AEGIS Automated Targeting for the MER Opportunity Rover

Tara Estlin, Benjamin Bornstein, Daniel Gaines, David R. Thompson, Rebecca Castaño, Robert C. Anderson, Charles de Granville, Michael Burl, Michele Judd, and Steve Chien

Jet Propulsion Laboratory, California Institute of Technology, Pasadena, CA 91109
e-mail: {*firstname.lastname*}@jpl.nasa.gov

Abstract

The Autonomous Exploration for Gathering Increased Science (AEGIS) system provides automated data collection for planetary rovers. AEGIS was uploaded to the Mars Exploration Rover (MER) mission Opportunity rover in December 2009 to provide automated targeting capabilities for remote sensing instruments. Geological targets for rover remote-sensing instruments, especially narrow field-of-view instruments (such as the MER Mini-TES spectrometer or the 2011 Mars Science Laboratory (MSL) Mission ChemCam Spectrometer), are typically selected manually based on imagery that has already been transmitted back to the operations team on Earth. AEGIS enables the rover flight software to analyze imagery *onboard* through the use of onboard data analysis techniques. This system allows the rover to autonomously select and sequence targeted observations in an opportunistic fashion. In this paper we provide an overview of the AEGIS automated targeting capability and describe how it is currently being used onboard the MER mission Opportunity rover.

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 storms, to continue collecting data. The 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.

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 100m in the entire mission, the MER rovers 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 (~12 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 in a reasonable time frame.

Communications bandwidth has not grown as fast as rover traverse range. As this trend in increased mobility continues, 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 observed or 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 narrow field-of-view instruments, must be selected manually based on imagery already on the ground with the operations team. Examples of these instruments include

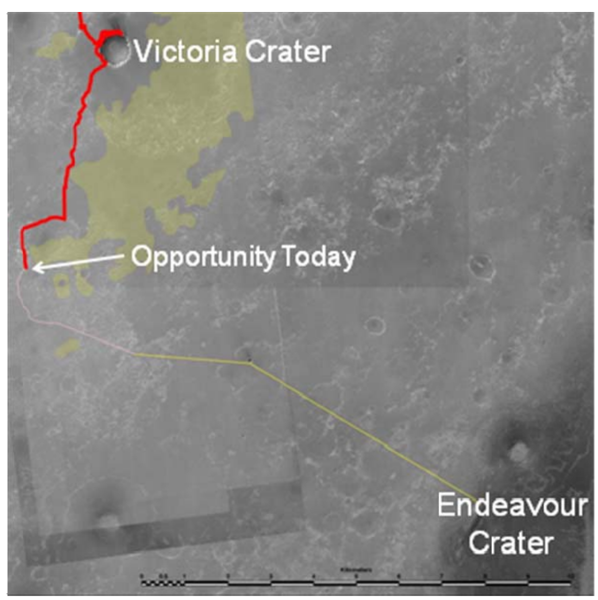

Figure 1. Planned path to Endeavour crater on Mars. Currently the Opportunity rover is en route to Endeavour, which is over 12 kilometers (or 7 miles) away from Opportunity's current position. To reach this target many long drives will be needed.

the MER Miniature Thermal Emission Spectrometer (Mini-TES) and the 2011 MSL Laser-Induced Remote Sensing for Chemistry and Micro-Imaging (ChemCam) spectrometer. AEGIS enables the rover flight software to analyze imagery onboard in order to autonomously select and sequence targeted remote-sensing observations in an opportunistic fashion. AEGIS operates by analyzing MER navigation camera images to identify terrain features of interest, such as rocks with certain characteristics. Scientists on the ground specify these target characteristics during the sequencing process. For example, scientists could request measurements of large rocks with light albedo (i.e, lightly colored). Once a target is identified onboard in a navigation camera image, its location is determined and a remote sensing instrument re-pointed to collect high-resolution followup data. AEGIS can run during and after traverses when the rover enters new terrain but has not yet downlinked images of the area 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 has passed checkout tests onboard Opportunity and is fully operational. Since the MER Mini-TES spectrometer is no longer functioning on Opportunity and is expected to have only a few remaining measurements on the Spirit MER rover, operators currently use AEGIS for targeting sub-framed, high quality (i.e., multiple filter) images with the MER panoramic cameras.

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 originally tested before upload to MER. Finally we provide sample results from AEGIS's use onboard the MER Opportunity rover.

2 AEGIS System Overview

The AEGIS system provides autonomous deployment of science instruments that target specific terrain features. A number of rover remote sensing instruments, such as the MER Mini-TES spectrometer, have a very narrow field-of-view and thus require selection of specific focused targets for sampling. Selecting targets for these instruments by mission personnel on Earth is currently a lengthy process. Typically operators will manually identify the targets in images that have already downloaded on a previous sol. These context images are collected with wide field of view 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 an end-of-day location, the rover performs only untargeted data collection until the context images can be analyzed and new measurement commands uplinked. At best this will happen on the next sol, but it may never happen if it's decided the rover should immediately proceed to a new location due to other goals or engineering constraints.

By analyzing image data onboard, AEGIS can autonomously select targets for these instruments and execute a set of measurement activities. These techniques could be used, for example, on the MSL mission to select targets for the ChemCam instrument to sample at the end of a long rover drive. For MER, AEGIS is being illustrated by taking additional measurements with the MER panoramic camera[1] in a quarter-frame high-quality (multiple filter) mode, which uses a 4 degree FOV.

AEGIS is run as part of the MER onboard flight software. 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. AEGIS was required to run using less than 4 MB of RAM to ensure other onboard processes were not impacted.

AEGIS was originally developed as part of a large autonomous science framework called OASIS (Onboard Autonomous Science Investigation System) [1]. OASIS is designed to operate onboard a rover identifying and reacting to serendipitous science opportunities. OASIS analyzes data the rover gathers, and then using machine learning techniques, prioritizes the data based on criteria set by the science team. This prioritization can organize data for transmission back to Earth or search for specific targets specified by the science team. If one of these 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

[1]The MER Mini-TES spectrometer is another example of a limited FOV instrument where the AEGIS automated targeting technology would be very 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 no longer functional.

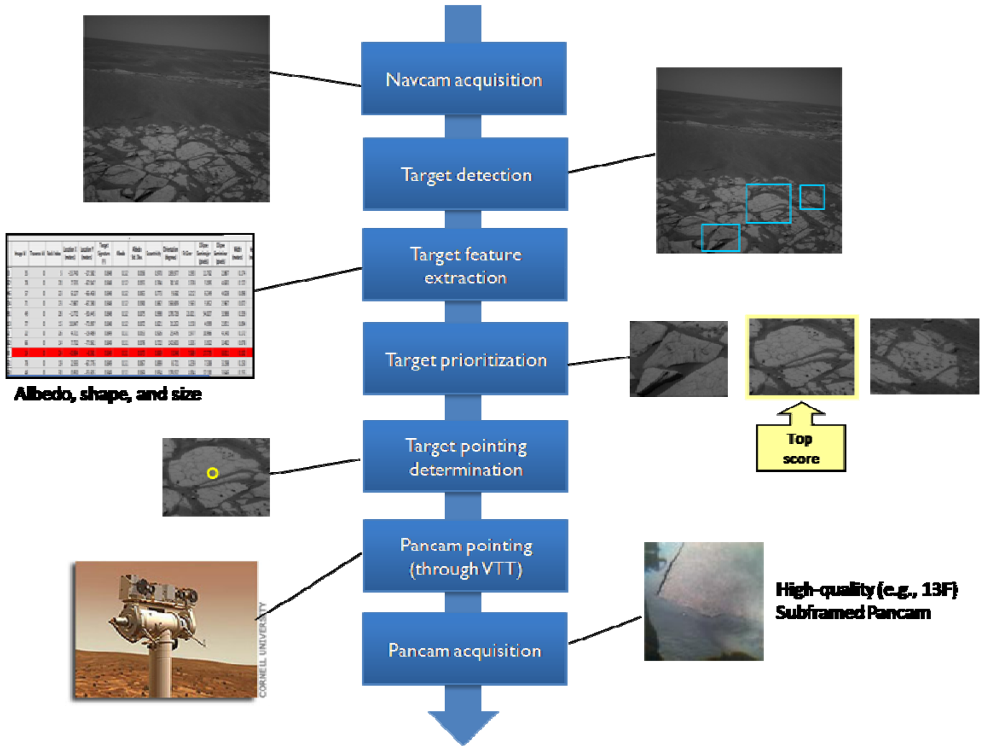

Figure 2. AEGIS Process Pipeline. When AEGIS is sequenced, the above series of steps is executed onboard the MER Opportunity rover. Parameters can be set during sequencing to specify navigation camera pointing, the "target rock signature" (e.g., rocks of large size and low albedo), and settings for the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

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 intensity edges in grayscale imagery. This algorithm is further detailed in Section 3.
- **Extract relevant target features:** AEGIS calculates a set of target features (or properties) for each candidate rock. These properties include measures of size, albedo, and shape.
- **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 signatures are "high albedo", "round shape", "large rocks with low albedo", 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 then acquires additional data with the panoramic cameras. The ground sequencing team can 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.

3 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

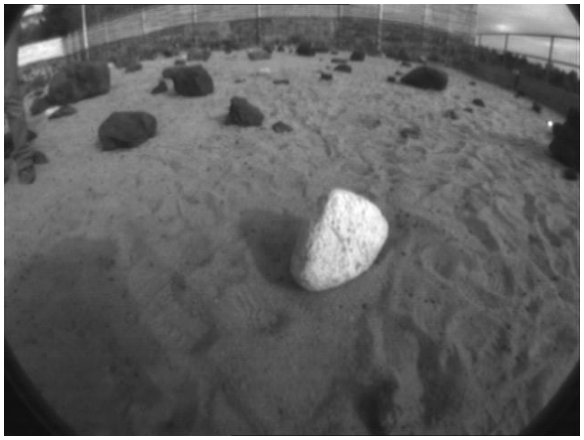

Figure 3. Detected rock of high albedo. Above is an image taken with the JPL FIDO rover hazard camera (in JPL Mars Yard) when searching for rocks of "high albedo" (or light-colored).

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 crater or other terrain features. Rockster initially locates partial boundary contours of targets using a procedure similar to the Canny edge detector. Specifically, it 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. Due to the limited processing capacity and memory available onboard the MER rovers, Rockster emphasizes techniques that could perform quickly and robustly in such an environment.

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. Potential T-junctions that were missed by the edge detector are identified and used to further split fragments into even smaller pieces. A gap-filling mechanism 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.

For automated targeting of limited FOV instruments, false detections are costly and high precision is important. Thus for this application, Rockster is 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. It could be judged for some application that is it more important to find a larger percentage of overall targets at the risk of returning additional false positive (i.e., non-interesting targets) more often. The sequencing team can also choose to limit target detection to specific rectangular subregions of the image. This might be useful for excluding image regions that contain the deployed rover arm, rover tracks, or other features that could generate spurious detections.

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 all as identifying atmospheric targets, such as clouds and dust-devils, in MER imagery [2].

4 Target Feature Extraction

Once candidate targets are identified, the AEGIS system computes numerical attributes corresponding to properties of each target image region.

Albedo: The albedo of a feature is an indicator of the reflectance properties of a surface. The reflectance properties of a rock can provide important information about its mineralogical composition. AEGIS measures albedo 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.

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. Another 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 [3,4].

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

Figure 4. AEGIS target selection in MER images. In a), AEGIS was directed to select a large rock as a potential target for a limited FOV instruments, such as the 2011 MSL ChemCam Spectrometer, to sample. In b) 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 rover navigation cameras. Autonomously selecting targets vs. blind sampling greatly increase the chances of accurately targeting a rock.

of the perimeter divided by the contour's pixel area [5].

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.

5 Target Prioritization

To prioritize candidate targets and select a top target, AEGIS recognizes pre-specified *target signatures* corresponding to particular feature attributes the science team identifies in advance. This algorithm enables scientists to efficiently and easily stipulate the value of each particular feature. Figure 3 shows an example using target signature in early tests with the JPL FIDO rover to look for rocks of high albedo. 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 β=0. We have:

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

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 images in the extreme 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 are easily detected as targets.

AEGIS selects the top target from the filtered candidates and uses it to acquire new targeted data. Figure 4 shows two potential targets selected in MER navigation camera image with selection criteria favoring size and shape. In both cases the AEGIS target finder is run in a mode that decreases the number of false positives (e.g., finding a shadowed area of sand vs. a rock) but finds fewer of the overall rocks in the image.

6 New Data Acquisition

Once AEGIS selects a top target, it commands additional science measurements. First, it selects a

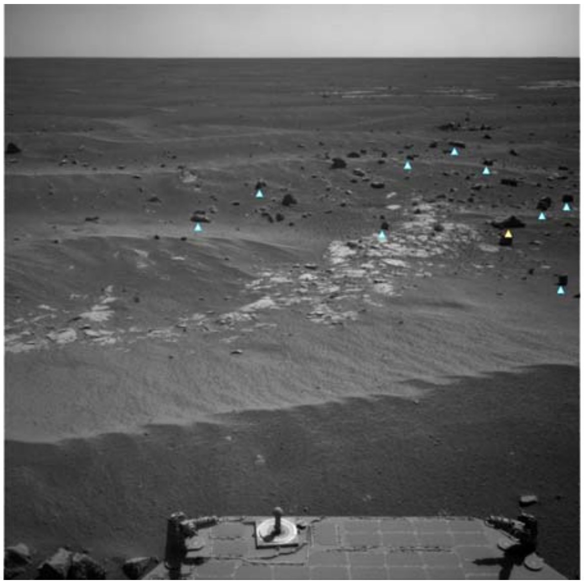

Figure 5. Top targets selected in MER navigation camera image. This figure shows the results of image analysis from the first full AEGIS run on the MER Opportunity rover. In this run, AEGIS was told to look for large, dark rocks. The yellow marker shows the top target and the blue markers show the top ten targets.

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 used to collect multiple filter, subframed images with both left and right cameras. This approach provides high quality images that are inexpensive 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 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 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 [2]. 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.

Figure 6. 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.

7 AEGIS Results on 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 were performed to exercise different AEGIS components. All of these checkout steps executed successfully. Figure 5 and 6 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 dark albedo. Figure 5 shows the top target selected (indicated by the yellow marker) and the top ten targets (indicated by the blue markers). Note that AEGIS does not downlink the detected target contours due to telemetry constraints.

Figure 6 shows the resulting panoramic camera

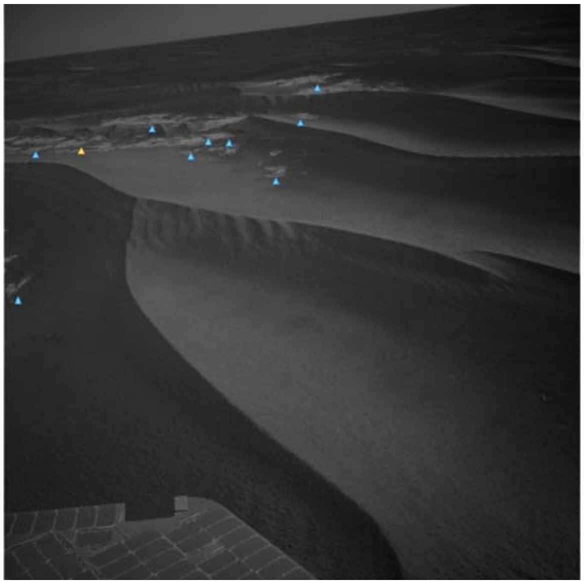

Figure 7. Top targets selected in MER navigation camera image. This figure shows the results of image analysis from an AEGIS run performed in April 2010 on the MER Opportunity rover. In this run, AEGIS was told to look for large rocks. The yellow marker shows the top target and the blue markers show the top ten targets.

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 May 2010, AEGIS has been run five times on Mars (Table 1) and has consistently chosen appropriate targets based on the specified target signature profile. On most runs AEGIS has either successfully detected either rock outcrop or loose rocks that were likely crater ejecta. For example, Figure 7 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 known behavior and in some cases can be arguably appropriate since interesting MER science discoveries have been made on disturbed material in rover tracks. However in this particular case, tracks were not expected in the image (the drive faulted out early and in an unexpected location) thus the track detections could be considered false positives. 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.

8 Related Work

Several systems have been developed and demonstrated for autonomous rover science operations. One such system was successful at autonomously identifying meteorites in Antarctica [6,7] and another provided techniques for analyzing field test data for the Marsokhod rover [8]. Another related system provided classification of features and survey techniques in association with the automated identification of life in the Atacama desert [9]. More recently researchers have investigated automated target selection for the upcoming 2013 ESA ExoMars rover mission [10]. Typically these systems have used different vision and data analysis techniques to detect and prioritize targets. For instance, [10] uses a graph-based growing algorithm to separate rocks from terrain and then edge detection techniques to identify when layering or bedrock is present. Another difference between AEGIS and these prior systems is that the integration with flight software and deployment to a space mission imposes challenging constraints on processing power and memory.

9 Future Work and Conclusions

In future work, we plan to expand AEGIS on several fronts. One 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 [11] to collect the final measurement. Another area of inquiry would use data collected from other types of instruments (such as spectrometers or ground-penetrating radar) to further enable AEGIS to select interesting science targets.

In summary, AEGIS enables targets of scientific interest to be automatically recognized 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.

Sol	Selection Profile	Rock detections	Non-rock detections	Notes
2138	*f = (inscribed circle radius)*	10	0	Checkout test. Outcrop detected.
2172	*f = (inscribed circle radius) – 0.5 (albedo)*	10	0	Checkout test. Loose rock detected.
2204	*f = (inscribed circle radius)*	0	10	Tracks detected. No rocks in scene.
2221	*f = (inscribed circle radius)*	10	0	Outcrop detected.
2247	*f = (inscribed circle radius)*	10	0	Outcrop detected.

Table 1: AEGIS runs on the Opportunity rover as of the time of publication. Note, each run can find a maximum of ten targets. 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 rover tracks were visible; all detected features were soil disruptions due to rover tracks.

10 Acknowledgement

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.

References

[1] R. Castano, T. Estlin, R. C. Anderson, D. Gaines, A. Castano, B. Bornstein, C. Chouinard, M. Judd, “OASIS: Onboard Autonomous Science Investigation System for Opportunistic Rover Science,” *Journal of Field Robotics*, Vol 24, No. 5, May 2007.

[2] R. Castano, T. Estlin, D. Gaines, A. Castano, C. Chouinard, B. Bornstein, R. C. Anderson, S. Chien, A. Fukunaga, and M. Judd, “Opportunistic Rover Science: Finding and Reacting to Rocks, Clouds and Dust Devils,” *Proceedings of the 2006 IEEE Aerospace Conference*, Big Sky, Montana, March 2006.

[3] A. W. Fitzgibbon, A. M. Pilu, and R. B. Fisher, R. B. "Direct least-squares fitting of ellipses". *Pattern Analysis and Machine Intelligence*, 21(5): 476--480, May 1999.

[4] R. Halir and J. Flusser, "Numerically stable direct least squares fitting of ellipses," *Proceedings of the 6th International Conference in Central Europe on Computer Graphics, Visualization and Interactive Digital Media*, 1: 125--132, February 1998

[5] L. M. Hentschel and N. W. Page, N.W., “Selection of Descriptors for Particle Shape Characterization,” *Particle & Particle Systems Characterization* 20 (1), pp. 25-38. 2002.

[6] L. Pedersen, “Robotic Rock Classification and Autonomous Exploration,” PhD thesis, Robotics Institute, Carnegie Mellon University, CMU-RI-TR-01-14, 2001.

[7] M. Wagner, M. D. Apostolopoulos, K. Shillcutt, B. Shamah, R. Simmons, and W. Whittaker, “The science autonomy system of the Nomad Robot,” *Proceedings of the International Conference on Robotics and Automation (ICRA 2001)*, Seoul, Korea May 2001.

[8] V. Gulick, R. Morris, M. Ruzon, and T. Roush, “Autonomous image analysis during the 1999 Marsokhod rover field test,” *Journal of Geophysical Research*, 106(E4) 2001.

[9] D. R. Thompson, S. Niekum, T. Smith, and D. Wettergreen, “Automatic detection and classification of features of geologic interest,”. *Proceedings of the 2005 IEEE Aerospace Conference*, Big Sky, MT, March 2009.

[10] M. Woods, A. Shaw, D. Barnes, D. Price, D. Long, and D. Pullan, “Autonomous science for an ExoMars Rover-like mission,” *Journal of Field Robotics*, 26:4, April 2009.

[11] M. Bajracharya, A. Diaz-Calderon, M. Robinson, M. Powell, "Target Tracking, Approach, and Camera Handoff for Automated Instrument Placement," *Proceedings of the 2005 IEEE Aerospace Conference*, Big Sky, Montana, March 2005.