A New Strategy to Land Precisely on the Northern Plains of Mars

Yang Cheng*, Andres Huertas*

*Computer Vision Group, Jet Propulsion Laboratory,
California Institute of Technology, USA
e-mail: {ycheng, Huertas}@jpl.nasa.gov

Abstract

During the Phoenix mission landing site selection process, the Mars Reconnaissance Orbiter (MRO) High Resolution Imaging Science Experiment (HiRISE) images revealed widely spread and dense rock fields in the northern plains. Automatic rock mapping and subsequent statistical analyses showed 30–90% CFA (cumulative fractional area) covered by rocks larger than 1 meter in dense rock fields around craters. Less dense rock fields had 5–30% rock coverage in terrain away from craters. Detectable meter-scale boulders were found nearly everywhere. These rocks present a risk to spacecraft safety during landing. However, they are the most salient topographic features in this region, and can be good landmarks for spacecraft localization during landing. In this paper we present a novel strategy that uses abundance of rocks in northern plains for spacecraft localization. The paper discusses this approach in three sections: a rock-based landmark terrain relative navigation (TRN) algorithm; the TRN algorithm feasibility; and conclusions.

1 Introduction

One critical capability of future solar system in situ exploration is the so-called pin-point landing (PPL), which is defined as the capability of landing a spacecraft within 100 meters of a targeted site. PPL provides safe and affordable access to high scientific targets, allows the highest science returns, and reduces risk to the spacecraft. Given the lack of a global positioning system (GPS) on other planetary bodies, PPL relies solely on computer vision based terrain relative navigation (TRN) technology. This technology recognizes the local terrain and locates the spacecraft within the local terrain frame by comparing a descent image with an on-board reference map.

A TRN system works as follows. From Earth, a landing site on the targeted body is selected using orbital imagery, and the appropriate landmarks (e.g. craters, rocks) within the landing ellipse are mapped. During descent, the lander's initial position with respect to the selected landing site is determined automatically on board in real-time, using the mapped landmarks. The lander is then guided to the landing site using continuous updates of lander position and velocity throughout the descent. There are three fundamental requirements for this new technology:

1. The system must be able to recognize the terrain (landmarks) reliably and repeatedly over a wide variety of environments. During spacecraft descent, considerable environmental variations, such as lighting angles, atmospheric conditions, viewing angle, and spacecraft altitude, can affect the appearance of landmarks. The key to success relies on defining a class of landmarks that has good invariance properties under variable environmental conditions.
2. The system must accomplish the task under extreme limitations imposed by the slow flight computer and the terminal descent time constraints. The projected CPU clock speed of the flight computer for a PPL mission will be one order of magnitude slower than a regular desktop. Generally, there is only a 60 second window of opportunity for spacecraft localization during the parachute stage, between heat shield jettison and powered descent. In order to obtain reliable spacecraft state information, several images as well as other on-board sensor data such as that from an IMU and altimeter must be processed during this period of time. In general, each cycle of spacecraft localization, which includes image acquisition, image processing and sensor data infusion, should be carried out in a few seconds. This is a very demanding requirement.
3. The system must be able to guide the spacecraft to land within 100 meters of the target. Many deficiencies and/or limitations of both hardware and software could degrade the system's performance. Landmark detection errors, base map errors (both position and elevation), sensor noise (imager, IMU, altimeter), and image and IMU sensor misalignment, etc., introduce unwanted uncertainties. By considering all these uncertainties, however, the system can be robust enough to provide a valid spacecraft state that meets or exceeds the PPL requirement.

A few TRN algorithms for PPL have been proposed to date. In [1][3], a correlation based TRN algorithm –Map and Image Alignment (MAIA) is suggested. In this approach, the base map is a piece of rectified orbital image. A correlation method is used to recognize the descent image within a base map. A high pass filter is applied to both data to increase robustness. Because this is a correlation method, the descent image must be rectified to the orientation and scale of the base map.

Thus, the method depends on the availability of an estimate of spacecraft altitude and attitude. Another drawback with this approach is that it is sensitive to lighting condition changes. In particular, when the sun angles differ greatly from those represented in the base map the correlation matching becomes less reliable.

Unlike the MAIA algorithm, the OBIRON algorithm does not depend on a complete map of the landing region. Rather, it carries a set of landmarks. Each landmark comprises a small patch surface elevation grid and an albedo map [7]. Instead of matching the descent image to a base map directly, OBIRON matches the descent image with a predicted image patch, which is rendered according to an *a priori* known spacecraft state, and descent camera model. A similar approach was also suggested in[2]. Like the MAIA, OBIRON algorithms it also needs very good *a priori* spacecraft state in order to render a terrain patch correctly. Both OBIRON and MAIA are currently too slow for real mission application.

Using natural landmarks, such as craters, is another option. A crater is very salient under different sun angles and viewing angles. A crater PPL TRN approach was suggested in [4][5]. This approach contains two algorithms: crater detection and crater matching. The crater detection algorithm uses image cues such as the rim edge, shape, and image intensity profile inside the crater to detect and fit crater the rim with an ellipse. The set of detected craters are then matched to corresponding craters stored in a database using crater position, size, orientation, and context information. This approach is much faster, less sensitive to lighting and viewing variations and, most importantly, it does not relay on *a priori* knowledge of spacecraft state. As a result, it is applicable at a very high altitude, well before the onboard altimeter locks on the surface. The principal drawback of this approach is that it works well only for scenes where craters are abundant.

During the Phoenix mission landing site selection process, the Mars Reconnaissance Orbiter (MRO) High Resolution Imaging Science Experiment (HiRISE) imagery revealed widely spread and dense rock fields in the northern plains of Mars (Figure 1). Automatic rock mapping and subsequent statistical analyses showed 30–90% CFA (cumulative fractional area) covered by rocks larger than 1 meter in diameter in dense rock fields around craters. Less dense rock fields had a 10–30% rock coverage, and 0–10% rock coverage in background terrain away from craters. Detectable meter-scale boulders almost always could be found nearly everywhere. In an area with 10% rock CFA, for example, there will be roughly three 1.5m diameter (or larger) rocks in a hectare ($100\times100m^2$) region Even though these rocks present a certain amount of risk to spacecraft safety during EDL, they are the most salient topographic features in this region, and represent the best landmarks for spacecraft localization.

In this paper we present a novel localization strategy using the abundance of rocks in the northern plains. The paper will discuss this approach in three sections:

- Rock-based landmark TRN algorithm
- Rock landmark TRN feasibility
- Conclusions

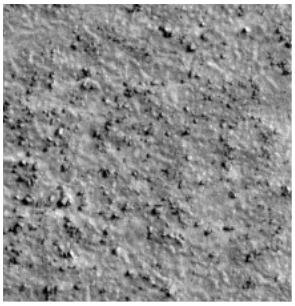

Figure 1: A small window of a HiRISE image shows a large number of rocks and boulders in this region. The image resolution is about 0.3 m/pixel.

2 Rock Landmark TRN Algorithm

2.1 Shadow Based Rock Detection

The shadow-based rock detection algorithm has been described in detail in [6][8], and we only summarize it here. It consists of four steps, illustrated in Figure2:

1) Image acquisition and state instantiation. Here, state includes environmental state (sun angle, shadow contrast) and a spacecraft state (altitude, pose and ground resolution).

2) Shadow Segmentation. Shadow regions are segmented by a modified Maximum Entropy Thresholding (gMET) algorithm[8]. In this method, the shadows are segmented by analysis of the histogram of a modified input image. The modified image is computed by adding a gamma-corrected image to saturate bright areas (non-shadows) at the image bit-depth.

3) Shadow Analysis. The aim here is to fit a "best-ellipse" as described in [8] to the shadow regions that are larger than a certain size (typically 5 pixels).

4) Rock Modeling. The parameters of the shadow elliptical approximation are combined with the Sun angle information to estimate shadow length and width, and rock height and location.

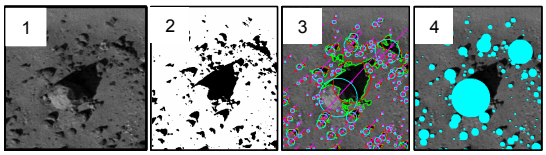

Figure 2: Shadow-based rock detection. 1) The input image; 2) Shadow segmentation; 3) Shadow analysis; 4) Rock model (position, footprint).

The rock detection algorithm and the measurements and statistics from these direct measurements were initially developed for real-time hazard detection during EDL. Later, it was applied to make thematic maps of candidate landing sites for the Phoenix mission. A large number of HiRISE images covering more than 1,500 km^2 of Mars northern plain were processed, and more than 10 million rocks larger than 1m in diameter were detected and mapped. Automatic rock detection results were scrutinized and compared to sample manual count results and to surface counts from images of previous lander missions. Eventually, the automatic counts were adopted by the Phoenix mission and the full set of automatic mapping results together with the statistics and maps derived from them were used in the landing site selection and certification task [8]. Figure 3 shows a set of rock detection results and an example of the thematic maps and statistics. The HiRISE image is PSP-001391-246; its size is 20048×10000. A total of 74,146 rocks were detected in this image. The low right figure illustrates the rock density map, which plots the number of rocks 1 meter in diameter of larger per hectare.

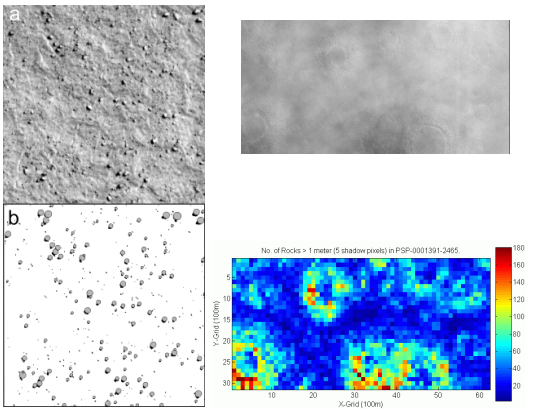

Figure 3: The modified rock detection was used successfully in mapping rocks/boulders from HiRISE image for Phoenix mission landing site selection.

2.2 Rock Landmark Formation

Logically, those algorithms used for star matching are natural choices for rock matching. However, the rock landmark matching for PPL has is a different problem domain. Unlike star matching, which aims to determine the s/c attitude, the rock matching TRN at here focuses on determining the s/c position, since the s/c attitude is estimated in very high precision by onboard star tracker and IMU. Thus, a different matching solution is suggested here.

For use in the Rock Landmark TRN algorithm, each rock is described by its location, an estimated diameter (the width of the shadow) and an estimated height (from shadow length and sun incident illumination angle). A rock landmark is defined as a constellation of a center rock and four neighboring rocks. The neighborhood zone is a circular region centered on the center rock. It is assigned a radius $\boldsymbol{r}$, proportional to the estimated height:

$$r = ch \tag{1}$$

where $\boldsymbol{c}$ is a constant and h is the estimated rock height.

Under this scheme, larger rocks are assigned larger neighborhoods. The estimated height was chosen as the measure of rock size because it is relatively insensitive to image resolution and sun angles.

The circular neighboring zone is divided into four quadrants (Figure 5). In each quadrant, the tallest rock is selected as the neighbor in this quadrant. Therefore, for each rock, there is a maximum of four neighbors, called a rock *quartet*. In total, five rocks form a rock landmark. The four clockwise bearing angles $Q(q_1, q_2, q_3, q_4)$ in each quarter, the estimated rock heights $H(h_1, h_2, h_3, h_4)$, and the distances to each neighbor $D(d_1, d_2, d_3, d_4)$ comprise the landmark descriptor $L(Q, H, D)$.

The rock landmark database over a landing site is built prior to a mission using orbital imagery. This database contains the rock landmark descriptors, their 3D position and few hash tables to facilitate fast searching.

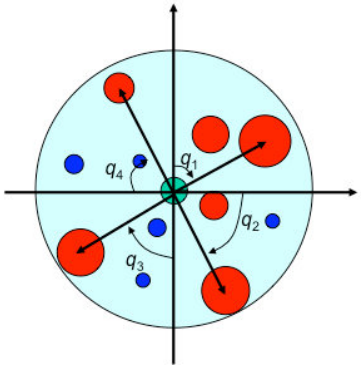

Figure 4: A rock landmark is defined as a set of four rocks centered at the green rock. The tallest neighbor in each quadrant is used to form the landmark.

2.3 Rock Landmark Matching Strategy

Ideally, any landmark L, in the database, should be related to its counterpart $L'(Q', H', D')$ by

$$\begin{pmatrix} q_1 \approx q_1' & q_2 \approx q_2' & q_3 \approx q_3' & q_4 \approx q_4' \end{pmatrix} \\ \left(\frac{h_1}{h_1'} \approx \frac{h_2}{h_2'} \approx \frac{h_3}{h_3'} \approx \frac{h_4}{h_4'} \right) \\ \left(\frac{d_1}{d_1'} \approx \frac{d_2}{d_2'} \approx \frac{d_3}{d_3'} \approx \frac{d_4}{d_4'} \right) \tag{2}$$

Because a rock landmark contains up to five individual rocks, the camera position can be estimated using a single landmark. Here we assume that the spacecraft attitude is known, therefore a unit ray to the i^{th}

rock can be calculated as:

$$r_i = {}^{w}R^{c} \frac{1}{(x_i^2 + y_i^2 + f^2)^{1/2}} \begin{pmatrix} x_i & y_i & f \end{pmatrix}^T \quad (3)$$

where ${}^{w}R^{c}$ is the rotation matrix from camera frame to local map frame and (x_i, y_i) is the rock position in the image frame.

The ray $\boldsymbol{r}_i$ should agree with the reconstructed unit vector between the rock's 3D position, P, and the camera center, C, as:

$$r'_i = (P_i - C) / \| P_i - C \| = (P_i - C)/d_i = \Delta P_i / d_i \quad (4)$$

where $\Delta P_i = \begin{bmatrix} dp_{i1} & dp_{i2} & dp_{i3} \end{bmatrix}^T = (P_i - C) \quad d_i = \| \Delta P_i \|$

Because r_i and r'_i should be parallel and therefore the cross product of the two vectors should be a null vector:

$$\begin{pmatrix} f_1 & f_2 & f_3 \end{pmatrix}^T = r_i \times \Delta P_i / d_i = [r_i \times] \Delta P_i / d_i = 0 \quad (5)$$

where

$$[r_i \times] \equiv \begin{pmatrix} 0 & -r_{i3} & r_{i2} \\ r_{i3} & 0 & -r_{i1} \\ -r_{i2} & r_{i1} & 0 \end{pmatrix}$$

Then, the maximum likelihood estimation process involves finding $\boldsymbol{C}$ that minimizes the following function:

$$\begin{aligned} J(\hat{C}, r) &= \frac{1}{2} \sum_{i=1}^{N} (f_{i1}^2 + f_{i2}^2 + f_{i3}^2) \\ &= \frac{1}{2} \sum_{i=1}^{N} \frac{1}{d_i^2} \Delta P^T [r_i \times]^T [r_i \times] \Delta P_i = \min \end{aligned} \quad (6)$$

The merit function is robust regardless the length of $(P_i\text{-}C)$. Therefore, we can fix or ignore d_i in each iteration and the partial derivatives of $\boldsymbol{J}$ over $\boldsymbol{C}$ (ignoring d) are:

$$G(\hat{C}, r) = \sum_{i=1}^{N} \frac{1}{d_i^2} [r_i \times]^T [r_i \times] (P - \hat{C}_i) = \begin{pmatrix} 0 & 0 & 0 \end{pmatrix}^T \quad (7)$$

We assign a uniform d_i (e.g. 1) in the beginning of the process and $\boldsymbol{C}$ is then:

$$\hat{C} = \left[\sum_{i=1}^{N} [r_i \times]^T [r_i \times] \right]^{-1} \sum_{i=1}^{N} [r_i \times]^T [r_i \times] \, P_i \quad (8)$$

With each newly computed C, we can update d_i between iterations, then:

$$\hat{C}_k = \left[\sum_{i=1}^{N} \frac{1}{d_{k-1}} [r_i \times]^T [r_i \times] \right]^{-1} \sum_{i=1}^{N} \frac{1}{d_{k-1}} [r_i \times]^T [r_i \times] \, P_i \quad (9)$$

The iterations stop when $\| \hat{C}_k - \hat{C}_{k-1} \| < t$. This is a very simple efficient solution and usually only takes a few iterations to converge.

The new estimated spacecraft position can be then used to help find additional landmarks by comparing r from (3) and r' from (4) as follows:

$$| (r \cdot r') - 1.0 | < t \quad (10)$$

As a result, the two landmarks are considered to correspond to each other.

We use the quartet bearings of a landmark as the primary signature to locate its counterpart in the landmark database. In order to achieve real-time operation, three ordered combinations of the quartet bearing angles, $Q_1(q_1, q_2, q_3)$, $Q_2(q_1, q_2, q_4)$, and $Q_3(q_2, q_3, q_4)$, are used to register the address of each landmark in the landmark database. Three 90×90×90 degree cubic matrices are used to store the address of each landmark (FigureFigure 5). In theory, a total 729,000 landmarks can be stored. In the case when two or more landmarks are collocated in one cell, we just move one to an empty cell of its 26 neighboring cells. If the rock detection process is perfect, a landmark in the descent image should have the same bearings as its counterpart in landmark database and the bearing angle indices can be used to quickly locate its counterpart without searching. In actuality, these landmark descriptors are not always repeatable. However, even if their repeatability is very low, a search of cells surrounding the matched element should return the correct match in a short time.

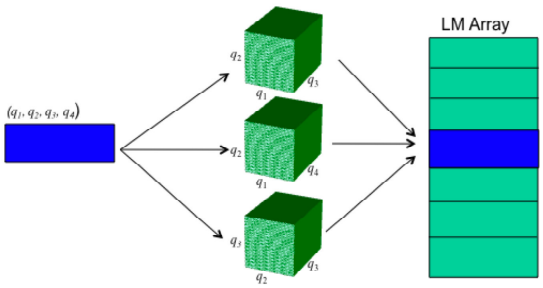

Figure 5: The landmark bearings are used as the primary signature for landmark matching.

The landmark-matching algorithm proceeds as follows:

1. Construct a rock landmark database of landing zone from orbital imagery.
2. Rectify a descent image to align it with the local map frame by using the onboard attitude reading.
3. Detect rocks from the descent image.
4. Construct rock landmark of the descent image.
5. Use quartet-bearing indices of a rock to locate its counterpart in the landmark database. In order to compensate for bearing angle uncertainty, the neighboring cells are also checked.
6. If a landmark candidate is located, the corresponding rock height and distance are checked using equation 2.
7. Estimate initial spacecraft position using equation (9)
8. Find additional landmarks using equations 9 and 10.

9. If the total number of matched landmarks is greater than a threshold, all matched landmarks are used for the final spacecraft position estimation. Search ends.
10. Otherwise check the next rock landmark.

3 Rock Landmark TRN Feasibility

Rock position, height and detection rate are the critical measurements for the proposed TRN algorithm. Both, how a rock casts its shadow and how accurate we can measure its dimensions directly influences TRN's performance. In general, the length of the shadow (which is a function of rock height, and shape), the atmospheric opacity and the sun elevation can all affect performance. For a TRN application, we are mostly interested in understanding the consistency of the rock detection algorithm under differing environmental conditions. We discuss these in detail next.

3.1 Atmospheric Opacity Effect

Reliable and robust shadow detection assumes that shadows are salient, which implies having a sharp contrast between the shadow and surrounding area. In natural scenes, however, there may be situations in which shadows are not clearly distinguishable. Consider, dark rocks, for example, or rock shadows cast on dark soil. These effects can be characterized by the albedo ratio between rocks and soil. Fortunately, surface albedo contrast is a trivial issue because it is captured by both orbital imagery and by the descent image, and has limited impact on rock detection consistency and performance. The second aspect is the contrast sensitivity of the shadow detector as a function of the optical depth (transparency) of the atmosphere. The Martian atmosphere, though thin in comparison with Earth's, does contain gases and aerosols that can attenuate the solar radiation. A dust storm, for example, can produce visible optical depth of 5 or more -- enough to completely block the Sun from view. As a consequence, shadows of rocks could be indistinguishable. On the other hand, optical depth rarely falls below 0.18, and during relatively clear Martian sols typical values for optical depth are between 0.2 and 0.5 [12].

Solar radiation at Mars has been well studied by the solar power community. The normalized net irradiance, f_{tot} (τ, z, A), is a function of altitude, total optical depth, τ, surface albeo, A, and solar zenith angle, z [11]. The total downward irradiance can be separated into its direct, f_{dir}, and diffuse, f_{diff}, components

$$f_{dir} = e^{-\frac{\tau}{\mu}}$$
$$f_{diff} = \frac{f_{tot}(\tau,\mu,A)}{(1-A)} - e^{-\frac{\tau}{\mu}} \quad (11)$$

where μ=cos z is the cosine of the solar zenith angle z and f_{tot} can be found or interpolated from two tables given in[11].

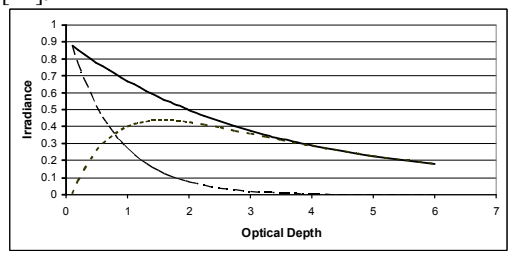

Figure 6: A solar irradiance curves (z = 55°) is a function of optical depth. Total irradiance (solid), direct irradiance (dashed) and diffuse irradiance (dotted).

In general, we are interested in the ratio between the f_{tot} and f_{diff} because it determines the contrast between the shadow, which receives diffuse radiation only, and the lit area, which receives both diffuse and direct radiation. The direct radiation declines sharply and monotonically with increased optical depth, while the diffuse components increase sharply and then fall off slowly with increasing optical depth [11] (Figure 6).

An image contrast simulator that emulates the physics of these effects was developed to assess shadow saliency. A representative sample of HiRISE image was used and the solar zenith and optical depth are 55° and 0.3 respectively. The shadows were segmented by applying the gMET algorithm (Figure 7). Since the solar zenith and the optical depth are known, the direct and diffuse irradiance curves can be readily obtained (Figure 6). These curves can be used for adjusting the image intensity of lit and shadowed regions. Special bilinear interpolation was applied to the shadow edge pixels because they are partially illuminated by both direct and diffuse irradiance. Figure 8 shows some simulated images under different optical depths. It clearly shows that if the optical depth increases, the image rock shadows become less salient and then the shadow segmentation becomes less reliable.

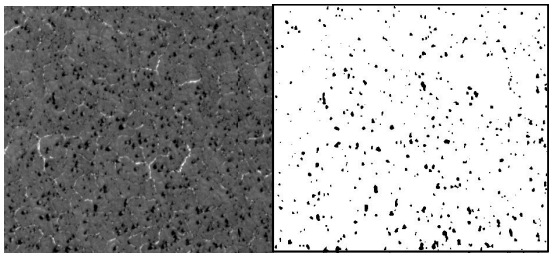

Figure 7: A piece of a HiRISE image and its shadow image are used in simulating atmospheric opacity effect.

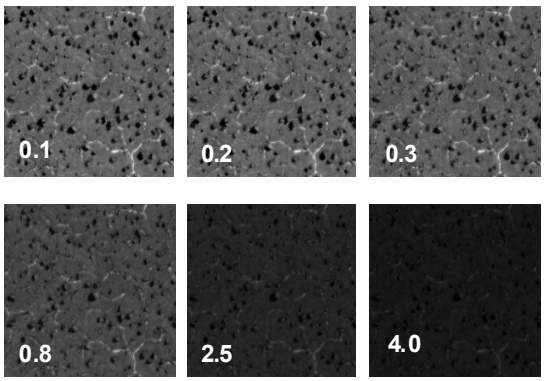

Figure 8: The simulated image contrast under different optical depth. The number in each image is the optical depth.

Rocks in each simulated images are detected and then compared with the rocks from the original image. We evaluated the effect from three aspects: detection rate, rock position accuracy and rock height consistency.

The top graph in Figure 9 shows the detection rates as function of optical depth. Between optical depths 0.1 to 0.9, the rock detection rate declines slowly from close 100% to 70%. Above 0.9, the detection rate declines sharply. This indicates that when the optical depth is greater than 0.9, the shadow based rock detection becomes less reliable.

The middle graph in Figure 9 shows rock position accuracy in pixels as function of optical depth. In general, rock position error caused by atmospheric opacity is within 1 pixel, which is adequate for spacecraft localization.

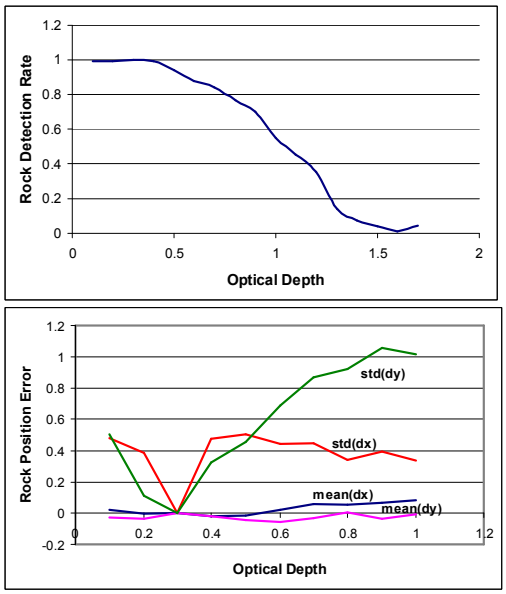

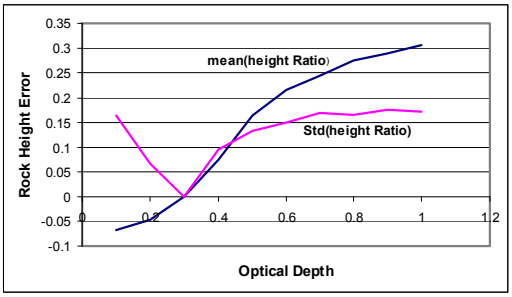

Figure 9: Rock detection consistency evaluation.

The bottom graph in Figure 9 shows rock height consistency under the influence of atmospheric opacity. The rock error is calculated as $e = (h - h')/h$, where h is the rock height estimated from the original image, h' is rock height from the simulated image. This figure shows that when the optical depth is less than 0.3, because the direct flux is much higher and the rock shadow is much sharper, the shadow segmentation will tend to delineate a larger shadow and as consequence, the estimated rock is systematically taller. On the other hand, when the optical depth is higher, the diffuse flux blurs shadow edges and the segmented shadow regions, thus the estimated rock heights are shorter. In a nominal atmosphere condition ($0.2<\tau<0.5$), the rock height error caused by optical depth is less than 15%.

3.2 Illumination Effect

In order to evaluate these measurements we have used the rock size-frequency distribution models proposed in [8] for Mars. These include rock abundance and size distributions for Mars missions going back to the Viking missions. Surface counts are the result of an extended effort to count and map the rocks at the landing sites of the Viking Landers (VL1 and VL2) using monocular and stereo tools applied to the images acquired by the lander sensors. Such mappings have been instrumental in the derivation of the rock abundance models[8][10][9]. For precise evaluation of height estimates we used three different HiRISE images (PSP_001501, PSP_001976, and PSP_002055). The local sun elevations of these three images are 39°, 32° and 33° respectively. The rocks in these three images are detected by the rock detection algorithm described above. In addition, rocks in one of these images were manually measured. Both manual and surface count curves agree with automatic detection curves regardless of the sun elevation. This indicates that if the sun elevation during descent is within about 10° of the orbital image, the effect on rock height estimation is quite limited.

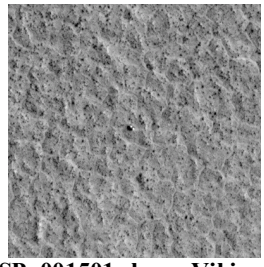

Figure 10: PSP_001501 shows Viking lander 2, the largest black dot in the middle of this image.

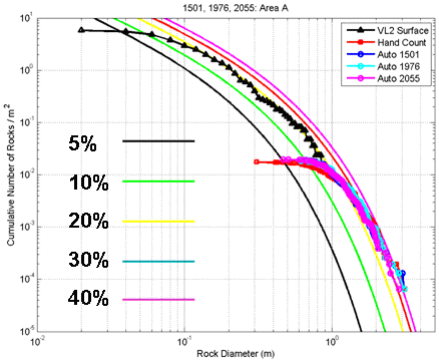

Figure11: Rock distribution curves of auto counted, hand counted and surface counted match each other. Rocks under 1m in diameter are not fully automatically or manually resolved, hence the curve roll off.

3.3 Rock Landmark Reliability

The previous section shows that variations in environmental conditions preclude perfect replication of rock detection behavior. For a TRN application, we are concerned with how much these inconsistencies affect the rock landmark formation described before.

We used the rocks collected from image PSP-001394 (Figure 2) to conduct a repeatability study. A rock landmark database was constructed from this data set. A second landmark database was constructed by adding Gauassian noise to the rock heights of the first database according to $h' = h + nh$, where n=N(0.2, 0.2) is randomly generated. The rock landmarks in both databases were compared to each other. If a rock and at least three of its neighbors are the same in both databases, we call it *repeated*.

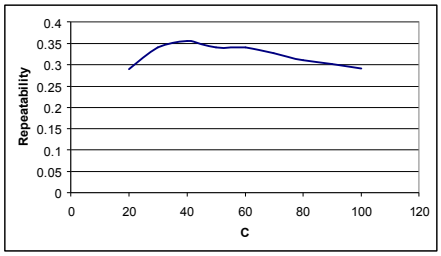

Figure 12: Rock landmark repeatability is about 30% in this test case.

Figure 12 shows the rock landmark repeatability for this test case. The c in the figure is the c in Equ. 1. Even though only about 30% landmarks are repeated, a few search iterations (< 20 times) are able to find the correct matching with >99% success probability.

3.4 Rock Landmark TRN Speed Performance

We evaluated the TRN speed on a Pentium 2.13 GHz under Windows XP. The speed of the rock detection algorithm is dependent on the image size, and the number of rocks in it. We tested its speed on a densely populated image. A total of 3,063 shadow regions are extracted in 14 msec. Of these, 1,250 are larger than 5 pixels in area and are therefore modeled as rocks, requiring another 24 msec. The rock landmark matching, including landmark formation and matching, takes about 30 msec. Therefore the total execution time from rock detection to final pose estimation is less than 70 msec. The same run on a 100 mHz processor would result in a expected execution time of less than 2 seconds, which is fast enough performance for a real mission application.

3.5 Rock Landmark TRN Reliability & Accuracy

A Monte Carlo simulation was conducted over the area of HiRISE image PSP-001391 (FigureFigure 3). The descent camera sensor has 1024 by 1024 pixels with 6^o field of view optics. The simulation started at 4000 m altitude and ended at 2000 m above ground with 500 m decrements.

At each altitude, 52 descent images with random attitudes and x and y positions were created using PSP-001391 as terrain texture map. In addition, an atmospheric opacity effect (τ=0.5) was added to these images.

Then the proposed rock TRN algorithm was used to estimate the spacecraft positions using these images. Figure 13 shows the position uncertainty of the proposed algorithm.

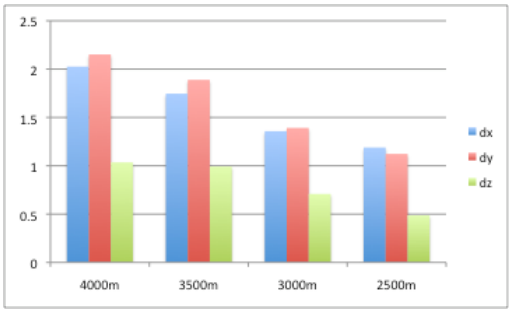

Figure 13: The position uncertainties in three axes are less than 3 meters in this study, which is excellent for pinpoint landing.

The success rates for the four altitudes are 0.80, 0.85, 0.92 and 0.98 respectively.

4 Conclusions

A novel rock landmark TRN algorithm is suggested in this paper. The advantages of this approach are:

1. It is simple and fast. One TRN cycle takes less than one second, and can therefore be carried out easily by currently available onboard computing power.
2. Its localization accuracy is less than 3 meters at 4 km above ground, which is adequate for PPL applications.
3. Its required storage space is much smaller than that of a high resolution image. For example, for a 20 by 20 km^2 landing site, a 2D 1m/pixel base map needs 1.2 gigabytes memory space, while the rock landmark database for the same area only needs less than 4 megabytes memory requirements.
4. It is shown to be very reliable on Mars northern plains terrain where abundant rocks are present.
5. It can handle large-scale uncertainty thus avoiding the need to rely on accurate altimetry data.

The algorithm, however, is applicable only to areas where a large number of rocks are present such as those on the northern plains of Mars

Acknowledgment

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.

References

[1] A. Johnson and J. Montgomery, "An Overview of Terrain Relative Navigation for Precise Lunar Landing," Proc. IEEE Aerospace Conference (Aerospace08), Big Sky, MT, March 2008.

[2] D. Adams, T. Criss, U. Shankar, "Passive Optical Terrain Relative Navigation Using APLNav", IEEE AC, pp. 01-22, 2008.

[3] A. Mourikis, N. Trawny, S. Roumeliotis, A. Johnson A. Ansar and L. Matthies, "Vision Aided Inertial Navigation for Spacecraft Entry Descent and Landing, IEEE Transactions on Robotics and Automation, vol. 25, no2, pp. 264-280, 2009.

[4] Y. Cheng and A. Ansar, "Landmark Based Position Estimation for Pinpoint Landing on Mars," IEEE Int. Conf. on Robotics and Automation, Barcelona, Spain, 2005.

[5] Adnan Ansar and Yang Cheng, "An Analysis of Spacecraft Localization from Descent Image Data for Pinpoint Landing on Mars and Other Cratered Bodies", PE&RS Vol. 71, Number 10, October 2005.

[6] Huertas, A., Y. Cheng, and R. Madison, "Passive imaging-based multi-cue hazard detection for spacecraft landing", IEEE Aerospace Conference, Big Sky, Mont, March 2006.

[7] R. Gaskell, et al, "Landmark Navigation Studies and Target Characterization in the Hayabusa Encounter with Itokawa," AIAA/AAS Astrodynamics Specialist Conference and Exhibit, Keystone, Colorado, 2006.

[8] Arvidson et al, Mars Exploration Program 2007 Phoenix landing site selection and characteristics, J. Geophys. Res., 113,

[9] Golombek et al, "Size-frequency distributions of rocks on the northern plains of Mars with special reference to Phoenix landing surfaces," J. Geophys. Res., 113,

[10] Golombek, M.P., "Rock statistics calculations for the MER landing sites", 3rd MER landing site workshop. March, 2002.

[11] Robert M Haberle et al, "Atmospheric Effects on the Utility of Solar Power on Mars" in *Resources of near-earth space*. Edited by John S. Lewis, Mildred S. Matthews, and Mary L. Guerrieri. Space Science Series. Tucson, London: The University of Arizona Press, 1993., p.845

[12] http://ccar.colorado.edu/asen5050/projects/projects_2001/benoit/solar_irradiance_on_mars.htm