Passive Imaging Based Multi-cue Hazard Detection for Spacecraft Safe Landing

Andres Huertas, Yang Cheng, Richard Madison
Jet Propulsion Laboratory, California Institute of Technology
4800 Oak Grove Drive, Pasadena, California 91109
Andres.Huertas|Yang.Cheng|Richard.W.Madison@jpl.nasa.gov

Abstract—Accurate assessment of potentially damaging ground hazards during the spacecraft EDL (Entry, Descent, and Landing) phase is crucial to insure a high probability of safe landing. A lander that encounters a large rock, falls off a cliff, or tips over on a steep slope can sustain mission-ending damage. Guided entry is expected to shrink landing ellipses from 100-300 km to ~10 km radius for the second-generation landers as early as 2009. Regardless of size and location, however, landing ellipses will almost always contain hazards such as craters, discontinuities, steep slopes, and large rocks. It is estimated that an MSL (Mars Science Laboratory)-sized lander should detect and avoid 16-150m diameter craters, vertical drops similar to the edges of 16m or 3.75m diameter crater, for high and low altitude HDA (Hazard Detection and Avoidance) respectively. It should also be able to detect slopes 20° or steeper, and rocks 0.75m or taller. In this paper[1,2] we will present a passive imaging based, multi-cue hazard detection and avoidance (HDA) system suitable for Martian and other lander missions. This is the first passively imaged HDA system that seamlessly integrates multiple algorithms—crater detection, slope estimation, rock detection and texture analysis, and multi-cues—crater morphology, rock distribution, to detect these hazards in real time.

TABLE OF CONTENTS

1. INTRODUCTION 1
2. PREVIOUS WORK 2
3. APPROACH 3
4. HOMOGRAPHICAL SLOPE ESTIMATION (HSE) ... 3
5. CRATER HAZARD DETECTION (CHD) 4
6. ROCK DETECTION (RD) 5
7. TEXTURE ANALYSIS (TA) 11
8. LANDING SITE SELECTION 11
9. REAL-TIME CONSIDERATIONS 12
10. FUTURE WORK 12
ACKNOWLEDGEMENT 12
REFERENCES 12

1. INTRODUCTION

With the current EDL (Entry, Descent, and Landing) capability, a spacecraft lands somewhere within a very large landing ellipse. The Viking landers (1976), Mars Pathfinder (1997), Mars Polar Lander (1999), and Mars Exploration Rovers (2003), had landing ellipses on the order of 100-300 km long. Guided entry is expected to shrink landing ellipses to ~10 km in radius for the second-generation landers as early as 2009. However, even if a landing ellipse is only a few kilometers long, it is very likely to contain hazards such as craters, discontinuities, steep slopes, and large rocks, regardless of how the ellipse is selected. A lander that encounters a large rock, falls off a cliff, or tips over on a steep slope can sustain mission-fatal damage.

Most importantly, the scientifically interesting sites are most likely near craters, ridges, fissures, and other relevant geological formations, the same features that can be considered hazards. To ensure a safe landing one must either choose a precise landing at a pre-selected safe site, or real-time HDA or a combination of both.

An HDA system can use an active sensor, such as scanning laser radar (LIDAR) or phased array terrain radar (PATR), or, it can use a passive sensor such as a camera. Active sensors tend to be favored because they can directly measure the depth of sensed terrain and are less sensitive to atmospheric opacity. Also, the algorithms to interpolate these data are relatively simple and fast. Active sensors however are costly, heavy (6-25 kg), and have high power requirements (40-200 w). Active sensors have low resolution (40x40-100x100), a narrow FOV (field-of-view) (15°-40°) and high volume (2-40 L). In contrast, flight-qualified passive sensors (cameras) are affordable and lighter (0.3 kg), consume less power (13 w), and have higher resolution (1024x1024), a wider FOVs (e.g. 120°), low volume (2 L) and a greater sensing range. Passive sensors however, work only during the daytime and algorithms may have seasonal requirements to handle dependencies on sun position and signal contrast attenuation from dust storms. Such storms on Mars, however, are seasonal and regional, and even under such conditions, the chance of high levels of obscuration are less than 1% [1].

The four types of hazards under consideration are shown in Figure 1. **Craters** of varied sizes are ubiquitous on places like Mars. Of most concern are the small impact craters (< 5 km in diameter). A lander can fall off a crater rim, tip over on a steep slope in the crater bowl, have limited communications while sitting at the bottom of the bowl or become trapped inside. Smaller missions that may land with

1———————————

[1] 0-7803-9546-8/06/$20.00© 2006 IEEE
[2] IEEEAC paper #1434, Version 2, Updated Oct, 25 2005

significant horizontal velocity could also damage themselves by impacting a steep crater wall. For an MSL-sized lander, safe landing guidelines require HDA to detect 16-150 m diameter craters. **Steep Slopes** are large, relatively non-level planar regions. Current guidelines dictate that slopes >20° should be detected and avoided. Landers can tip over, take damage and pin parts in unusable positions. Steep slopes may mislead a radar altimeter that interprets first-return as vertical. Underestimates like these can cause propellant mis-management. **Discontinuities** in terrain are areas where elevation changes significantly over a short distance, as in a cliff or ridge. A small lander with significant horizontal velocity could impact a discontinuity or tumble over it and impact the ground below. A larger lander could tip while landing on a discontinuity or while deploying mechanisms after landing. MSL-type safe landing guidelines indicates discontinuities similar in size to the wall of a 16 m-diameter crater (from high altitude) or a 3.75 m-diameter crater (from low altitude) should be detectable.

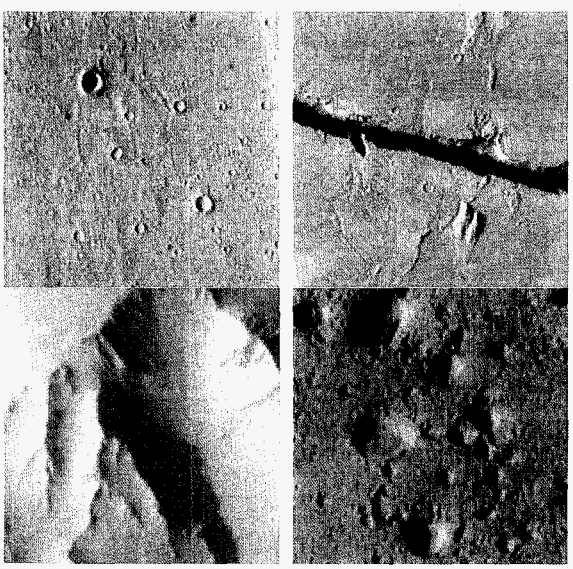

Figure 1. Craters, Discontinuities, slopes, and rocks.

Rocks that are not large enough to tip a lander are still a concern. A legged-lander may high-center on a large rock during landing, causing an unstable landing and may impact the lander underbelly and the foot stabilizers. MSL-type HDA must detect rocks 0.75m tall and safe landing guidelines require 99% detection rates and < 1% false alarm rates.

Current vision-based HDA algorithms are slow and exhibit incomplete detections. In this paper we present an integrated HDA system that improves on both aspects. The algorithms for detection of craters and slopes have been described in detail elsewhere [2, 3] and are only summarized here. The more recent developments on rock detection are described in detail using laboratory imagery, and earth imagery of similar characteristics as those on Mars.

2. Previous Work

The only vision system used for planetary spacecraft safe landing is the Descent Image Motion Estimation system (DIMES) for the Mars Exploration Rovers (MER) mission [4]. DIMES uses two descent images, and IMU and an altimeter to estimate horizontal velocity of descending spacecraft [4, 5]. No space mission has however attempted autonomous, vision-guided safe and precise landing, and we are not aware of a complete vision-based safe-landing system which can satisfy, say, Martian EDL requirements. However, several relevant developments have been reported. They are illustrated and compared in Table 1. The speed performance column in the table does not represent a formal comparison but gives an idea of the order of magnitude of performance. A visual method supported by an angular sensor to provide position measurements was suggested in [6]. A texture analysis (TA) scheme for autonomous helicopter safe landing was reported in [7]. A structure-from-motion (SFM) method for rock and slope HDA was proposed in [6, 7]. SFM reconstructs 3-D surface topography from multiple descent images in four steps. First, feature "windows" are selected in the first image. Second, matching windows are located in later images. Third, the spacecraft motion between the images is estimated independently. Fourth, triangulation is used to determine the 3-D position of the feature windows that accounts for their 2-D motion under the estimated spacecraft motion. Landing hazards such as the craters, steep slopes, and rocks, can be detected in the resulting 3-D surface map. Current approaches of SFM have computational costs that make unsuitable for reconstruction of natural terrain during landing; it typically can take several seconds to process a pair of images. Another deficiency is that the algorithms cannot recover the surface model at the focus of expansion of a moving camera

A homography-based slope estimation (HSE) scheme was suggested in [10] by one of the authors and it is summarized below in Section 4. Shape-from-shading (SFS) [11] recovers surface shape from an image and a known lighting direction. Under consistent illumination, a Lambertian surface with constant albedo reflects intensity as a function of the angle θ between surface normal and lighting direction. One variation of SFS maps image intensity to θ using a scene of known geometry, and then uses the map to quickly recover θ at each pixel in later images. It then finds surface normals at each pixel that maximize some constraint, such as surface smoothness, subject to the θs. Finding surface normals can be slow and may underestimate the texture of a surface. In order to reconstruct the surface well, the surface albedo must to be uniform, which is not the case for Mars [12. 13]. The large albedo variation of Martian surface and slow speed makes the SFS a less desirable option for EDL. However, some of its concepts are useful to help detect slopes.

In summary, these approaches suffer from low speed and incomplete detection. The high resolution, say 1024x1024,

of a camera facilitates hazard detection and image based HDA algorithms typically require two images to reconstruct a surface model. Current flight processors are slow and to handle such a large data volume in real-time will require some innovation. Another drawback of current algorithms is that they focus on detecting one type of hazard, whereas the Multi-cue HDA (MC-HDA) safe landing system described here detects all types of potentially fatal hazards.

Table 1. State-of-the-Art in Hazard Detection

Hazard/ Method	Crater	Slope	Discont.	Rock	Coverage	Speed
SFM	Y	Y	Y	Y	N	<0.015Hz
HSE	N	Y	N	?	N	<0.1 Hz
CHD	Y	N	?	N	Y	<0.1 Hz
TA	N	N	Y	Y	Y	~1.0 Hz
SFS	Y	Y	Y	N	Y	<0.001Hz
MC-HDA	Y	Y	Y	Y	Y	>0.5 Hz

Lastly, the detection of shadows from overhead imagery has been shown to be very useful to detect and verify detection of building structures [14, 15] from monocular and multiple aerial images. We are not aware of a system that use shadows to detect natural rocks from aerial images.

3. APPROACH

The Multi-Cue Hazard Detection system operates in four stages as illustrated in Figure 2:

A. Survey stage. During early parachute descent, the system starts analyzing descent imagery. Initially only large, salient, craters are quickly located. Next, the intensity gradient across these craters is compared against an ideal crater model [16] to calibrate a "slope-from-shading" and "intensity-to-angle" map. This map in effect relates surface topography to pixel intensities.

B. Regional hazard detection stage. During late parachute descent or ensuing powered descent, the system detects regional hazards, such as large craters, discontinuities, and steep slopes. This involves three algorithms. First, the system uses the intensity-to-angle map calibrated in the survey stage to identify pixels on steep slopes facing toward or away from the sun. These areas may represent ridges, hills, or deformed craters. Second, the system carries out crater detection to identify craters. It uses the crater slope model and rock distribution models to identify hazards on the outer slope of the craters. Results of the first two algorithms are used to mask off detected hazards, reducing the search scope (and processing time) of the detection algorithms. Third, the system employs image homography techniques to estimate the surface flatness and slopes of potential landing sites. From these three algorithms, the best landing region, typically much bigger than the landing site, is selected and then the spacecraft can proceed to maneuver toward the best region.

C. Local hazard detection stage. In the final descent, the system detects local hazards such as small craters and rocks. This involves three algorithms. First, texture analysis looks for discontinuities and rocks in areas that have not yet been identified as hazards. This step uses image histograms, crater and rock distribution models, and similar metrics found and carried over from the first two stages, saving computation time. Second, rock detection algorithms look for large rocks. The first uses a combination of texture analysis and the intensity-to-slope map for detecting steep rock faces. The second detects individual rocks from their shadows. Third, HSE searches supposedly hazard-free areas to identify additional slopes that slope from shading and crater detection did not locate earlier. To improve performance potential landing sites can be chosen and ranked before some algorithms are applied to eliminate unacceptable candidates.

D. Site selection stage. A site selection algorithm typically chooses the "safest" site on a hazard map. Previous work [17, 18] has investigated combining multiple information sources into a hazard map for this purpose. Our system does not produce a hazard map. It identifies candidate sites early and aims to maintain speed performance by selectively using algorithms to eliminate hazard areas and selecting only promising areas for further processing.

4. HOMOGRAPHY SLOPE ESTIMATION (HSE)

HSE estimates the slope of surface patches seen in a pair of images using the homography transform, depicted in Figure 3, an example of the HSE of the Mars Exploration Rovers (MER)-A descent images. The homography coordinate transform describes how translation, rotation, and perspective projection of a planar surface patch modify the shape and location of patch from the first image to the second. The transform parameters encode the surface normal of the patch and the motion of the camera between images (See Figure 3.)

An HSE algorithm [8] has been developed with support from the JPL's Mars Technology Program. The algorithm operates in three steps. First, relatively featureless (*i.e.,* planar) patches, which should make good landing sites, are selected from one image.

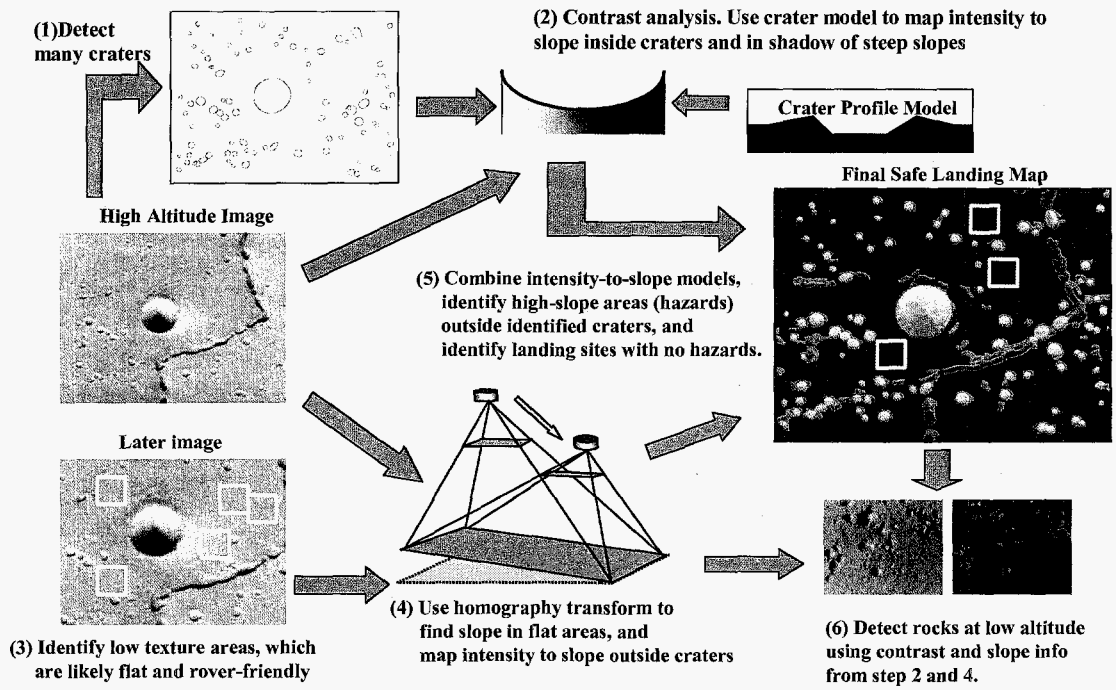

Figure 2. The Multi-Cue Hazard Detection algorithm operates in multiple stages.

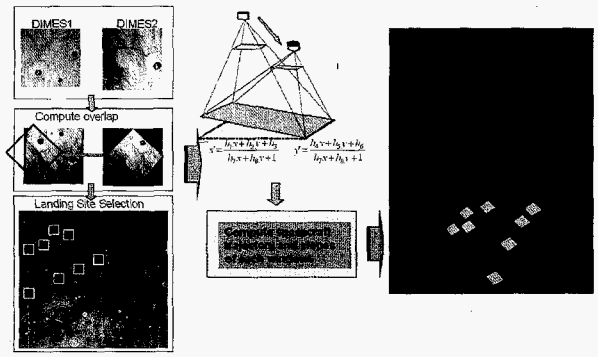

Figure 3. Homography-based slope estimation (HSE)

Secondly, a Levenberg-Marquardt iterative algorithm [19] is applied to locate the patches in the second image and the homography transforms that describe their new shape and position. Third, the parameters of each homography transform are analyzed using multiple homographies to extract spacecraft motion (both rotation and translation) and the surface normal at each patch.

HSE is the fastest method available for slope estimation but it is still rather slow (0.3 Hz in the current implementation) and is unstable at the focus of expansion. Under ideal circumstances, it can effectively determine the slopes of selected patches. Unlike SFM, it does not generate a dense surface map, so it cannot be used to detect other hazards

HSE estimates the slopes of several patches. The most suitable landing site is selected by considering their slopes, local texture, the SSD (sum of squared differences) matching residual as well as spacecraft maneuvering capability. On the other hand, because the relative motion between two images is also determined, the spacecraft horizontal velocity can be also obtained. This scheme can be considered a next generation of the DIMES system.

5. Crater Hazard Detection (CHD)

CHD recognizes craters in an image. A crater, in general, is a bowl shaped depression created by collisions or volcanic activity. A typical crater in an image has an elliptical rim surrounding a bright-to-dark shading pattern dictated by the lighting angle and the crater's topography. These distinguishing characteristics are used extensively in crater detection.

One of the authors participated in the development of a crater detection algorithm [9, 10] for spacecraft optical navigation. The algorithm has been tested successfully for autonomous orbit determination. On imagery from MOC (Mars Overhead Camera), Odyssey, NEAR Eros, and others, its overall detection rate is better than 94% and false alarm rate is less than 5%. Most falsely detected craters are those topographic features that are similar to craters. The

algorithm consists of five stages [9]. A sample result is shown in Figure 4. First, Canny edge detection locates image edges that could correspond to crater rims. Second, edges that may belong to the same crater are grouped. Third, an ellipse is fit to each edge group. Fourth, ellipses are refined using the original, intensity image. Fifth, a confidence value is assigned to each crater.

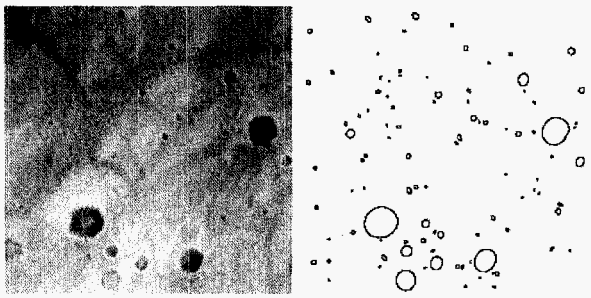

Figure 4. An crater detection example on a MER-A descent image.

The crater detection algorithm has also been shown to be an ideal solution for crater landmark recognition and matching, when a landmark map is available. It can be used for spacecraft localization either during orbiting or while descending [2, 9, 10, 20]. Here it is applied for crater hazard detection for safe landing. The detected craters not only provide the locations of crater hazards but also provide image photometric cues such as contrast between lit and shadowed areas as discussed earlier. These are useful to the rock detection and large discontinuity detection algorithms.

6. ROCK DETECTION (RD)

A number of rock detection techniques are currently under study at JPL to show the feasibility of rock detection based on shadows, on texture analyses, on elevation mapping in visible imagery, and also based on contrast between rocks and soils in thermal imagery. These algorithms can be applied at different altitudes in a cooperative manner. Here we focus on the use of shadows to detect the rocks. Shadows in visible imagery enable reliable rock detection over a range of solar incidence angles. Image texture analysis may enable detecting rocks for a broader range of sun angles if needed.

The rock detection algorithm has two phases: shadow detection and rock modeling. We have tested two algorithms for shadow segmentation and two algorithms for rock modeling. The height of the rocks is the critical measurement and can be determined form the length of the shadow region along the direction of illumination as a function of the sun incidence angle. The resulting rock map can then be analyzed for safe landing site selection.

6.1 Shadow Segmentation

The first step is to extract the shadow regions from the image. Many algorithms for region segmentation have been reported in the literature over the years. Some of these algorithms attempt to group pixels based on criteria that defines the level of uniformity of the pixels in the regions. Others use gray-scale morphology to segment specific regions from the background. At the other end, a number of thresholding techniques have been developed to segment the images into foreground and background classes. All these technique iterate in some form to achieve the result. We have studied several techniques and have selected two to implement and test: K-means clustering which iterates on the image intensities and Maximum Entropy Thresholding (MET) which iterates on the image histogram.

6.1.1 K-means

The K-means clustering algorithm [21, 22] is one of the simplest unsupervised learning algorithms that solve the well known clustering problem. It differs from the more general hierarchical clustering methods in that, the number of desired clusters, K, is given in advance. Unfortunately there is no general theoretical solution to find the optimal number of clusters for any given data set. A simple approach is to compare the results of multiple runs with different K classes and choose the best one according to a given criterion. The goal is to divide the image pixels into K clusters such that some metric relative to the centroids of the clusters is minimized. Various metrics to the centroids that can be minimized, for example, the maximum distance to its centroid for any pixel; the sum of the average distance to the centroids over all clusters, the sum of the variance over all clusters, and the total distance between all pixels and their centroids. The metric to minimize and the choice of a distance measure will determine the shape of the optimum clusters. K-means produces different clusterings depending on the initial cluster assignments. In our experiments so far we randomly assign each pixel to one of the K clusters. The position of the K centroids is determined as the value of the metric to minimize. In summary, the algorithm aims to minimize an objective function, in our case, a squared error function:

$$J = \sum_{j=1}^{k} \sum_{i=1}^{n} \left\| x_i^{(j)} - c_j \right\|^2 ,$$

where $\left\| x_i^{(j)} - c_j \right\|^2$ is a chosen distance measure between a pixel $x_i^{(j)}$ and the cluster center c_j, is an indicator of the distance of the n pixels from their respective cluster centers. Although it can be proven that the procedure will always terminate, the K-means algorithm does not necessarily find the most optimal configuration, corresponding to the global objective function minimum. The algorithm is also sensitive to the initial randomly selected cluster centers. Typically the

K-means algorithm is run multiple times to reduce this effect. The strong saliency of the shadows in our imagery however, makes multiple runs unnecessary.

6.1.2 gMET

Maximum Entropy Thresholding (MET) is an automatic thresholding technique based on the maximum entropy of the image histogram [23, 24]. MET maximizes the inter-class entropy. Entropy measures the uncertainty of an event and is defined as:

$$S = -\sum p_i \cdot \log_2(p_i),$$

where p is the probability of a pixel grayscale value, i, in the image. The *apriori* entropy of the image is then:

$$H_T = -\sum_{i=0}^{255} p_i \cdot \log_2(p_i)$$

Let b_0 represent the class of background pixels and b_1 the class of shadow pixels. Then we have that:

$$H_{b_0}(t) = -\sum_{i=0}^{t} \frac{p_i}{p(b_0)} \cdot \log_2 \frac{p_i}{p(b_0)} \text{ and}$$

$$H_{b_1}(t) = -\sum_{i=t+1}^{255} \frac{p_i}{p(b_1)} \cdot \log_2 \frac{p_i}{p(b_1)}, \text{ where}$$

$$p(b_0) = \sum_{i=0}^{t} p_i, \quad p(b_1) = \sum_{i=t}^{255} p_i \text{ with}$$

$$p(b_0) + p(b_1) = 1.$$

Now, let the information between the two classes be:

$$\Phi(t) = H_{b_0}(t) + H_{b_1}(t)$$

$$= \log[p(b_0)p(b_1)] + \frac{H(t)}{p(b_0)} + \frac{H_T - H(t)}{p(b_1)}, \text{ with}$$

$$H(t) = -\sum_{i=0}^{t} p_i \cdot \log_2(p_i).$$

The MET algorithm selects the threshold t^* at which $\Phi(t^*)$ is maximum and work well when the image has a bi-modal histogram, that is a single zone of high entropy. Since this condition is not likely to be always true for the scenes under consideration we apply the following transformation to the image intensities:

Let $\boldsymbol{I}$ be a graylevel image. The image $\boldsymbol{I_g}$ given by $I_g = I^\gamma + I$ now consists of an image where the shadow regions have been enhanced by gamma correction (hence the term gMET) and the background has been saturated. Such an image has a bi-modal histogram with a single high entropy region in the histogram. An example is illustrated in Figure 5. A small 320x360 portion of one of our large test images and its histogram are shown in Figure 5a and Figure 5b respectively. The modified image and its histogram are shown in Figure 5b and Figure 5d respectively.

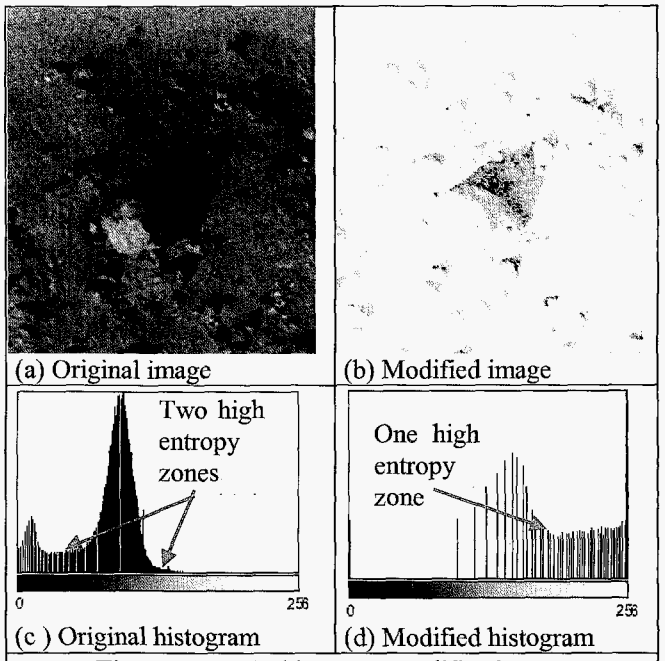

Figure 5. gMET histogram modification.

The shadows detected in our example in Figure 5 by 5-means and by gMET are shown in Figure 6a and Figure 6b respectively. In order to determine the appropriate number of clusters for K-means and the level of gamma correction for gMET we selected eight images from a data set (see Section 6.3.2 below) for eight different sun angles and registered them manually. We also constructed manually a ground truth from one of the images to determine rock detection and false alarm rates. Many small rocks are not present in the ground truth and the registration among the eight images is not perfect. The data however has been useful to conduct preliminary evaluations. We carried out rock detection experiments varying the number of clusters from 2 to 9 for K-means, and the level of gamma correction from 2.0 to 9.0 for all eight scenes and used the shadows

detected to model the rocks. The rock models derived from the shadows were compared against a ground truth (shown later in Section 6.3.2) to estimate detection and false alarm rates. The results were used to determine an appropriate number of clusters for K-means shadow segmentation, and the appropriate gamma value for gMET shadow segmentation. Figure 7 shows the ROC curves for all eight images of the same scene, with varying illumination angles, using 5-means for shadow extraction.

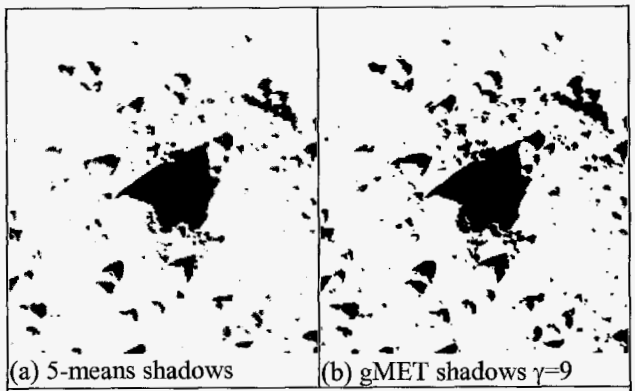

Figure 6. Analyses of the performance evaluation of the rock detection technique suggests using 5-means or 6-means and a gamma value between 4 and 9. These wide choices illustrate the non-criticality of these parameters.

The area under the ROC curves for all scenes and cluster combinations is shown in Figure 8. It suggests using five clusters performs best for this dataset. The areas under the ROC curves for the gMET method suggest gamma corrections of 6.0 to 9.0 for high contrast images. The selected parameters correspond to the Mars Hill data set described below in more detail in Section 6.3.2. The algorithm behaviors indicate that the sensitivity to the parameters is not crucial, and we have applied these settings to scenes in two other datasets.

None of the tested datasets however include imagery corrupted by noise of having dynamic range and contrast degradations due to sand storms or other atmospheric phenomena. We have manually produced images having such contrast variations in order to determine the compensation needed for K-means and gMET parameters. Figure 9 illustrates a moderate, a severe, and a very severe contrast alteration of the image shown earlier in Figure 5, and the shadow regions extracted from them by our two algorithms with the adjustments needed to compensate. Note that contrast information is available early through the first two stages of decent described in Section 3 above.

6.2 Rock Modeling

The critical measurements for EDL are the rock heights. Note that the aim is not to model the rocks in terms of a perfect and accurate delineation of the rock volume nor even an accurate delineation of its largest 2-D horizontal cross-section. The goal is to derive a very good estimation of its height together with a good approximation of its location, and a reasonable approximation of its 2-D horizontal cross-section.

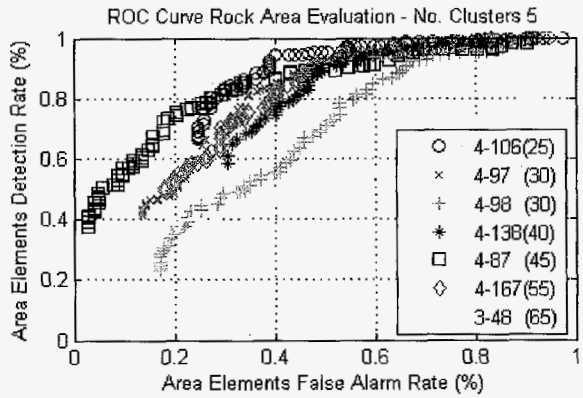

Figure 7. ROC curves for experiments using five clusters. The legend designates the image labels and the sun elevation angles in parentheses.

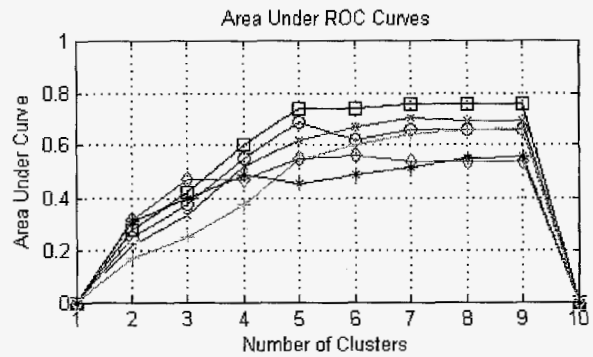

Figure 8. The area under the ROC curves illustrate a steady performance staring at 5 clusters.

6.2.1 Shadow Analysis

Figure 10 illustrates the rock-shadow measurements that can be made and inferred from the shadow of a rock. The farthest point from the rock laying on the shadow boundary along the direction of illumination is cast by the highest point on the rock. These corresponding pair (blue dots in Figure 10) gives the true shadow length. The shadow boundary corresponds to the projection of the shadow casting boundary, the tangential profile on the surface of the rock. The highest point on the rock lies on this boundary (between the pink dots.) Part of the shadow may lie on the rock itself, part on the adjacent terrain and part on other rocks.

The width of a rock is given by the sum of the distances (w_1 and w_2 in Figure 10) from the sun ray passing through the highest point and the two farthest points in the shadow boundary in the direction orthogonal to the sun ray. The region corresponding to the self-shadow is usually not measurable (green text) as the contrast of the rock boundary

in the shadow is likely to be low. The actual center (of mass) of the rock is also not measurable directly. Our algorithm measures only the length of the shadow, and the width of the shadow for large shadow regions that may correspond to very large rocks or boulders.

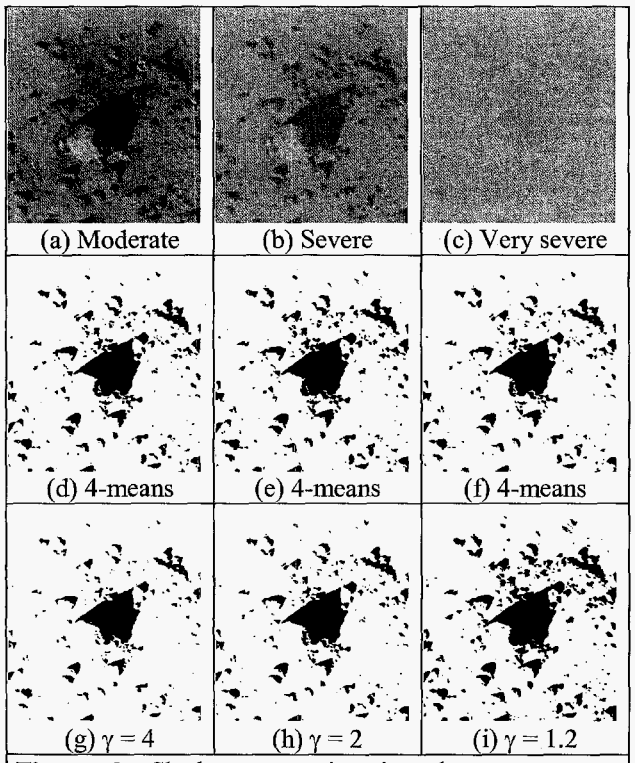

Figure 9. Shadow extraction is robust to contrast variation. Shadow contrast can be estimated during descent form higher altitude images.

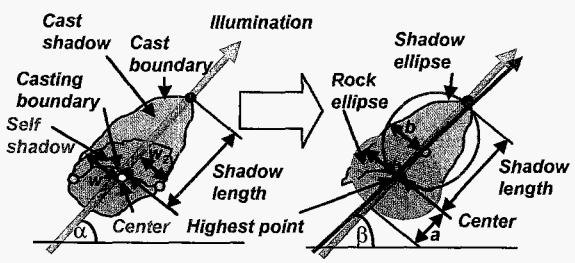

Figure 10. A rock, its shadow and their models

To improve performance, for most rocks we approximate the shadow region by a best-fitting ellipse (Figure 10, right), and use its parameters to derive the rock model. A best-fitting ellipse [25] equates the second order central moments of the ellipse to those of the distribution of the pixels in the shadow region, and thereby effectively defines both the shape and size of the ellipse. The projection of the appropriate ellipse axis onto the sun ray passing through its center gives the shadow length. This may over- or under estimate shadow lengths of small rocks, which is not crucial. For large rocks however we are comparing this elliptical approximation with the shadow length actually measured from analysis of the shadow region itself. The shadow width (and hence the rock width) is given by the length of the other ellipse axis. In a detailed analysis, the center of the rock, and hence its position, can be estimated by offsets from the highest shadow casting point. These offsets currently are a function of the distance between the points where the w_1 and w_2 measurements are made. In the simpler model the center of the rock is given by the extreme of the shadow ellipse (Figure 10, right). The rock itself is modeled by an ellipse with one axis accurately determined by the shadow width, and the other estimated by a simpler heuristic, a fraction of the rock width. The length pf the rock is not measurable directly and is estimated by a heuristic that is a function of the rock width. In simpler terms, the rock cross-section can be made to correspond to a circle having a diameter equivalent to the rock width. The rock height is given by the shadow length as a function of the sun incidence angle.

In order to evaluate these measurements we are looking at the rock size-frequency distribution models that have been developed at JPL by Golombeck [25] and his colleagues. These include rock abundance and size distributions for Mars missions going back to the Viking missions. For precise evaluation of height estimates we are constructing datasets and ground truth references at our Mars Yard facility. Height is the critical measurement and is affected by the accuracy at which we can measure shadow length. In general the length of the shadow is affected by light diffraction as a function of rock height and the roughness of the shadow casting profile. Secondary error effects are due to the line-spread function of the optics which blurs intensity boundaries.

6.3 Experiments and Results

We have tested our algorithms successfully with descent imagery from the near landing on the EROS asteroid, aerial imagery of "Mars Hill" in Death Valley, California, and test imagery from our Gantry facility. Next we present a representative sample.

6.3.1. EROS Results

A sample 492x392 image from the NEAR descent image of the EROS asteroid is shown in Figure 11. The ground resolution is about 0.024 mpp. The sun incidence angle is 45°. Rocks>400 pixels are labeled hazards in red color. The shadow regions extracted by 5-means and gMET, shown in Figures 11b and 11c, leading to similar rock models, as shown in Figures 11d and 11e. The K-means algorithm is well established and gMET is new but much faster. We typically run both in all our tests as real-time considerations are most important.

6.3.2. Mars Hills Experiments

Aerial images of Mars Hill were acquired in 1989 to evaluate landing hazard detection algorithms. The site is useful because it has minimal vegetation and it has a rock

distribution similar to that seen at the Viking 2 Mars landing site. We selected and registered eight images of the site for our tests and evaluations. Figure 12 shows one of such images. Processing these images at ¼ resolution yields the typical results shown in Figure 13. Small 320x360 windows from this set were used in evaluation experiments. The window from the image in Figure 12 is shown in Figure 14a. The ground truth reference constructed by hand from this image is shown in Figure 14b.

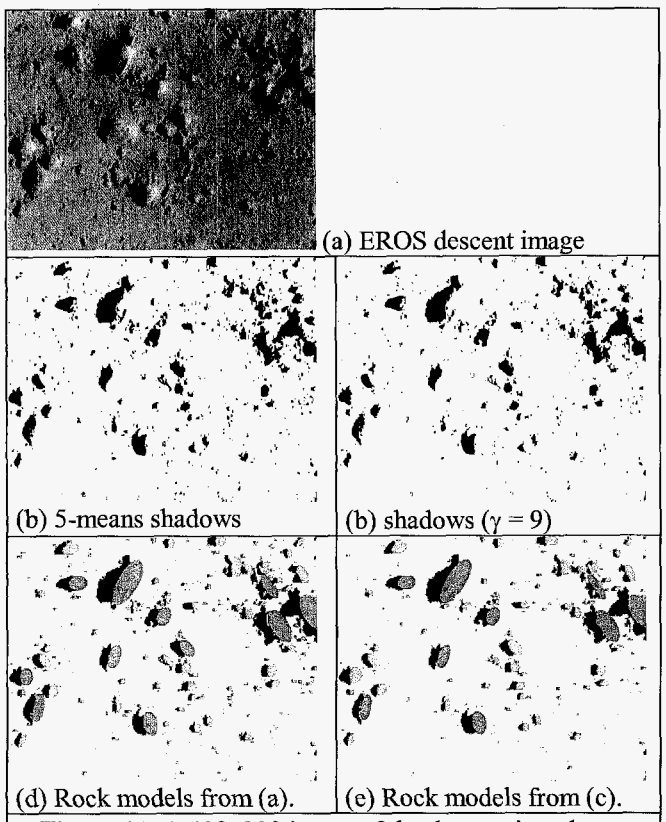

Figure 11. A 492x392 image of the descent into the EROS asteroid. Potential hazards are shown in red.

In our experiments for algorithm evaluation and rock detection performance evaluations we processed windows from eight images having different illumination and sun incidence angles. These help determine the effects of shadow blending and shadow region sizes as a function of the sun angles and the subsequent effects on the rock models generated. In general, shadow blending is not a major problem for sun incidence angles > 25° and < 75°. According to previous studies [26] the expected rock size-distribution above 1.5 m follows a steep decrease with increasing diameter. In all our experiments we ran K-means and gMET using eight combinations of clusters and gamma corrections. In all cases the darkest cluster from K-means corresponded to the shadow regions. For gMET the dark class always corresponded to the shadow regions. Figure 15 shows representative results from our experiments. In Figure 15a the large rock is aligned with the small rock along the direction of illumination. If the shadows are blended, the height of the boulder may be overestimated. If they are not, the height of the boulder may be underestimated. In Figure 15b the shadow of the boulder falls on smaller rocks with no consequences. In Figure 15c, the long shadows from a high sun incidence angle blend the boulder's shadow with that of the rocks in the shadow and with shadows from nearby rocks. Additional analysis of the shapes of the (large) shadows may be necessary to break the shadow regions. In the case illustrated the rock model is oversized in 2D but its height is accurate. In Figure 15d the illuminated surface of dark rocks may blend with their self-shadows resulting in overestimated rock heights. This condition can be handled by examining the shape of the shadow casting profile in more detail. Precise ground truth under construction will help quantify these effects better.

Figure 12. 1600x1600 Rock field image from a set of eight. The large boulder on the top spans 80 pixels and a small 320x360 window used in our examples below.

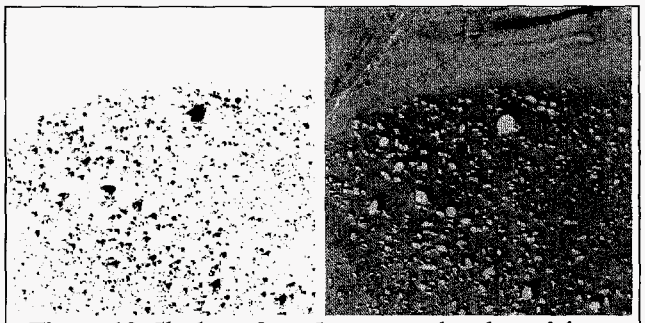

Figure 13. Shadows from 5-means and rock models

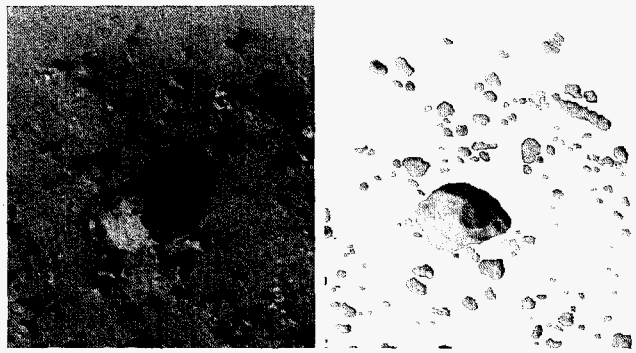

Figure 14. Portion of image in Figure 12.

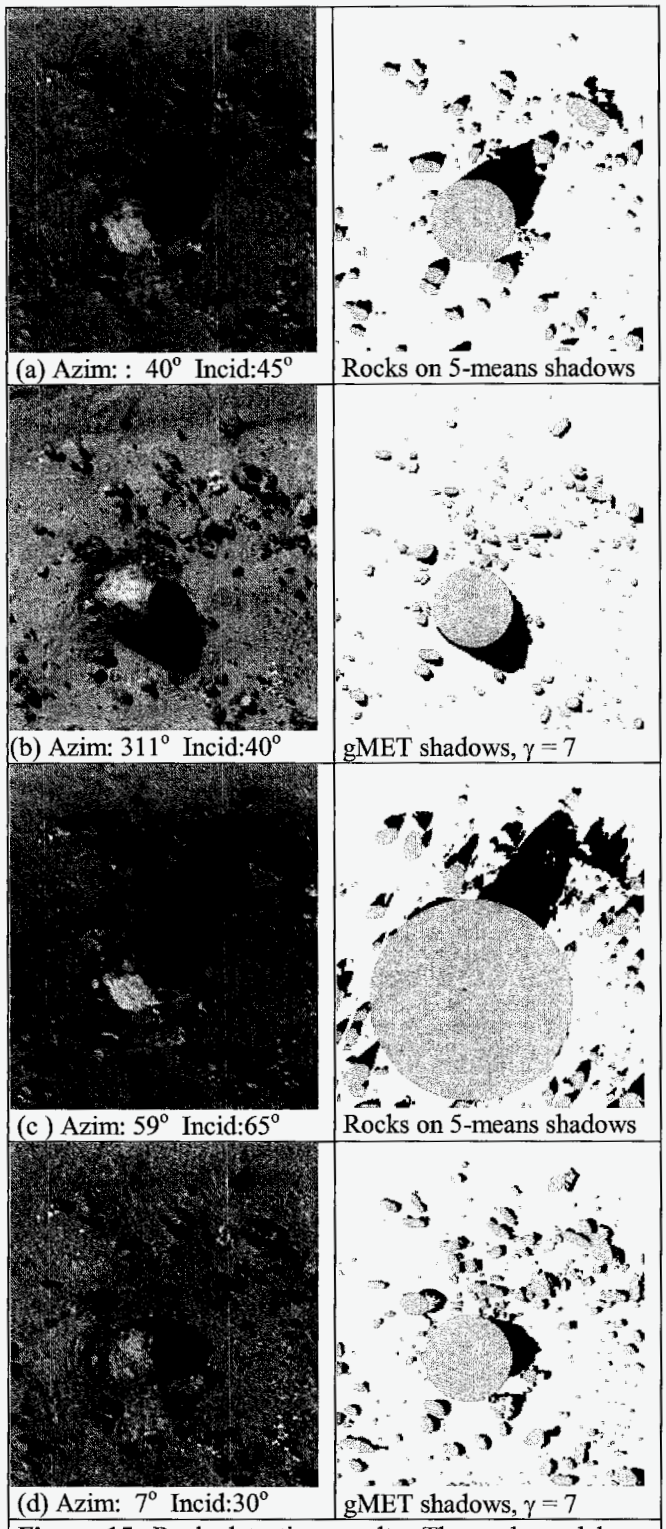

Figure 15. Rock detection results. The rock models are superimposed on the shadow regions detected

6.3.3 Gantry Experiments

The Gantry data set consists of eight 1024x768 images of a 1.0x0.75 m scene illuminated artificially to simulate sun incidence angles from 20° to 80°. We also acquired laser scan data at 0.05° resolution for an initial evaluation of measured rock heights. The laser data turned noisy but useful to start characterizing height measurements accuracy. Other datasets are being currently constructed outdoors for this purpose. Figure 16a shows one of the images. The scene contains 47 small rocks (a few centimeters in diameter and height) having a variety of shapes, appearance and contrast. The ground truth was constructed by delineating the rock regions by hand and registering these with the laser scan data.

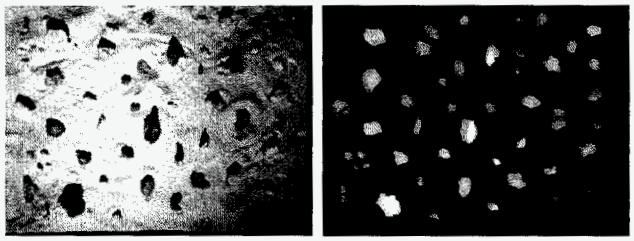

(a) Gantry image at 1/2 resolution (2 cm/pixel) artificially illuminated with a 40° incidence angle.

(b) Ground truth. Pixel values encode height in mm. Higher points are brighter.

Figure 16. Forty-seven small rocks of varying shape, height, appearance, and contrast on sandy disturbed terrain.

Figure 17a and Figure 17c illustrate 100% detection of the rocks using 5-means and gMET shadow detection respectively. In Figure 17b and Figure 17c we show overlays of the rock models and the ground truth. Note that only the tree buried rocks on the top left quadrant of the image are severely underestimated in size due the very small shadows they cast. The remaining of the rock are detected and positioned well. The red models correspond to small rocks and other marker objects not present in the ground truth. These typically will be filtered out on size.

We are developing several performance evaluation methods to study rock detection performance in terms of global cumulative area element detection and false alarm rates, and in terms of individual rock detection and false alarm rates. Figure 18 illustrates some of these measurements for the result shown in Figure 17a. The plots show an example of absolute differences in height, width, and area for each rock.

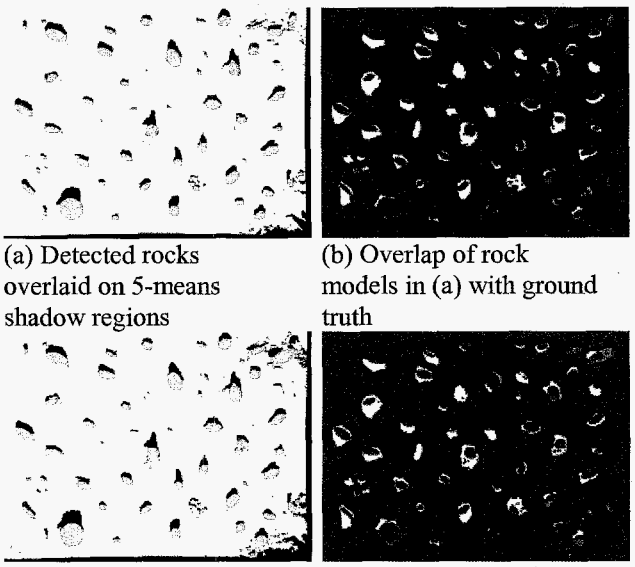

(a) Detected rocks overlaid on 5-means shadow regions

(b) Overlap of rock models in (a) with ground truth

(c) Detected rocks overlaid on gMET shadow regions

(d)) Overlap of rock models in (c) with ground truth

Figure 17. In spite of the difficult scene, 100% of the rocks in the ground truth are detected. The footprint of the detected models illustrate good positioning.

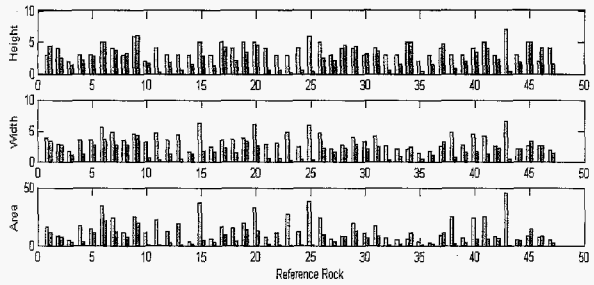

Figure 18. Rock height is the critical estimate. Larger rocks produce very good estimates. The ground truth is shown in cyan bars and the measured estimates are shown in magenta. The average error in height for his example is 1.2 cm and 0.8 cm for the K-means and gMET shadows respectively.

7. Texture Analysis (TA)

We do not expect that a single natural hazard detection algorithm perform perfectly in terms of detection and false alarm rates. More than one algorithm however can cooperate to deliver very high detections. Our system uses in addition a TA algorithm to provide a protective layer between the hazard detection and final safe landing site selection. TA identifies smooth areas in the image, which are likely to have fewer rocks (see Section 8 below) or other discontinuities, and can therefore increase the confidence on safer landing sites. Texture analysis can take many forms, such as thresholding on intensity variance, on spatial frequency coefficients, or maximum "blob size" over windows in an image. A simple texture analysis such as windowed image variance can efficiently find the smoothest area in an image – useful for avoiding rocks – but it cannot necessarily detect other hazards such as craters and slopes. A significant advantage of TA is that it can be implemented to run very fast thus affording multiple runs if necessary.

8. Landing Site Selection

The selection of a landing site is a continuous process. When the spacecraft is at a high attitude, between 10 and 2 km, the image resolution is not sufficient do detect small rocks. Large hazards such as craters and terrain discontinuities however can be detected at such altitude and a few very large terrain patches that are crater and discontinuity free are selected. The patch selection criteria are based on low texture and a safe distance from large hazards. As additional descent imagery becomes available, these patches should be located in overlapping portions of the images to enable surface slope estimation. Figure 19 illustrates an example from our Gantry dataset.

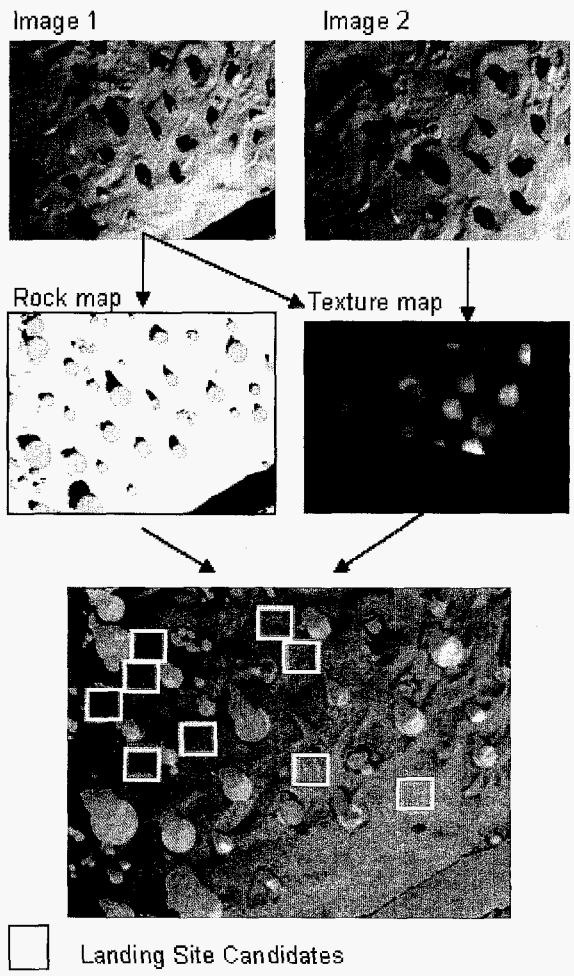

Figure 19. The landing site is selected by considering the detected rocks, local texture, overlap between the two descent images. When the surface slopes are estimated, the final landing site can be selected from many patches.

9. REAL-TIME CONSIDERATIONS

Earlier tests of shadow detection using k-means clustering runs in 0.22s on 768x764 pixel imagery on a 2.2 GHz Pentium 4 PC. We expect that this can be sped up by accelerating k-means by using kd-tree data structures [27]. The kd-tree data structure is used to reduce the large number of nearest-neighbor queries issued by the traditional algorithm. Sufficient statistics are stored in the nodes of the kd-tree and then an analysis of the geometry of the current cluster centers results in great reduction of the work needed to update the centers. The technique described in [27] claims exact behavior as the traditional algorithm and suggest methods to initialize the K-means starting centers efficiently. Potential speedups for MET type algorithms are described in [28] however we anticipate that the current algorithm can be implemented to run efficiently.

At the MC-HAD system level hazards must be detected and characterized within the few seconds allowed for EDL. A number of factors however are taken into consideration. First, we take advantage of the much higher resolution and wider field of view than that available to from other sensors thus starting the image analysis early even if the actual landing site is not in view. Second, we have available considerable information on Mars geomorphology to provide *apriori* information that can be used to refine parameters and assist future missions there. We also have the rock abundance and size-frequency distribution models [26], which are helpful to fine tune rock detection algorithm strategies. On the computing side, processed and labeled portions of the image can be skipped by subsequent processes. On the strategy side, we allow algorithms to cooperate and reinforce each others strengths.

10. FUTURE WORK

Plans are underway to construct additional data sets at our Mars Yard facility and the appropriate "Ground truth" rock distributions to quantitatively assess performance. We would like to be able to detect rocks 5 pixels across representing 1 m diameter rocks at 1 km altitude. This is consistent with conclusions of the MRSR study [29]. Since it is clear that large rocks are detected reliably, we can use rock diameter thresholds of 10 pixels in the full resolution image to approximate ground truth, then reprocess the images at half resolution to quantify performance against the results from the full resolution image. We also expect that the system will be fast enough to enable multiple looks, which will increase the overall detection probability.

ACKNOWLEDGEMENT

The research described in this paper was carried out at the Jet Propulsion Laboratory, California Institute of Technology, funded by Mars Program under NASA Mars Exploration Program Advanced Technologies NRA *OSS-03-01*.

REFERENCES

[1] L. K. Tamppari, "Dust in the Mars Atmosphere", Dec. 10, 2002.

[2] Y. Cheng and A. Ansar, "Landmark based position estimation for pinpoint landing on Mars", IEEE Conf. on Robotics and Automation, 2005, pp. 4481-4486.

[3] A. Johnson and L. Matthies, "Precise Image-Based Motion Estimation for Autonomous Small Body Exploration,:" Proc, 5th Int'l Symph. Artificial Intelligence, Robotics and Automation in Space (iSAIRAS '99), June 1999, pp 627-634.

[4] Y. Cheng, A. Johnson, and L. Matthies, "MER-DIMES: A Planetary Landing Application of Computer Vision," IEEE Conference on Computer Vision and Pattern Recognition (CVPR 2005), June 2005.

[5] Y. Cheng, J. Goguen, A. Johnson, C. Leger, L. Matthies, M. San Martin, and R. Willson, "The Mars Exploration Rovers Descent Image Motion Estimation System," IEEE Intelligent Systems 19(3), pp. 13-21, May/June 2004.

[6] O. Amidi, T. Kanade, and K. Fujita, "A visual odometer for autonomous helicopter flight", Robotics and Autonomous systems, 28:185-193, 1999.

[7] P. Garcia-Pardo, G. Sukhatme, and J. Montgomery, "Towards Vision-Based Safe Landing for an Autonomous Helicopter." Robotics and Autonomous Systems, Vol. 38, No. 1, pp. 19-29, 2001.

[8] Y. Cheng, A. Johnson, L. Matthies, and A. Wolf, "Passive Imaging Based Hazard Avoidance for Spacecraft Safe Landing", I-SAIRAS, 2001.

[9] Y. Cheng, A. Johnson, C. Olson and L. Matthes, "Optical Landmark Detection for Spacecraft Navigation" 13th Annual AAS/AIAA Space Flight Mechanics Meeting, 2003.

[10] Y. Cheng and J. Miller "Autonomous landmark Based Spacecraft Navigation", 13th Annual AAS/AIAA Space Flight Mechanics Meeting, 2003.

[11] R. Zhang, P. Tsai, J. Cryer and M. Shah, "Shape from Shading: A Survey," IEEE Trans. PAMI, Vol. 21, No. 8, 690-706 Aug 1999.

[12] P. Christensen, The Martian surface layer" *Mars*, University of Arizona Press, 1992.

[13] P. Christensen, "Morphology and composition of the surface of Mars: Mars Odyssey THEMIS results", *Science*, Vol 300, pp 2056-2061.

[14] A. Huertas and R. Nevatia, "Building Dtection in Aerial Images using Shape and Shadows," Proc. 8th Int. Joint Con. On Artificial Intelligence, Karsruhe, Germany, 1983.

[15] A. Huertas, A. Lin and R. Nevatia, "Detecton of Buildings from Monocular Images," In Automatic Extraction of Man-Made Objects from Aerial and Space Images," Ed. A. Gruen, O. Kuebler, P. Agouris, Vrinhauser Verlag, Switzerland 1995.

[16] D. Bernard and M. Golombek, "Crater and Rock Hazard Modeling for Mars Landing", Proc. AIAA Space 2001 Conference, Albuquerque NM, Aug 2001.

[17] A. Howard and H. Seraji, "A Rule-Based Fuzzy Safety Index for Landing Site Risk Assessment," 9th International Symposium on Robotics and Applications, Orlando, FL., June 2002.

[18] A. Howard, "A Novel Information Fusion Methodology for Intelligent Terrain Analysis," Proc. IEEE Int. Conf on Fuzzy Systems, Honolulu, HI, May 2002.

[19] P. Gill, W. Murray, and M. Wright, "The Levenberg-Marquardt Method." §4.7.3 in Practical Optimization, London Academic Press, pp. 136-137, 1981.

[20] A. Ansar and Y. Cheng, "An Analysis of Spacecraft Localization from Descent Image Data for Pinpoint Landing on Mars and other Cratered Bodies" Journal of Photogrammetric Engineering and Remote Sensing, Special Issue on Mapping Mars, Accepted for publication.

[21] J. MacQueen, "Some Methods for classification and Analysis of Multivariate Observations," Proceedings of 5-th Berkeley Symposium on Mathematical Statistics and Probability", Berkeley, University of California Press, 1:281-297.

[22] A. Jain and R. Dubes, *Algorithms for Clustering Data*, Prentice Hall, 1988.

[23] P. Sahoo, S. Soltani, K. Wong and, Y. Chen, "A Survey of Thresholding Techniques," *Computer Vision, Graphics, and Image Processing*, Vol. 41, pp.233-260, 1988.

[24] J. Kapur, P. Sahoo and A. Wong, "A New Method for Gray-Level Picture Thresholding Using the Entropy of the Histogram," CVGIP(29), No. 3, March 1985, pp. 273-285.

[25] H. Cramer, *Mathematical Methods of Statistics*, Princeton Landmarks in Mathematics and Physics, Princeton University Press. Princeton, NJ, 1999.

[26] M. Golombeck, A. Haldemann, N. Forsberg-Taylor, E. DiMaggio, R. Schroeder, B. Jakosky, and R. Matijevic,"Rock-size frequency Distributions on Mars and Implications for Mars Exploration Rover landing Safety and Operations," Jpurnal of Geophysical Research, Vol. 108, No. E12, 2003, pp. 27-1:27-23.

[27] D. Pelleg and A. Moore, "Accelerating Exact k-means Algorithms with Geometric Reasoning," Proc. Conference on Knowledge Discovery in Databases, 1999.

[28] M. Fleury, L. Hayat, and A. Clark, "Parallel entropic auto- thresholding," Image and Vision Computing, Vol. 14, No. 4, 1996, pp. 247-263.

[29] R. Burke (ed.), "Report of the Mini-Workshop on orbital Imaging Requirements for the Mars Rovers Sample Return Mission," Jet Propulsion Laboratory Technical Report No. D-6442, May 15.1989.

BIOGRAPHY

Andres Huertas *is a Senior Member of the Technical Staff at the Jet Propulsion Laboratory. He earned his B.S. degree in Systems and Computing Engineering from Universidad de los Andes in 1974, and his M.S. and Engineer degrees in Computer Science in 1978 and in Electrical Engineering in 1984 from the University of Southern California. Prior to joining the Computer Vision Group at JPL in 2002, he worked on computer vision research since 1978 at the Image Processing Institute and at the Institute for Robotics and Intelligent Systems at the University of Southern California. At JPL, he works on computer vision techniques for autonomous off-road navigation, and on techniques for hazard detection for safe landing of spacecraft on planetary surfaces.*

Dr. Yang Cheng *is a Senior Member of Technical Staff at the Jet Propulsion Laboratory. He earned the Ph.D. in Geographic Information Processing at the University of South Carolina, where he studied remote sensing and parallel computation. In 1993 he moved to the Oak Ridge National Laboratory, where he studied remote sensing and large scale geographic information processing. In 1999, he came to Jet Propulsion Laboratory.*

His expertise covers areas of robotic vision, cartography, map projection, photogrammetry, remote sensing and parallel computing. He was a main developer of the Descent Image Motion Estimation System (DIMES) and Visual Odometry for the NASA Mars Exploration Rovers (MER) mission. He currently works on vision algorithm development for small body navigation, spacecraft safe and pin point landing and others.

Dr. Richard Madison *is a Member of the Technical Staff in the Computer Vision Group at the Jet Propulsion Laboratory. He earned his B.S. degree in Engineering at the Harvey Mudd College in 1992, and his M.S. and PhD degrees in Electrical and Computer Engineering in 1994 and 1997 from Carnegie-Mellon University. He works on computer vision related infrastructure tasks for the Mobility and Robotics Systems Section at JPL.*