

AEGIS Automated Targeting for the MER Opportunity Rover

TARA A. ESTLIN, BENJAMIN J. BORNSTEIN, DANIEL M. GAINES, ROBERT C. ANDERSON, DAVID R. THOMPSON, MICHAEL BURL, REBECCA CASTAÑO, AND MICHELE JUDD

Jet Propulsion Laboratory, California Institute of Technology

The Autonomous Exploration for Gathering Increased Science (AEGIS) system enables automated data collection by planetary rovers. AEGIS software was uploaded to the Mars Exploration Rover (MER) mission's Opportunity rover in December 2009 and has successfully demonstrated automated onboard targeting based on scientist-specified objectives. Prior to AEGIS, images were transmitted from the rover to the operations team on Earth; scientists manually analyzed the images, selected geological targets for the rover's remote-sensing instruments, and then generated a command sequence to execute the new measurements. AEGIS represents a significant paradigm shift --- by using onboard data analysis techniques, the AEGIS software uses scientist input to select high-quality science targets with no human in the loop. This approach allows the rover to autonomously select and sequence targeted observations in an opportunistic fashion, which is particularly applicable for narrow field-of-view instruments (such as the MER Mini-TES spectrometer, the MER Panoramic camera, and the 2011 Mars Science Laboratory (MSL) ChemCam spectrometer). This paper provides an overview of the AEGIS automated targeting capability and describes how it is currently being used onboard the MER mission Opportunity rover.

Categories and Subject Descriptors: I.2.1 [**Artificial Intelligence**]: Applications and Expert Systems – *Medicine and Science*; I.2.6 [**Artificial Intelligence**]: Learning; I.2.9 [**Artificial Intelligence**]: Robotics – *Autonomous Vehicles*; I.2.10 [**Artificial Intelligence**]: Vision and Science Understanding – *Intensity, color, photometry and thresholding; Shape*; I.4.0 [**Image Processing and Computer Vision**]: General – *Image processing software*; I.4.6 [**Image Processing and Computer Vision**]: Segmentation – *Edge and feature detection, Pixel classification*; I.4.7 [**Image Processing and Computer Vision**]: Feature Measurement – *Size and shape*; I.4.9 [**Image Processing and Computer Vision**]: Applications;

General Terms: Algorithms, Experimentation

Additional Key Words and Phrases: data analysis, spacecraft autonomy, autonomous science

1. INTRODUCTION

The Mars Pathfinder (MPF) and Mars Exploration Rover (MER) missions have demonstrated that mobile rovers are a viable and productive option for exploring the surface of other planets. The MER rovers have traveled over many kilometers of terrain and survived harsh planetary conditions, including Martian winters and major dust

This research was supported by the NASA New Millennium Program, the NASA Mars Technology Program, the JPL Interplanetary Network Development Program, and the NASA Intelligent Systems Program.

Authors' addresses: Tara Estlin, M/S 301-260, Jet Propulsion Laboratory, Pasadena, CA 91109; Benjamin Bornstein, M/S 306-463, Jet Propulsion Laboratory, Pasadena, CA 91109; Daniel M. Gaines, M/S 301-260, Jet Propulsion Laboratory, Pasadena, CA 91109; Robert C. Anderson, M/S 321-220, Jet Propulsion Laboratory, Pasadena, CA 91109; David R. Thompson, M/S 306-463, Jet Propulsion Laboratory, Pasadena, CA 91109; Michael Burl, M/S 306-463, Jet Propulsion Laboratory, Pasadena, CA 91109; Rebecca Castano, M/S 168-527, Jet Propulsion Laboratory, Pasadena, CA 91109; Michele Judd, MC 220-46, California Institute of Technology, Pasadena, CA 91125.

Permission to make digital/hard copy of part of this work for personal or classroom use is granted without fee provided that the copies are not made or distributed for profit or commercial advantage, the copyright notice, the title of the publication, and its date of appear, and notice is given that copying is by permission of the ACM, Inc. To copy otherwise, to republish, to post on servers, or to redistribute to lists, requires prior specific permission and/or a fee.

© 2001 ACM 1073-0516/01/0300-0034 \$5.00

storms, to continue collecting data. The mission's extensive scientific observations have uncovered profound new insights into Mars' current and past environment, the history of its rocks and the various roles and abundances of water (Squyres and Knoll, 2005).

Surface rovers offer scientists the ability to move around a planetary surface and explore different areas of interest. Advances in rover mobility have increased daily traverse range, and with it the opportunity for scientific discovery. While the

Sojourner rover traveled a distance of approximately 100 meters in the entire mission, the two MER rovers (Spirit and Opportunity) have now traveled over 27 kilometers combined. Long traverses have become commonplace on the Opportunity rover. Currently Opportunity is trying to reach a new scientific target – the Endeavour crater (shown in Figure 1). The distance to Endeavour from Opportunity's current position (~11 kilometers) is more than half the 20 kilometer distance Opportunity has traveled since landing in 2003. Many long drives will be used to reach this target over a several year time period.

Unfortunately, communications bandwidth has not grown as fast as rover traverse range. As rover traverse distances continue to increase with each mission, the quantity of data that can be returned to Earth per meter traversed is reduced. Thus, much of the terrain the rover visits on a long traverse may never be examined by scientists. This paper discusses a system developed to autonomously recognize and characterize high value science targets during and after drives without requiring large amounts of data to be transmitted to Earth.

The Autonomous Exploration for Gathering Increased Science (AEGIS) system provides automated targeting for remote sensing instruments on the Mars Exploration Rover (MER) mission. Currently, targets for remote sensing instruments, especially




Fig. 1. MER Opportunity rover location in route to Endeavour crater on Mars. In September of 2008, the Opportunity rover left Victoria Crater to begin a 19 kilometer journey to Endeavour crater. As of September 2010, Opportunity has completed approximately half of its journey. To reach Endeavour, long drives are being emphasized.

narrow field-of-view instruments, must be selected manually based on wide angle imagery already on the ground with the operations team. Examples of narrow field-of-view instruments include the MER Miniature Thermal Emission Spectrometer (Mini-TES), the MER Panoramic Camera (especially when collecting sub-framed images) and the 2011 Mars Science Laboratory (MSL) Laser-Induced Remote Sensing for Chemistry and Micro-Imaging (ChemCam) spectrometer. AEGIS enables the rover flight software to autonomously analyze imagery onboard, select a target based on scientist input, and carry out remote-sensing observations in an opportunistic fashion.

AEGIS operates by analyzing MER navigation camera images to identify terrain features of interest, which are typically rocks with certain characteristics. Scientists on the ground specify these target characteristics in the command sequence sent to the rover. For example, scientists could request measurements of large rocks with high spectral reflectance. Once a target is identified onboard in a navigation camera image, its location is determined and a remote sensing instrument is re-pointed to collect high-resolution follow-up data. Currently, AEGIS is run at the end of traverses when images of the new terrain have not yet been downlinked to Earth. This capability is especially useful for multi-sol (i.e., multi-day) plans where a drive occurs on the first sol and only untargeted remote sensing can be performed on the second and third sols since another communication cycle with Earth has not yet occurred.

AEGIS was uploaded to the MER Opportunity rover in December 2009. The system passed checkout tests onboard Opportunity and is fully operational. On Opportunity, AEGIS is used to collect high-resolution, multi-spectral images of selected targets using the narrow field-of-view Panoramic camera. AEGIS's autonomous targeting capabilities enable the rover to collect a sub-framed image tightly focused on the target of interest and the immediate surrounding area. As a result, AEGIS provides the ability to collect high quality data on a target of interest while conserving downlink volume. Further if no target is found to match the scientist's interest, than no data is collected

In this paper, we provide an overview of the AEGIS automated targeting system. We first describe the general steps used by AEGIS to select new targets and autonomously collect new data of those targets as part of the MER onboard flight software. Next, we discuss each of the components in more detail and describe how AEGIS was integrated and tested with the MER flight software. We then discuss results from AEGIS's use onboard the MER Opportunity rover. Finally we discuss related work and mention some potential future directions for AEGIS expansion.

2. AEGIS SYSTEM

2.1 Overview

The AEGIS system enables autonomous operation of science instruments that target specific terrain features, especially rocks with certain properties. A number of rover remote sensing instruments have a very narrow field-of-view and thus require selection of specific focused targets for sampling. Selecting targets for these instruments has traditionally been a lengthy process. Typically scientists would manually identify the interesting targets in context images that have already downloaded on a previous sol. These context images are collected with wide field-of-view (FOV) cameras such as the MER navigation cameras, which have a 45 degree FOV, or the MER panoramic cameras in a full-frame low-resolution (single filter) mode, using a 16 degree FOV. After reaching the end of a traverse, the rover performs only untargeted data collection until the context images can be downlinked, analyzed, and new measurement commands uplinked. At best this will happen on the next sol and it may require the rover to remain at the same location for several sols. Further it may never happen if it is decided the rover should immediately proceed to a new location due to other goals or engineering constraints.

AEGIS was designed to provide additional image data for the mission scientists. By analyzing image data onboard, within the power and computational constraints, AEGIS can autonomously select targets for onboard instruments and execute a set of measurement activities. The capabilities of AEGIS are presently being demonstrated on the MER Opportunity rover by taking additional measurements with the Panoramic camera in a quarter-frame high-quality (multiple color filter) mode, which uses a 4 degree FOV.

The MER Mini-TES spectrometer is another example of a limited FOV instrument where the AEGIS automated targeting technology would be beneficial. The Mini-TES has a FOV of 8-20mrad. Unfortunately, since the MER rovers have been in operation for over six years, the Mini-TES instrument on the Opportunity rover is not currently functional. For this reason, AEGIS has used MER panoramic cameras exclusively. On the future MSL rover mission, AEGIS could be used to select targets that the ChemCam spectrometer instrument should sample at the end of a long rover drive.

AEGIS is run as part of the MER onboard flight software, which imposes strict computational and resource constraints. All AEGIS components run onboard the MER 20 MHz RAD6000 flight processor, which has an early PowerPC instruction set, with 128 MB of RAM and 256 MB flash memory. Even though it processes full-frame images of over 1 MB each, AEGIS was required to run using less than 4 MB of total RAM to

ensure other onboard processes were not impacted. Time efficiency was another important limitation on the complexity of image analysis algorithms, since operations requiring a fraction of a second on a modern commercial processor could easily take tens of minutes on the Mars surface. Operator-modifiable time limits allow controllers to specify the maximum duration of AEGIS processing; any run exceeding this allowed time will terminate early without affect.

AEGIS was originally developed as part of a large autonomous science framework called OASIS (Onboard Autonomous Science Investigation System) (R. Castano, et al., 2007). OASIS is designed for use onboard a rover to identify and react to serendipitous science opportunities by analyzing the data collected by the rover and then using machine learning techniques, prioritizes the data based on criteria set by the science team. This prioritization can organize the data for transmission back to Earth or search for specific targets specified by the science team. If one of these specified targets appears, the system attempts to act on the new science opportunity by taking new instrument measurements. The AEGIS technology focuses on this second task of using onboard data analysis to acquire new instrument data on science targets, typically rocks, which have been identified in an opportunistic fashion.

AEGIS performs seven major steps to autonomously acquire new targeted data on an interesting science target. These steps are shown in Figure 2 and described below:

- **Acquire an image with the MER navigation camera:** Scientists and other sequence team members select image parameters, such as the pointing direction and resolution, during the AEGIS sequencing process. The navigation camera is typically pointed at a terrain area where potential science targets may be in view.
- **Analyze the navigation camera image for potential terrain targets:** Targets for AEGIS typically correspond to rocks. AEGIS uses an algorithm called ROCKSTER to look for enclosed boundary contours (defined by intensity edges) in grayscale imagery. This algorithm is further detailed in Section 2.2.
- **Extract relevant target features:** AEGIS calculates a set of target features (or properties) for each candidate rock. These properties include measures of size, reflectance, shape, and rock location.
- **Prioritize targets and select top target:** This component uses a prioritization algorithm to analyze rock property data and determine a top candidate. Scientists provide a “target rock signature” in the command sequence. This signature specifies what property values are of interest in the local terrain. Example




Fig. 2. AEGIS Process Steps. When AEGIS is sequenced, the above series of steps is executed onboard the MER Opportunity rover. Parameters can be set during sequencing to specify navigation camera pointing, the “target rock signature” (e.g., rocks of large size and high reflectance/albedo), and settings for the panoramic camera (e.g., what filter set to use).

signatures are “high reflectance”, “round shape”, “large rocks with high eccentricity”, etc.

- **Determine 3D target pointing requirements:** After identifying the best scoring candidate rock, AEGIS selects a center point on the target using an inscribed circle method.
- **Point remote sensing instrument:** AEGIS points the panoramic cameras at the new target using the resulting center point.
- **Acquire new data:** AEGIS acquires additional data with the panoramic cameras. The ground sequencing team can pre-select the exact filters and other imaging parameters to use for each individual run. Typical command sequences take a quarter-framed, multiple filter image with both left and right cameras. The rover downlinks these opportunistic images with other standard data products.

The next few sections provide additional details on these system components.

2.2 Terrain Target Detection

AEGIS uses the ROCKSTER algorithm to identify a set of targets in the initial navigation camera image. ROCKSTER identifies edge segments in grayscale imagery and searches for objects with an enclosed boundary. Such objects typically correspond to rocks when looking at the Mars terrain but could also correspond to small craters or other terrain features. ROCKSTER initially locates partial boundary contours of targets using a procedure similar to the Canny edge detector (Canny, 1986). Specifically, ROCKSTER calculates the intensity gradient 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, including: spurious contours from the sky-ground boundary (horizon line) and texture within individual rocks and the background. ROCKSTER attempts to resolve these problems by splitting the initial contours into low-curvature fragments.

A gap-filling mechanism joins nearby contour fragments whose endpoints lie within a predefined radius. The final step is to regroup the edge fragments into coherent contours, which is accomplished through background flooding. Figure 3 shows a high-level view of the process that ROCKSTER uses to detect and generate usable target contours.

ROCKSTER distinguishes a variety of Martian geologic features including outcrop, cobbles, boulders, and sediment. While detecting many of these features is a natural fit to ROCKSTER image segmentation approach, the detectability of outcrop may be less obvious. Outcrop is often distinct in intensity from surrounding regolith, making for well-defined edges. However, since portions of each outcrop can be obscured by small amounts of soil debris, detected edges have a tendency to fragment into smaller edge segments. The ROCKSTER gap-filling procedure rejoins many of these disconnected segments. In contrast, stereo imaging and range-based techniques pose a number of detection challenges, as outcrop is often flush with the surrounding terrain and therefore does not present significant depth discontinuities.

For automated targeting of limited FOV instruments, false detections are costly and high precision is important. Thus for this application, ROCKSTER is typically run in a mode that reduces false positives; however, this also has the effect that fewer overall targets are found. This behavior is a trade-off that can be adjusted depending on the application, i.e., for some applications it may be more important to find a larger percentage of the true targets despite the higher risk of returning some false positives




Fig. 3. AEGIS Target Detection using ROCKSTER. This picture shows the major steps performed when AEGIS analyzes an image for specific targets. The goal of target detection is to search for objects in the image with an enclosed boundary. These objects typically correspond to rocks. Steps consist of locating partial edge contours, applying a set of morphology operations to clean up and smooth out detected contours, flood filling the image to separate out rock targets, and then identifying final target contours. Since the MER rovers have limited processing capacity and available memory, ROCKSTER relies on techniques that can run quickly and robustly in such an environment.

(i.e., non-interesting targets). The sequencing team can also choose to limit target detection to specific rectangular subregions of the image. This option is useful for excluding image regions that contain the deployed rover arm, rover tracks, or other features that could generate spurious detections.

Due to the limited processing capacity and memory available onboard the MER rovers, ROCKSTER relies on techniques that can perform quickly and robustly in such an environment. Image preprocessing, in particular, smoothing reduces the total number of edge elements detected. Considering fewer edges saves considerable computational effort in downstream gap-filling and contour following. ROCKSTER also tracks a number of internal space and time complexity measures related to the overall segmentation computation. As a measure of last resort, if these limits are exceeded, ROCKSTER and AEGIS terminate gracefully, but prematurely. While not ideal for meeting science objectives, monitoring and bounding the usage of precious onboard computational resources that correlate directly to mission timelines and available power, is essential.



Fig. 4. Examples of AEGIS Target Features. These images show examples of different rock features that are extracted by AEGIS. On the left are rocks of different reflectance/albedo. On the right are rocks of various shapes. A set of numerical attributes is collected for reflectance and shape (as described in the text). Other feature categories include size and target location.

Although AEGIS is using terrain targets identified from monocular grayscale imagery, the overall approach is not tied to any particular type of target, data source, or instrument type. For example, the general OASIS system has been applied to analyze spectrometer data as well as identifying atmospheric targets, such as clouds and dust-devils, in MER imagery (Castano, et al., 2006; A. Castano et al., 2007).

2.3 Target Feature Extraction

Once candidate targets are identified, the AEGIS system computes numerical attributes corresponding to properties of each target image region. Some examples of these properties are shown in Figure 4.

Surface reflectance: The surface reflectance (or albedo) of a target is an indicator of the integrated reflectance properties of a target's surface. The reflectance properties of a rock can provide important information about its mineralogical composition. AEGIS measures surface reflectance by computing the mean gray-scale value of the pixels within the target. Note that this value can be affected by shadowing so the calculation does not provide a perfect measure of physical surface albedo. However, it provides some useful information about surface properties. It has proved useful for discriminating between shaded rocks protruding above the sediment, and flat rock outcrop that generally appears brighter to the sensor. AEGIS calculates additional moments of the pixel intensity distribution including variance, skew and kurtosis; these higher moments serve as a rough proxy for texture. Examples of rocks with varying reflectance are shown in Figure 4.

Size: One of the most important properties of rocks on the surface is their size, which can be used to identify sorting and geologic contacts. Several features describe the target size. The pixel area of the rock is one simple measure. AEGIS also calculates the radius of the largest inscribed circle that fits within the contour. It computes this latter measure efficiently using an image distance transform. A third measure of size is the length of the semimajor and semiminor axes of the best-fitting ellipse. AEGIS fits an ellipse to the rock's outline using a least-squares criterion (Fitzgibbon, et al., 1999; Halir and Flusser, 1998). Unfortunately, AEGIS does not have access to range data on the acquired Navigation camera images. Thus stereo information could not be used to determine true target size, though it has been used in other versions of the system.

Shape: Although the shape of a rock is complex and often difficult to describe, significant geologic information can be extracted from this property to better understand provenance (source of material) and environmental conditions. Various shape parameters are used to classify rocks in terrestrial studies, including elongation (or aspect ratio), ruggedness (or angularity), and surface area. AEGIS uses the eccentricity of the fit ellipse as well as a ruggedness score based on the square of the perimeter divided by the contour's pixel area (Hentschel and Page, 2002).

Pixel Location: The x and y coordinates of the ellipse centroid and inscribed circle are also treated formally as features. Incorporating these values in the feature vector lets operators favor or exclude candidates based on their position in the image. (Note that masking out the rover deck when it appears in images is performed as a separate process.)

2.4 Target Prioritization

Once features are computed for identified targets, AEGIS next prioritizes candidate targets and selects a top target. To guide the prioritization process, AEGIS uses a pre-specified target signature, corresponding to particular feature attributes (e.g. prefer rocks that are large in size and have high reflectance), provided by the MER science team during ground sequencing. This algorithm enables scientists to efficiently and easily stipulate the importance of each particular feature. For each run of AEGIS, scientists can specify one target signature but can change this signature each time the system is run. AEGIS gives each candidate target a score f corresponding to a weighted sum of up to two feature values x_1 and x_2 . Two coefficients α_1 and α_2 control whether the algorithm prefers high or low feature values, while a weighting coefficient β describes the comparative importance of the second feature. Note that one need not specify a second feature at all, in which case $\beta=0$. We have:

$$f = \alpha_1 x_1 + \beta \alpha_2 x_2, \alpha_i \in \{-1, 1\}, \beta \in [0, 1]$$

Earlier versions of AEGIS allowed an arbitrary number of features to be used in the above function and multiple t , however due to limitations on the number of parameters available for the AEGIS commands on MER, the number of features for each run was limited to two for the MER deployment of AEGIS.

Scientists can also specify a filter that removes from consideration any target where the value of a specified feature falls outside a threshold (either above or below). This is useful for excluding targets that are likely spurious detections, such as very small rocks or objects in the distant background. The sequencing team sets all feature selection and filter parameters manually and can change them each time AEGIS is used. In addition to feature-based filtering, AEGIS can optionally remove all contours that intersect the known locations of the rover deck. Deck masking projects a polygonal model of the rover solar panels and High Gain Antenna into the image. Operators can instruct AEGIS to ignore any contour that intersects this image area. This step is important because the solar panels and antenna contain many closed features and sharp edges that may be detected and labeled as targets.

AEGIS selects the top target from the filtered candidates and uses it to acquire new targeted data. Figure 5 shows two potential targets selected in MER navigation camera images with selection criteria favoring large size (on the left) and round shape (on the right). In both cases the AEGIS target finder is run in a mode that decreases the number of false positives (e.g., a shadowed area of sand) but finds fewer of the overall rocks in the image. Several software tools and documentation manuals are available to help the sequencing team select appropriate target signature and filter parameters for different AEGIS runs. Further several standard profiles have been developed that can be used to find common occurrences (such as outcrop or loose surface rock).

2.5 New Data Acquisition

Once AEGIS selects a top target, it commands additional science measurements. First, it selects a center point on the target using the center point of the largest inscribed circle that fits within the target contours. The software then points the MER panoramic left and right cameras at the target.

The type of follow-up observations of the top target are defined in advance by the ground team, but can be set each time the AEGIS capability is used. Typically AEGIS is




Fig. 5. AEGIS target selection in MER images. On the left, AEGIS was directed to select a large rock as a potential target for a limited FOV instrument, such as the 2011 MSL ChemCam Spectrometer, to sample. On the right, AEGIS was directed to select a round rock as a potential target. These images were taken on the MER mission by the MER Opportunity and Spirit rovers and were used during AEGIS development to test the AEGIS software.

used to collect multiple filter, subframed images with both left and right panoramic cameras. This approach provides high quality images that are smaller in size due to subframing and thus less expensive to downlink. The resulting panoramic camera images are downlinked with other standard MER data products for that sol.

The ground team pre-allocates resources such as power, time, and onboard memory whenever AEGIS is scheduled to ensure that sufficient resources exist onboard to collect the new data. This approach ensures that the correct resources and required rover states are fully verified before any extra activities are commanded. However, since it is possible that AEGIS might not find any unfiltered, unmasked targets (e.g., if an image contains only dunes), a more optimal approach would be to only allocate resources only if a good target was actually found. The original AEGIS system contained an automated planning and scheduling system that could perform this allocation dynamically, only if and when a new science target is found (Estlin, et al., 2008). Due to computing requirements, the planning component was not used in the final version of AEGIS for the MER mission. However, we hope to use it in future versions of AEGIS.

3. INTEGRATION WITH MER MISSION FLIGHT SOFTWARE

We faced a number of constraints and challenges when integrating AEGIS with MER mission flight software. Due to funding and time limitations, we did not have the option of creating a new version of MER flight software that included AEGIS, which is how

past new software technology updates to the mission had been performed. Instead, we developed AEGIS as a standalone software module that co-opts an existing onboard autonomy technology to enable AEGIS to be commanded and executed. We also delegated part of the AEGIS response to a third onboard autonomy technology, Visual Target Tracking (VTT), for mast motion and pointing. The standalone nature of AEGIS and its use of and delegation to other established MER autonomy technologies, afforded a more manageable integration and test effort.

AEGIS is designed as a standalone software object module that is loaded into rover memory prior to each use and unloaded from memory after each use. This loading process takes less than 30 seconds to complete and the unloading process is performed automatically through normal operations, since the MER Opportunity rover computers are shutdown and restarted several times each sol. Once the AEGIS software is loaded, it leverages previous MER autonomy technology developments (S. Hayati, et al., 2007). In particular, it uses the same interface as for the MER dust-devil detection software (A. Castano et al., 2008), which shares several similarities with AEGIS including providing a rover autonomy capability and analyzing Navigation camera imagery for targets of interest. The particular sequencing command used to execute AEGIS has 33 algorithm parameter command arguments that we are able to map to the specific needs of the AEGIS algorithm, including specifying science criteria for prioritizing targets. We were also able to use previously designed telemetry data, to report AEGIS run statistics and up to 10 image locations that bound detected targets.

Figure 6 highlights AEGIS' interaction with MER flight software components. Although AEGIS does not actually drive the rover, the Mobility module is well suited for running autonomy algorithms such as AEGIS. The design of MER flight software, including the roles of flight software tasks and task priority levels, allows the Mobility task to run computationally intensive algorithms, such as hazard detection and visual odometry, without compromising the ability of other flight software tasks from meeting time-critical deadlines. When AEGIS is used, the operations team prepares a sequence specifying where to point the navigation camera in order to collect the image that AEGIS will process along with a specification of the type of targets the science team is interested in for this run. When the AEGIS command is executed by the rover, the MER Mobility module coordinates with other flight software components to acquire the appropriate navigation camera image and then runs the AEGIS target detection algorithm passing in a pointer to the acquired image. In addition, the Mobility module provides convenient access to memory-management capabilities used by AEGIS.




Fig. 6. AEGIS Integration with MER Flight Software. AEGIS is run as part of the mobility module and has to interact with several other flight software components when the algorithm is executed. For example, AEGIS interacts with the Visual Target Tracking module, to point the Panoramic Camera at the top target.

In implementing AEGIS for MER, we were able to leverage previous MER autonomy technology developments (S. Hayati et al., 2007). In particular, it uses the same interface as for the MER dust-devil detection software (A. Castano et al., 2008), which shares several similarities with AEGIS. They both process navigation camera images and they both produce data products to identify objects of interest that were found. As such, we were able to make use of the existing command and data handling facilities that were originally developed for the dust-devil detection software. This helped reduce the effort in deploying AEGIS on MER. The particular sequencing command used to execute AEGIS has 33 algorithm parameter command arguments that we are able to map to the specific needs of the AEGIS algorithm, including specifying science criteria for prioritizing targets. We were also able to use previously designed telemetry data, to report AEGIS run statistics and up to 10 image locations that bound detected targets.

While running, AEGIS interacts with a few other flight software components. AEGIS uses the Timing service to periodically check the current spacecraft time. When the operations team sequences AEGIS, they can specify how long AEGIS is allowed to run or a specific time by which AEGIS must complete. This is an important feature for allowing an autonomous algorithm to be used operationally as it gives the operations team the ability to set hard constraints on the timing of AEGIS which allows them to schedule AEGIS around other activities including other science observations, communication passes or planned rover shutdowns as well as to ensure the algorithm

complies with resource allocations made on the ground. If AEGIS detects that it has approached its deadline, it will give up its search for targets and exit gracefully. AEGIS queries the Imaging flight software component to determine the current viewing angle of the navigation camera. Using this information AEGIS is able to mask out regions of the image, as explained in Section 2.4.

One of the last steps of the AEGIS process is to acquire a follow-up measurement of the top target with a high-resolution, narrow field-of-view, science instrument. Unless the top AEGIS target happens to be centered in the science instrument's field of view, this follow-up measurement requires some motion on the part of the rover mast. To acquire the follow-up measurement, AEGIS leverages another piece of onboard autonomy technology, a Visual Target Tracking (VTT) system (Kim, et al., 2009).

In order to center the MER Panoramic camera on a select target, AEGIS uses the following process. First, AEGIS determines a two-dimensional image (pixel) location for its top target. This location is communicated to the onboard VTT subsystem as a seed (initial) point. After AEGIS seeds VTT, the AEGIS target detection algorithm completes. The onboard sequence containing the call to the AEGIS command then calls VTT to acquire a stereo image pair and compute a three-dimensional location and distance-to-target, from the two-dimensional seed point. If VTT succeeds in determining the target in space, a pointing sequence is called which uses the distance-to-target and known rover mast geometries to center the target in the panoramic camera field-of-view. With the panoramic camera centered on the top AEGIS target, additional measurements, dictated by a sequence of commands, can be acquired. By placing the follow-up imaging in a sequence, the operations team can easily change the type of follow-up imaging acquired for the selected target on each run. For example, if, on a given day, data is extremely limited, the operations team can prepare a follow-up sequence that uses only a small number of panoramic camera filters and employ higher levels of compression. If data is more plentiful, the follow-up sequencing could include many filters and make use of loss-less compression of the images.

In addition to integrating the AEGIS with flight software components, the AEGIS detection algorithm itself was optimized in order to be effectively deployed on Opportunity. In particular, the algorithm was modified to run more efficiently and to require less memory. The original version of the target detection algorithm consumed 64 MB of memory, but was ultimately reduced to 4 MB to fit within our allotted onboard memory budget.

4. SYSTEM TESTING ON EARTH

AEGIS was extensively tested before being approved for upload to the MER mission Opportunity rover. The AEGIS test plan included several levels of testing at increasing levels of system integration. At the lowest level were unit tests, used to verify the behavior of individual functions, followed by regression tests, and rover system-level tests with both avionics simulators and eventually full rover hardware. The AEGIS code was also regularly scanned for errors with the Coverity PreventTM (Bessey et al., 2010) static source code analysis tool. The entire source code underwent three separate peer reviews by AEGIS developers, MER flight software developers, and JPL flight software and autonomy technology developers who were not familiar with the software intricacies of either AEGIS or MER.

Unit tests were automated and run regularly as part of the normal development cycle and nightly build process. Whenever functions were updated, all unit tests were run to identify any negative ripple effects that could have been caused by the code change. As each bug was discovered, a unit test was written that fails in order to prove the existence of the bug. The passage of this unit test then proved the bug has been fixed and also guarded against the inadvertent reintroduction of the same bug at a later time. When AEGIS software development was finished, low-level functionality was verified by 348 unit tests.

Regression tests, also automated and run as part of the nightly build process, were used to verify correct end-to-end algorithm behavior. To assemble the regression test set, we collected 246 MER Navigation camera images. The images represent a broad sampling of the various geologic terrain types both MER rovers have encountered throughout their long missions (e.g. cobbles, dunes, craters small and large, outcrop and regolith, etc.). Before regression testing began, we ran a series of performance tests and parameter tunings on this image suite. Since some system parameters could not be changed once the software was uploaded, we used this image set to tune certain parameters so the system could operate robustly on future images. By including images from a broad range of terrain types, we attempted to safeguard the system from being overly tuned for any one specific terrain category. Once we were satisfied with AEGIS' performance on each image, we fixed the algorithm parameter settings and recorded the expected results. In addition to running regression tests on our build server and development workstations, we adapted our software test harness to work with the MER avionics simulators. This proved invaluable in quickly identifying subtle issues caused by compiler and hardware platform differences.



Fig. 7. MER full hardware testbed which is located at the NASA Jet Propulsion Laboratory. A number of high fidelity tests of the AEGIS software were run on this testbed.

When algorithm development was complete and the AEGIS had been well tested in isolation, we “froze” (did not modify) the AEGIS code in order to proceed with system-level tests. The purpose of these tests was to verify the AEGIS algorithm in the context of MER mission flight software and the entire rover hardware system. The first set of tests was in the MER avionics simulator environment and second with full rover hardware situated in a building-sized Martian-like terrain sandbox. The full hardware testbed is shown in Figure 7. Tests conducted at this level were of the highest possible fidelity. Each test was driven by onboard sequences similar to those used daily by the MER tactical operations teams. This is also when we developed our initial Martian surface checkout sequences (see Section 5). Test input was obtained from both onboard sequences and Navigation camera hardware. Test output was encoded as standard MER data product telemetry and flowed through a UHF downlink from the engineering rover hardware to our test workstation. In this environment, we verified AEGIS either functioned correctly or degraded gracefully and in either case, without harming the rover system. We tested a variety of situations including nominal operations, off-nominal parameters, reduced memory availability, stress tests, and multiple types of preempted and interrupted operations. After a MER project software review board examined our test

plans, procedures, and results, we were given final approval to begin planning and coordinating the uplink of AEGIS software to the MER mission Opportunity rover.

5. RESULTS ON MER OPPORTUNITY ROVER

The AEGIS software was uploaded to the MER mission Opportunity rover in December of 2009. Over the next few months a series of checkout steps was performed to exercise different AEGIS components. All of these checkout steps executed successfully. Figures 8 and 9 show the result of the final checkout, which was run in March 2010 and was the first time all AEGIS components were exercised together onboard the Opportunity rover. During this run, Opportunity was located near the Mars Concepción Crater. The analyzed navigation camera image contained a scattering of loose crater ejecta, providing a number of potential targets. For this run, the sequencing team parameterized AEGIS to look for targets of large size and low reflectance. Figure 8 shows the top target selected (indicated by the yellow marker) and other targets that comprise the top ten ranked targets (indicated by the blue markers). Note that AEGIS does not downlink the detected target contours due to telemetry constraints.

Figure 9 shows the resulting panoramic camera image taken of the top target. Overall these results were excellent. All system components were run successfully with no errors. All top ten targets were good selections that matched the specified target signature. Further, the resulting panoramic images nicely captured the top target, showing that the automated pointing and data acquisition components of the systems were operating correctly.

As of March 2010, AEGIS has been run sixteen times on Mars (Table 1) and has regularly chosen appropriate targets based on the specified target signature profile. On most runs AEGIS has successfully detected either rock outcrop or loose rocks that were likely crater ejecta. For example, Figure 10 shows the target selection results of sol 2221 where the system identified a number of rock outcrop targets in an area primarily dominated by sand dunes. In one run (sol 2204) there were no rocks in the navigation camera image but the system did detect targets consisting of compressed sediment features in rover tracks. Detecting tracks as potential targets is a known behavior and in some cases can be arguably appropriate since interesting MER science discoveries have been made on disturbed material in rover tracks. In this particular case, tracks were not expected in the image (the drive faulted out early during a turn maneuver and in an unexpected location) thus the track detections could be considered false positives.




Fig. 8. Results of image analysis from the first full AEGIS run onboard the MER Opportunity rover. Shown is the Navigation camera image taken for this run. At this time, Opportunity was located near the Mars Concepción crater and the area imaged contains a scattering of crater ejecta. AEGIS was told to look for large, dark rocks. The yellow marker shows the top target and the blue markers show the other targets that were ranked in the top ten.



Fig. 9. Panoramic false-color image of top AEGIS target. This figure shows the resulting composite panoramic camera image, taken during an AEGIS run. The component images are one-quarter subframe field of view, taken with the left camera of the stereo panoramic camera through filters admitting wavelengths of 750 nanometers, 530 nanometers and 430 nanometers. The false color used in this composite makes differences between rock materials easy to see.

Sol	Selection Profile	Rock Detections	Non-rock Detections	Notes
2138	Large size	10	0	Checkout test. Outcrop detected.
2172	Large size + dark reflectance	10	0	Checkout test. Loose rock detected.
2204	Large size	0	10	Tracks detected. Image contained mainly soil.
2221	Large size	10	0	Outcrop detected.
2247	Large size	10	0	Outcrop and loose rock detected.*
2278	Large size + dark reflectance	7	0	Outcrop detected.*
2290	Large size + light reflectance	10	0	Outcrop detected.
2292	Large size + dark reflectance	0	0	Duration limit reached.
2304	Dark reflectance	9	1	Outcrop and dark patch of soil detected.
2312	Light reflectance	10	0	Outcrop detected.
2313	Light reflectance	10	0	Outcrop detected.
2325	Dark reflectance	0	2	Image contained mainly soil. Top target was false positive of rover cable pulls.
2332	Dark reflectance	0	0	Image contained mainly soil. No targets found.
2407	Large size + dark reflectance	0	0	Image contained mainly soil. No targets found.
2428	Dark reflectance	1	7	Image contained mainly soil. Top target was false positive near the horizon.
2550	Large size + dark reflectance	10	0	Targets detected include several large boulders and outcrop.

Table 1. AEGIS runs on the Opportunity rover as of the time of publication. Each run acquires a sub-framed, multiple filter Panoramic camera image of the top target. Returned telemetry includes information on the top ten targets found in each image. Most runs have successfully detected pieces of outcrop and loose rocks (e.g., crater ejecta). On one run (sol 2204) no rocks were present in the scene, but soil disruptions due to rover tracks were visible and detected as potential targets. On several runs (sols 2332 and 2407), only soil was visible in the image and AEGIS correctly returned a result of no targets found. False positives have also occurred (sols 2325 and 2428), especially when AEGIS is allowed to consider very small targets (e.g., less than 20 pixels in size). Runs with an * indicated that no panoramic images were acquired due to a preceding drive error, which prevents AEGIS from articulating the mast.

On several runs (sols 2332 and 2407), only soil was visible in the image and AEGIS correctly returned a result of no targets found. False positives have also occurred. On sol 2325, cable pulls on the rover deck were detected when they slightly extended outside of the deck masking polygon in the image field-of-view. On sol 2428, AEGIS was allowed to consider very small targets (less than 25 pixels), and found a set of very small false positive targets. In both these runs, no significant rocks were in the scene and AEGIS parameters have been modified since to prevent similar occurrences.




Fig. 10. Top targets selected in MER navigation camera image. This figure shows the results of image analysis from an AEGIS run performed in April 2010 on the MER Opportunity rover. In this run, AEGIS was told to look for large rocks. The yellow marker shows the top target and the blue markers show the top ten targets.

During these runs, AEGIS has also shown its ability to not only collect interesting science data, but to also save valuable operations time. On sol 2550, AEGIS detected several large boulders, which were of high interest to the science team. Because the system had already collected a multiple color filter Panoramic image of these targets, the operations team was able to devote more time in the next sol's plan to driving and other science activities that would have been shortened if the AEGIS data has not been collected. Future runs should provide additional data with which to further evaluate system performance. The MER mission plans to continue the use of AEGIS to enable automated targeting of the MER panoramic camera.

This paper is intended to report on our experience and results gathered during actual flight usage. Related publications discuss comparisons of ROCKSTER to other approaches for image rock detection (Thompson and Castano, 2007) and a study that compared a separate rock finding approach (explored as part of previous work) to a strawman approach for using blind sampling to select targets (Castano, et al., 2006; Castano, et al., 2007).

6. RELATED WORK

AEGIS builds on a foundation of related work in autonomous rover science systems. Terrestrial platforms have demonstrated classification of terrain types or features in analog environments, as well as automatic followup utilizing cameras and spectrometers. One early system autonomously identified meteorites in Antarctica (Pedersen, 2001; Wagner, et al., 2001). Here, the Nomad rover segmented dark rocks against an ice background and guided the robot to perform followup measurements with an arm-mounted spectrometer. Another early system provided techniques for analyzing field test data by the Marsokhod rover (Gulick, et al., 2001). These experiments demonstrated rock detection on a soil background using shadows as a cue to infer three-dimensional shapes. Other research focusing on the problems of feature extraction and prioritization includes work by Roush (2004) and Dunlop (2006).

Other experiments have focused on autonomous science during longer traverses. A field campaign demonstrated the utility of rover science autonomy during long over-the-horizon drives, as part of a broader survey to characterize the distribution of life in the Atacama Desert (Smith et al., 2007). Later work at Amboy Crater demonstrated rock detection and spectrometer followup as well as adaptive survey techniques, (Calderón et al., 2008, Thompson, et al., 2008). More recently researchers have investigated automated target selection for the upcoming 2018 ESA ExoMars rover mission (Woods, et al., 2009; Pugh et al., 2010). Woods, et al., (2009) use a graph-based growing algorithm to separate rocks from terrain and then edge detection techniques to identify when layering or bedrock is present.

In this context, AEGIS now contributes the first deployment of autonomous rover geology to a planetary rover mission. The integration with flight software and deployment to a space mission is a key advance since it imposes challenging constraints on processing power and memory. Finally, AEGIS provides an important example of how autonomous target prioritization has been used in practice by scientists, and how the processing and followup has been integrated into standard mission operations.

A separate image processing approach has been using on the MER rovers to detect dynamic atmospheric phenomena, such as dust-devils, in rover images (Castano, et al., 2008). However, this approach to event detection is quite different from analyzing images for terrain features. The dust-devil detector looks for differences between a series of images to detect areas where motion has occurred between images. A large challenge of this process is to detect what often are subtle feature changes in the presence of significant image noise. Further, the dust-devil detector only performs image

prioritization by deciding what images should and should not be downlinked. In comparison, AEGIS detects a large range of static terrain targets in single images, prioritizes these targets based on geological features, and enables the acquisition of new data on targets of high priority.

7. FUTURE WORK

In future work, we plan to expand AEGIS on several fronts. As already mentioned, AEGIS could be applied for automated targeting of the MSL ChemCam spectrometer instrument. We are currently working with the MSL ChemCam team to determine how AEGIS could be best applied and/or expanded to select high quality targets for this instrument. For instance, one extension could be to incorporate the evaluation of new target features, such as texture, layering, or color. Use of these features has been explored in past work on the full OASIS system (Castano, 2007), but has not yet been flown.

We would also like to expand AEGIS to evaluate onboard when resources are available to acquire additional opportunistic measurements. For work with the JPL research rovers, the CASPER automated planning system has been applied to perform online sequence modification in support of opportunistic science (Estlin, et al., 2008). Currently on AEGIS, resources must be pre-allocated by the ground, before the capability is executed, regardless of whether an interesting target is found. A more efficient approach would be to allocate resources onboard only if a top target is determined. This extension would also enable the capability to be run more frequently since resource levels could be evaluated at run-time.

Another new area of work would use AEGIS to select targets for close-contact instruments (such as the MER Microscopic Imager). These instruments are typically located on a rover arm and require close proximity to the target of interest. AEGIS could select the initial target and use existing rover technology to autonomously drive and place a rover instrument (Bajracharya, et al., 2005) to collect the final measurement. A further area of inquiry would use data collected from other types of instruments (such as spectrometers or ground-penetrating radar) to enable AEGIS to select interesting science targets.

CONCLUSIONS

In summary, AEGIS enables autonomous recognition of scientifically interesting targets in MER rover navigation camera imagery. These targets can then be successfully characterized without requiring a communication cycle with ground. New measurements

with the MER Panoramic cameras can be acquired during or immediately after a long drive, and before images of the rover's current location have been acquired and analyzed by ground. AEGIS was uploaded to the MER Opportunity rover in December 2009. The system has been successfully checked out onboard Opportunity and is fully operational. AEGIS has been run a number of times on the surface of Mars and has consistently picked out appropriate targets. The MER mission plans to continue the use of AEGIS to enable automated targeting of the MER panoramic camera. Future runs should provide additional data that can be used to further evaluate AEGIS' performance.

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 New Millennium Program, the Mars Technology Program, the JPL Interplanetary Network Development Program, and the Intelligent Systems Program. We would also like to thank the Mars Exploration Mission for their support of this work.

REFERENCES

- MAREK, W., AND TRUSZCZYNSKI, M. 1989. Relating autoepistemic and default logics. In *Proceedings of the 1st International Conference on Principles of Knowledge Representation and Reasoning*, Toronto, Canada, May 1989, H. BRACHMAN AND R. REITER, Eds. Morgan Kaufmann, San Mateo, CA, 276-288. Received August 2000; revised March 2001; accepted May 2001
- BAJRACHARYA, M., DIAZ-CALDERON, A., ROBINSON, M., AND POWELL, M., "Target Tracking, Approach, and Camera Handoff for Automated Instrument Placement," *Proceedings of the 2005 IEEE Aerospace Conference*, Big Sky, Montana, March 2005.
- BESSEY, A., BLOCK, K., CHELF, B., CHOU, A., FULTON, B., HALLEM, S., HENRI-GROS, C., KAMSKY, A., MCPEAK, S., ENGLER, D. 2010. A few billion lines of code later: using static analysis to find bugs in the real world, In *Communications of the ACM* 53:2, 66–75.
- CALDERÓN P., F., D. R. THOMPSON, D. WETTERGREEN, "Autonomous Rover Reflectance Spectroscopy with Dozens of Targets," *Proceedings of the International Symposium on Artificial Intelligence, Robotics, and Automation in Space (i-SAIRAS)*. Pasadena, California, 2008.
- CANNY, J. , "A Computational Approach to Edge Detection," *IEEE T-PAMI*, vol. 8, 679-698, 1986.
- CASTAÑO, A., FUKUNAGA, A., BIESIADECKI, J., NEAKRASE, L., WHELLY, P., GREELEY, R., LEMMON, M., CASTAÑO, R., AND CHIEN, S. 2008. Automatic detection of dust devils and clouds on Mars. In *Machine Vision and Applications* 19:467–482.
- CASTANO, R., ESTLIN, T., ANDERSON, R.C., GAINES, D., CASTANO, A., BORNSTEIN, B., CHOUINARD, C., AND JUDD, M., "OASIS: Onboard Autonomous Science Investigation System for Opportunistic Rover Science," *Journal of Field Robotics*, Vol 24, No. 5, May 2007.
- CASTANO, R., ESTLIN, T., GAINES, D., CASTANO, A., CHOUINARD, C., BORNSTEIN, B., ANDERSON, R.C., CHIEN, S., FUKUNAGA, A., AND JUDD, M., "Opportunistic Rover Science: Finding and Reacting to Rocks, Clouds and Dust Devils," *Proceedings of the 2006 IEEE Aerospace Conference*, Big Sky, Montana, March 2006.

- DUNLOP, H. "Automatic Rock Detection and Classification in Natural Scenes," Carnegie Mellon Masters Thesis, CMU-RI-TR-06-40, 2006.
- ESTLIN, T., GAINES, D., BORNSTEIN, B., CASTANO, R., RABIDEAU, G., JOHNSTON, M., CHOUINARD, C., JUDD, M., ANDERSON, R.C., AND NESNAS, I. "Supporting Increased Autonomy for a Mars Rover," *Proceedings of the International Symposium on Artificial Intelligence, Robotics and Automation in Space (iSAIRAS 2008)*, Los Angeles, CA, Feb 2008.
- FITZGIBBON, A.W., PILU, A. M., AND FISHER, R. B. "Direct least-squares fitting of ellipses". *Pattern Analysis and Machine Intelligence*, 21(5): 476--480, May 1999.
- GULICK, V., MORRIS, R., RUZON, M., AND ROUSH, T., "Autonomous image analysis during the 1999 Marsokhod rover field test," *Journal of Geophysical Research*, 106(E4), 2001.
- HALIR, R. AND FLUSSER, J., "Numerically stable direct least squares fitting of ellipses," *Proceedings of the Sixth International Conference in Central Europe on Computer Graphics, Visualization and Interactive Digital Media*, 1: 125--132, February 1998.
- HAYATI, S., RANKIN, A., KIM, W., LEGER, P., CHIEN, S. AND ALI, K. 2007. Advanced Robotics Technology Infusion to the NASA Mars Exploration Rover (MER) Project, In *IFAC Symposium on Intelligent Autonomous Vehicles*, September 2007.
- HENTSCHEL, L. M. AND PAGE, N.W., "Selection of Descriptors for Particle Shape Characterization," *Particle & Particle Systems Characterization* 20 (1), 25-38. 2002.
- KIM, W.S., NESNAS, I. A., BAJRACHARYA, M., MADISON, R., ANSAR, A.I., STEELE, R.D., BIESIADECKI, J.J., AND ALI, K.S., "Targeted Driving Using Visual Tracking on Mars: From Research to Flight," *Journal of Field Robotics*, 26, 3, March 2009, 243-263.
- PEDERSEN, L., "Robotic Rock Classification and Autonomous Exploration," PhD thesis, Robotics Institute, Carnegie Mellon University, CMU-RI-TR-01-14, 2001.
- PUGH, S., BARNES, D., PULLAN, D., AND TYLER, L.. "Knowledge based science target identification system (KSTIS)." *Proceedings of the International Symposium on Artificial Intelligence, Robotics and Automation in Space*, 2010.
- ROUSH, T. "Essential autonomous science inference on rovers (EASIR)," *Proceedings of the IEEE Aerospace Conference*, Big Sky, Montana, 2004.
- SMITH, T., THOMPSON, D. R., WETTERGREEN, D. CABROL, N. A., WARREN-RHODES, K. AND WEINSTEIN, S. "Life in the Atacama: Science Autonomy for Improved Data Quality." *Journal of Geophysical Research* 112, Dec 2007.
- SQUYRES, S. W., AND KNOLL, A.H., Sedimentary rocks at Meridiani Planum: Origin, diagenesis, and implications for life on Mars, *Earth and Planetary Science Letters*, 240, p. 1-10, 2005.
- THOMPSON, D.R. AND CASTANO, R. "A Performance Comparison of Rock Detection Algorithms," *Proceedings of the 2007 IEEE Aerospace Conference*, Big Sky Montana, 2007.
- THOMPSON, D. R., SMITH, T., WETTERGREEN, D. "Information-optimal selective data return for autonomous rover traverse science and survey," *Proceedings of the IEEE International Conference on Robotics and Automation*, 2008.
- WAGNER, M., APOSTOLOPOULOS, M. D., SHILLCUTT, K., SHAMAH, B., SIMMONS, R., AND WHITTAKER, W., "The science autonomy system of the Nomad Robot," *Proceedings of the International Conference on Robotics and Automation*, Seoul, Korea May 2001.
- WOODS, M., SHAW, A., BARNES, D., PRICE, D., LONG, D., AND PULLAN, D., "Autonomous science for an ExoMars Rover-like mission," *Journal of Field Robotics*, 26:4, April 2009.