to gather further data on rocks that either conform to the pre-set specifications by 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. HARDWARE INTEGRATION

Integration with CLARAty

To test OASIS with rover hardware, our approach has been to integrate with the CLARAty robotic architecture [3]. CLARAty is a unified and reusable robotic software architecture that simplifies the integration of new technologies onto different robotic platforms. For this testing, the OASIS rockfinder and planning and scheduling software were directly integrated with CLARAty and handled interaction with other key CLARAty elements that were required to run with rover hardware. Other pieces of OASIS (such as feature extraction) are planned for

integration in later years. Through this integration, the OASIS system interacted with several technologies provided by CLARAty, which operate on-board the rover. These software components are locomotion, position estimation, navigation and stereo vision processing.

The locomotion module coordinates a set of motors to steer and drive the rover's wheels, enabling the rover to move across the terrain. The FIDO rover, which was used for current development and testing of OASIS, is a six wheeled vehicle with all wheel steering and driving enabling the rover to turn in place, drive in straight lines and in arcs and perform crab maneuvers in which the rover moves sideways while facing forward.

As the rover traverses across the landscape, position estimation software attempts to estimate the rover's current location and orientation relative to a global reference frame. Varying sand texture results in different amounts of wheel sinkage, which in turn, results in variance in the effective wheel radius of the rover. In addition, the rover often encounters rocks of varying sizes, shapes and textures which results in wheel slippage as it drives over the rocks. The position estimation algorithm in the current system uses an IMU (Inertial Measuring Unit) to estimate rover attitude (roll, pitch, and heading) and wheel odometry to estimate linear velocity. The IMU helps detect slippage that could not be detected with wheel odometry alone.

OASIS uses the estimated position to determine the locations of rocks in the world so that, if an interesting rock is identified, a science alert can be created with the appropriate coordinates. Estimating rover position is a challenging problem given the nature of the terrain. While position estimation error can be quite difficult to deal with for many aspects of autonomous software, in practice, the OASIS system has been found to be highly robust to this type of error. Typically, estimation error gradually accumulates as the rover traverses. OASIS can usually detect rocks in an image, identify rock locations and alert the planner of interesting rocks before significant position estimation error has accumulated. Therefore, when the planner acts on the science alert, even if the global position estimate is off, the current position estimate is highly accurate relative to the estimated locations of the interesting rock.

Navigation software is used to guide the rover to a goal location while avoiding obstacles. OASIS uses the Morphin navigation system developed at Carnegie Mellon University (CMU) [4]. As part of the hazard avoidance algorithm, the navigator acquires stereo images from the rover hazard cameras to detect obstacles in the rover's path. The front hazcam images (as opposed to images from the rear hazcams) are also passed to the OASIS rock detection process, running in a separate thread.

Both the navigator and the OASIS rock detection algorithm make use of the JPL stereo vision package to acquire range data for the hazard camera images. The navigator uses stereo processing to generate a point cloud representing the terrain around the rover. OASIS uses range data to determine the locations of identified rocks. Rocks are first located in a coordinate frame relative to the center of the rover and then translated into the global reference frame using the current position estimate.

OASIS component run times on FIDO

As part of our testing, we gathered preliminary statistics on the run time of key OASIS components. While we have not spent significant time optimizing the performance of these components, the numbers provide a general idea of current performance and provide a reference to track future improvements. In a representative run, the OASIS rockfinder processed 11 hazcam images resulting in 5 science alerts being sent to the planner. The rockfinder was run on a 233 Mhz Pentium processor running VxWorks 5.5 with 128 MB of RAM. Rockfinder took an average of 53 seconds to process each image and found about 9 rocks per image with a total of 103 rocks being identified for the 11 images. The OASIS feature extraction component was run on a 930 Mhz Pentium processor running Linux 2.4 with 256 MB of RAM. For the current tests, albedo and shape information were extracted for each rock. For the 11 images processed, feature extraction averaged 0.5 seconds per image. The planning and scheduling component ran on a 2.5 GHz Pentium processor running Linux 2.4 with 1 GB of RAM. In handling the 5 science alerts, the planner spent an average of 6 seconds generating a plan for each alert.

Lessons learned

Integrating and testing with hardware and required control software (e.g., for navigation) introduced a number of challenges. In this section, we will briefly outline a few of the key lessons learned during this process. One primary challenge was that the testing and demonstration of our onboard science system required a number of underlying components to properly perform and support our software. An important lesson for running testing with rover hardware is that a significant amount of time must be allocated to tune and test supporting software and hardware components. Though the majority of supporting components were provided through CLARAty, many had not been run through extensive testing in the environment that we wished to demonstrate OASIS. In particular, many of the elements had been individually tested but not tested working as a completely integrated system.

Some key components that we spent a large amount of time testing for our use include the FIDO rover cameras, navigation software and position estimation software. This process was further complicated when supporting software had large numbers of tunable parameters, such as the

Morphin navigator, which has over 40 user-settable parameters. And since some supporting software had only been tested in certain environments or exercised in unit testing, our testing process also uncovered several problems (or bugs) in supporting software that had not been previously discovered. In the future, we hope to have more access to the people who developed the supporting technology and have direct support in correctly tuning their application for our needs. However, since this will not always be the case, we expect future work with hardware testing to consistently require time for tuning supporting technology.

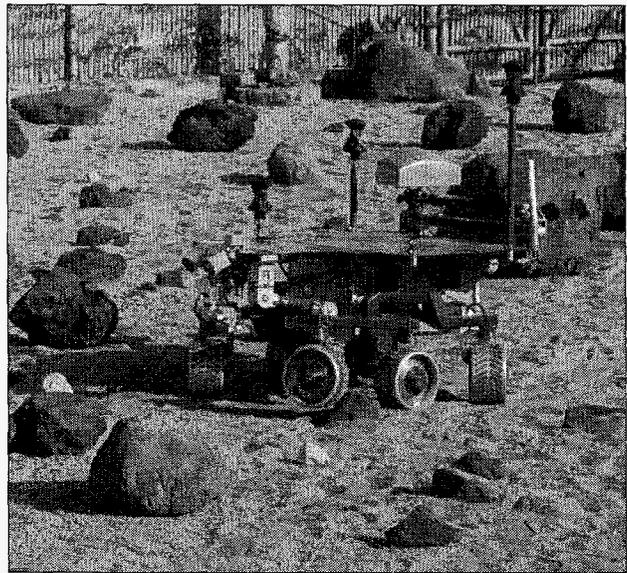


Figure 6 FIDO Rover in Mars Yard at JPL, October 2004. OASIS has identified an interesting rock for the rover to investigate (see white rock directly in front of FIDO). The green flags show the human observers where the regularly-scheduled science events should occur.

Another lesson learned was the value of performing a large number of tests in simulation, even with a relatively simple simulator. To test the ability of the planning system to handle science alerts under varying conditions, we invested in a testing infrastructure that allowed large numbers of tests to be run offline using a simple simulator. This simulator tracked items such as resource usage and rover position, but did not simulate higher fidelity items such as obstacle avoidance or stereo image processing. During these tests certain factors were varied such as the number of science alerts, available resources, rover speed, etc. This allowed the planner to be exercised on a large variety of situations and allowed many bugs to be caught before hardware testing was even performed. Furthermore, problematic situations that were discovered in hardware testing could often be easily reproduced and debugged by using a simulator. Simulated testing also provided an easy

way to run large numbers of tests that would have been impossible to run exclusively on hardware, due to the time required to complete them. Overall, our testing in simulation made our system significantly more robust to variations in resource usage, rover position, activity execution time, and frequency of science alerts. Such robustness is particularly important when handling rover operations since factors such as traverse time, power usage and even possible science opportunities are difficult to predict. We should note that although testing using a simple simulator was very valuable, it would have also been valuable to perform some testing with higher fidelity simulation. For instance, when testing with rover hardware, obstacle avoidance software often performed differently depending on variations in rover position, sun angle, etc. Since we could not easily re-produce this behavior in our simple simulator, we had to rely on our hardware tests to ensure this behavior did not cause problems or unexpected situations for the OASIS software.

4. ROCKIT RESULTS

Performance Metrics

The performance of the rock finding algorithm is measured by comparing the results of the algorithm to hand labeled images, where the rocks have been manually identified. Two types of comparison have been used in scoring the labels returned by the rock finder. Candidate rocks produced by the rock finder are compared with the hand labeled rocks based on either the distance between the centers or the number of overlapping pixels.

Rock labels produced by the rock detector are automatically scored against hand labeled data and a web page showing the results is automatically generated. This web interface allows browsing of the images, hand labeled rocks, and the detected rocks as well as summary statistics for a set of images.

Scoring, or the matching of detected rocks to hand labeled rocks, establishes the number of True Positive (TP), False Positive (FP), and False Negative (FN) detections. Center match scoring is accomplished by computing the centroid of the polygon that marks each rock. The pair-wise distances between the labels returned by the detector and the hand labeled distances are sorted and matches are determined by a greedy algorithm. Pixel overlap matching is similar to center matching except that the number of overlapping pixels between detected and hand labeled rocks is used instead of distance. For all both scoring methods, detections that are matched with hand labeled rocks are considered true positives, detections which do not match are false positives, and hand labeled rocks which do not match detections are false negatives.

The performance metrics of Recall and Precision are calculated for each of the two scoring methods. Recall is the fraction of hand labeled rocks found by the detector, $\text{Recall} = \text{TP} / (\text{TP} + \text{FN})$. Recall captures how well the algorithm does at detecting rocks with no penalty for false detections, that is, what percentage of the actual rocks are identified. Precision is the fraction of total rocks detected that appear in the hand labeled set,

$$\text{Precision} = \text{TP} / (\text{TP} + \text{FP}).$$

Precision captures the degree to which the detector finds only rocks; precision is the percentage of objects identified by the rock finder as rocks that are truly rocks. A high precision directly corresponds to a low false alarm rate. Figure 6 shows the precision results for 110 Mars Yard hazcam images taken by FIDO.

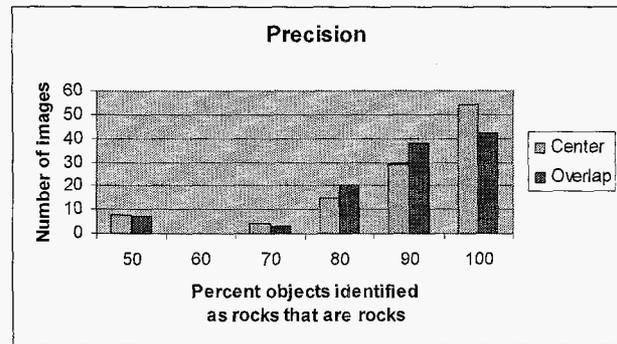


Figure 7 Precision numbers for a total 110 of FIDO hazcam images and 2942 total rocks analyzed.

The two scoring methods show similar results. The average precision per image using the center matching method was 89%, while the average precision using the overlap method was 87%. 83 images had higher than 90% precision. Our recall results will follow in a future publication.

5. CONCLUSIONS AND FUTURE WORK

Future Work

OASIS will continue development work on RockIT, the autonomous rock-finder, and will address occlusion identification and correction. Additional RockIT work includes the preconditioning of illuminated and shadowed regions of the image for the stereo algorithm.

Since gray-scale images were used in the rock finder, when a rock occludes another, it appears that these two rocks are a single rock. The distinction is apparent when range

information is used, as there will be a discontinuity of the depth map. This process requires a stereo-based segmentation of the scene to handle the three-dimensional nature of the occlusion vs. the two-dimensional objects needed to obtain intensity or texture segmentation. Stereo relies on texture matching to perform correspondences between images, which eventually lead to triangulations and range estimation.

The data analysis component will be expanded to perform atmospheric analysis based on image frames. In the images, we will analyze the sky to estimate atmospheric conditions, particle densities and atmospheric rates of change.

In the planning and scheduling component of OASIS, a more difficult opportunistic science scenario will be explored. OASIS can currently re-task the rover to turn in place and take a grayscale image of the new science target generated by the onboard data analysis unit. Future work will have the rover leave its traverse path and approach the new science target to take a science measurement.

For this type of request, the planner will evaluate what new science observations can be performed given current rover resources and state. This process will involve weighing new science observation priorities against priorities of already scheduled operations and interacting with onboard navigation software to evaluate new path options. To this end we will employ optimization techniques to evaluate potential schedule changes and generate a schedule that meets as many high priority science requests as possible.

Conclusions

Autonomous rovers can never replace human geologists on the surface of another planetary body. Humans have remarkable pattern recognition skills and can quickly accumulate large amounts of geological data over a wide area of coverage. Steve Squyres, Athena Science Payload Principal Investigator for the MER mission, has stated that human geologists take 40 seconds to evaluate a geological feature that a rover would take a day to complete⁵.

But human geologists are not yet on Mars and have yet to establish a sustained lunar presence. Rovers, therefore, are useful assistants to geologists, and autonomous capabilities that can maximize the science data gathered and transmitted back to Earth can only enhance that usefulness. The goal of the OASIS project is to provide scientists with these autonomous capabilities. Significant progress has been made in developing autonomous pattern recognition skills for the rover and integrating these skills with other control systems. Validation of these skills is ongoing, and has been greatly assisted by the development of the RockIT GUI.

⁵ AAS conference in Pasadena, California (11/16/2004) during the Carl Sagan Memorial Award Lecture. He based this comment on a field test using JPL's FIDO rover.

As humans finally reach Mars, smart rovers will still be useful assistants, but their role will morph into a more collaborative one – where geologists can interact directly with the rover. Rovers can be especially valuable for exploring dangerous areas like the narrow crevasses of caves or canyon systems, where the risk to human life is great. Autonomous analysis of science data will continue to play an important role in these new exploration environments.

6. ACKNOWLEDGEMENTS

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.

We would like to thank our sponsors, the Intelligent Systems Program of CICT, and the Mars Technology 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), Matt Golombek (JPL), Rob Sullivan (Cornell) and Kathy Weitz (Planetary Science Institute, Tuscon).

REFERENCES

- [1] 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.
- [2] 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.
- [3] I. A. Nenas, A. Wright, M. Bajracharya, R. Simmons, T. Estlin, Won Soo Kim, "CLARAty: An Architecture for Reusable Robotic Software," *SPIE Aerosense Conference*, Orlando, Florida, April 2003.
- [4] C. Urmson, R. Simmons, I. Nenas, "A Generic Framework for Robotic Navigation," *Proceedings of the IEEE Aerospace Conference*, Montana, March 2003.

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 dissertation in the area of computer vision. Dr. Castaño has been advancing the state of the art in onboard science analysis methods for the past five years and has been



lead author on numerous publications in this field. From 1999 - 2001, Dr. 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 Millennium Program's ST6 project. Dr. Castaño is also the Team Lead of the Onboard Autonomous Science Investigation System (OASIS) project and the subtask lead for Data Analysis on the automated Multi-rover Integrated Science Understanding System (MISUS). Her research interests include machine learning, computer vision and pattern recognition.

Michele Judd, P.E., OASIS Research Manager and senior technical staff member of JPL's Earth and Space Sciences 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. She currently manages two additional research tasks in earthquake modeling and simulation, "Numerical Simulations for Active Tectonic Processes: Increasing Interoperability and Performance" and "Complexity Computational Environments: Data Assimilation SERVO Grid." A nationally certified Quality Manager and licensed professional engineer, Ms. Judd has over 15 years of experience in project management, both within JPL and for Mobil Oil.



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 from Tulane 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. 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. 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,

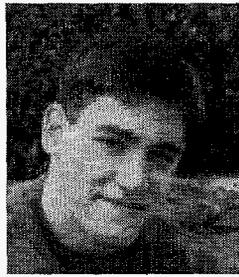


Dr. Gaines was an Assistant Professor in Computer Science at Vanderbilt University. His work at JPL is primarily focuses on planning and execution techniques for planetary exploration rovers.

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.



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.



Tim Stough is a member of the Data Understanding Group and a Research Programmer. His interests include machine learning and computer vision methods. Mr. Stough received both his bachelors and masters of electrical and computer engineering concentrating in artificial intelligence from Purdue University. His work at JPL involves computer vision algorithm development, integration, and testing.



Dr. Kiri Wagstaff is a research scientist with the Machine Learning Systems group at JPL. She received her Ph.D. from Cornell University in 2002 on the subject of how to incorporate domain knowledge as constraints for data clustering algorithms, and she is currently working on an M.S. in Geology at the University of Southern California. She co-chaired the ICML 2003 Workshop on Machine Learning for Autonomous Space Applications. She is the PI for a grant from the NASA Earth Science Technology Office entitled "Interactive Analysis of Heterogeneous Data to Determine the Impact of Weather on Crop Yield," a co-PI for "Knowledge-Enhanced Discovery Systems (KEDS)" (NSF), and a co-investigator for "Classification of Martian Surface Units and Representative Regions of Interest" (NASA). Her research interests include classification, clustering, and exploratory discovery methods, with a focus towards adapting these techniques for onboard computing environments.

