

Abstract. Missions to Mars already under study call for very long autonomous traverses (on the order of hundreds of meters per sol) during which there will be limited contact with Earth. The Onboard Autonomous Science Investigation System (OASIS) is a technology for increasing science return during rover traverses by prioritizing data onboard, and identifying and reacting to unanticipated science opportunities[1]. By prioritizing data for downlink onboard, it is expected that the set of images selected for downlink by OASIS will have a consistently higher scientific interest than any set of equal number of images of the same traverse obtained using random or periodic sampling. Thus, OASIS can be used to increase the science returned from a long traverse.

OASIS prioritizes data and identifies science opportunities by extracting information from image data and then using this information to prioritize the data and identify potential new science opportunities. Currently the information extracted from the images is the location of the rocks in the image and several properties of these rocks such as size, albedo, texture and shape. In this paper we give an overview of the role of rock detection in the OASIS system and describe the functionality of the rock detector. We then present results on representative images.

Introduction. Rovers offer scientists the ability to move around a planetary surface and explore different areas of interest. The farther the rover can travel, the greater the opportunity exists for increased scientific discovery. Downlink bandwidth available to a mission, however, does not typically increase proportionally to the distance travelled. This means that with a greater distance traveled per sol, there will be less data (fewer bits) per meter of traverse returned to Earth. The OASIS system is one mechanism for ensuring that the data with high scientific value reaches scientists on the ground.

The OASIS system. OASIS consists of an information extraction module, a data analysis and prioritization module and a planning and scheduling module.

The information extraction module enables extraction of features of interest from collected images of the surrounding terrain. This module includes a perception component and a feature analysis component. The perception component identifies objects present in the image. It receives images from the rover cameras and outputs areas of the image corresponding to different aspects of the scene such as the sky, the soil and rocks. The output is used by a feature analysis module to extract features for the identified objects or regions. The rock finder is one aspect of the perception component. In addition to extracting the features of the rocks, results from the rock finder can be used to estimate rock distributions, create rock maps and count rocks in a region.

The data analysis and prioritization module uses the information from the feature extraction module to assess the scientific value of the planetary scene and to generate new science objectives that will further contribute to this assessment.

This component consists of three separate prioritization algorithms that analyze the collected data and prioritize the rocks. The results from these three algorithms are then fed into a unified prioritization algorithm that provides a prioritized list of images for downlink. A new set of observation goals is also generated to gather further data on rocks that were ranked as high priority.

The planning and scheduling module enables dynamic modification of the current rover command sequence (or plan) to accommodate new science requests from the data analysis unit. This component uses a continuous planning approach to iteratively adjust the plan as new goals and/or faults occur.

Rock Detection. The perception module of OASIS identifies objects present in the image and maps them to symbolic objects, i.e., it interprets the image and describes, at a very high level, its contents. In general, this task is extremely complex and there is no methodology to perform this interpretation for an arbitrary image. In every case, the success of a program in interpreting an image depends on the constraints on the subject, sensor, and environment. Consider the case of face detection where subject (e.g., human face), sensor (e.g., camera) and environment (e.g., imaging conditions) are well determined. The constraints reduce issues with varied image resolution, face pose, illumination, expression, and other parameters. Even with these restrictions, multiple sensor modalities are often used to facilitate the recovery of some critical features, e.g., color simplifies the identification of lips, stereo range eases the determination of the nose, etc. These considerations highlight some of the challenges encountered in designing a robust algorithm for detecting rocks.

The task of interpreting a Mars scene is highly unconstrained. Although it is true that many Mars scenes consist exclusively of sky, soil and rocks, it is also true that there is no strong constraint on the environment, the sensor, or the subject. With respect to the environment, the images may be taken under any ambient illumination: rocks may have a long shadow or exhibit no shadow at all. With respect to the cameras, the images may be taken with any focal length covering anything from a wide angle, which leads to poor rock resolvability, to a very narrow angle, where a single rock might cover a large portion of the image. Likewise, the resolution of the images is not fixed: some images may have high resolutions (e.g., pan-cam) while others may have low resolutions (e.g., haz-cams). Finally, with respect to the subject, rocks may appear in any part of the image, have any size, albedo, shape, texture; they may be completely contained within the image or not, might be present in groups or isolated, might occlude other rocks and might be (and probably are) covered by dust that makes them appear more similar to a dirty sand patch than to a rock. To these difficulties, we must add that, as OASIS is being designed to run autonomously, it is imperative that the parameters of the rock detector do not depend on the image, regardless of the conditions under which this was acquired, i.e.,

the system must not require any parameter tuning (a liability of the previous rock detector used by Oasis[2]). Given these observations, it is clear that the success rate of a rock detector will vary widely, from complete success to complete failure, depending on the particular environmental conditions, sensor and subjects.

From the many sensor modalities expected to be available in a rover, we have chosen to initially base the rock detector on analysis of intensity of single greyscale images. Thermal infrared, suitable for detection of large rocks due to their thermal inertia, would restrict the algorithm to work on images from the pan-cam, the only set with a filter wheel; this same argument applies for the use of color. The use of stereo is highly desirable for detection of large rocks but it is unsuitable for detection of small rocks and pebbles; also, it could not be applied to imagers that do not come in stereo pairs, as is the case for the microscopic imager (MI). In contrast, rock detection on single greyscale images applies directly to analysis of low-res haz-cam, high-res nav-cams, any spectral band (or combination) of the pan-cam and the MI camera.

The detection of small rocks is carried out by finding small closed shapes in the image. The image is initially normalized, filtered with an edge preserving smoother[3] and its edges are enhanced using unmask sharpening. The edges of the resulting image are detected using both a Sobel and a Canny edge detectors [4]. For each result, we search for small closed shapes (which presumably correspond to small homogeneous regions, i.e., pebbles) using an edge-walker. The results from both detectors are combined and output as a list of contours of the found shapes.

The detection of large rocks is an extension of the detection of small rocks. On high resolution images, the exact contours of a large rocks might be very complex and the albedo of the rock might not be homogeneous due to both variations in the rock composition and accumulations of dust or sand. However, all this detail is present because of the scale at which the rock was imaged. If we reduce the resolution of the image in half, the rock will lose its detail but it would still be recognizable as a rock. Thus, the strategy of the rock detector consists of building a pyramid of images with the original image as the base of the pyramid [5]. At every level of the pyramid, each image has half the width and height of the image immediately below it, i.e., image resolution is reduced as we move toward high levels of the pyramid while it is increased as we move down, toward low levels of the pyramid. Thus, large rocks present in the high resolution image (i.e., at the base of the pyramid) are mapped into small rocks at high levels of the pyramid, where they can be found using the small-rock detector.

The detection of large rocks uses the detection of small rocks in images where the resolution has been reduced, sometimes substantially. To recover the contour of the large rock, at high resolution, we proceed as follows. Starting with the highest level of the pyramid (lowest resolution), we double the size of each rock found and map it to the level of the pyramid immediately below it. For each rock thus magnified, we can have two situations. On one hand, the rock might have been detected at both levels in which case we keep the contour de-

tected at the highest resolution level and dismiss the other. On the other hand, the rock detected in the low resolution image (high level of the pyramid) had not been detected in the high resolution image (low level of the pyramid), in which case, the low resolution rock needs to be added to the list of rocks of the high resolution image; in this case, the low resolution approximate contour is refitted to the high resolution image using active contours (a.k.a. snakes), an energy-minimization procedure whose function is to reshape an approximate contour toward a real contour using an image as a guide [4]. This procedure is done once at each level of the pyramid, starting with its highest level. The final result is a set of full-resolution contours of rocks of all sizes.

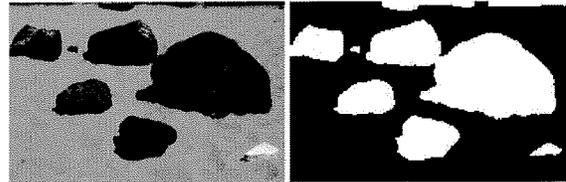


Figure 1: Mars yard scene and its 2-D map of detected rocks

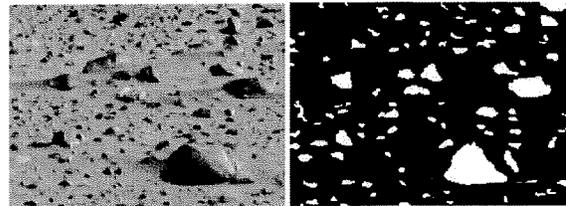


Figure 2: Gusev crater scene and its 2-D map of detected rocks (Public image PIA04997, NASA/JPL/Cornell)

Results: Figure 1 shows a sample image and the 2-D map of the rocks found. There are several points worth noting. First, this scene was processed automatically, with no tuning of any type; normalization and filtering takes care of photometric variations and the scale-based search accounts for variations in rock size. Second, the program is looking for homogeneous patches and thus it errs when there are shadows attached to dark rocks or when there is too much sand covering a light rock. These issues can be solved using texture analysis but this restricts the solution to the analysis of high resolution images. Third, the program does its search on the image, not on a 3-D estimate of the world, and thus, it cannot handle occlusions. This can be solved using stereo but this restricts the solution to the analysis of scenes imaged using stereo rigs. In spite of these limitations, the program is successful at finding rocks in many scenes that exhibit a low-to-medium rock density and where the rocks can be easily distinguished from the ground. In particular, the images from Spirit show that these are characteristics of the scenes imaged at the Gusev crater (e.g., Fig. 2).

References: [1] Anderson R. et al, LPSC (2003), [2] Gor V. et al, AIAA (2001) [3] Tomasi C. and Manduchi R., ICCV (1998), [4] Trucco E. and Verri A., "Introductory Techniques for 3D Computer Vision", Prentice-Hall (1998), [5] Rosenfeld A., "Multiresolution Image Processing", Springer-Verlag (1984),