INTEGRATION OF ORBITAL AND GROUND IMAGE NETWORKS FOR THE AUTOMATION OF ROVER LOCALIZATION

Ju Won Hwangbo, Kaichang Di, and Rongxing Li
Mapping and GIS Laboratory
Dept. of Civil & Environmental Engineering & Geodetic Science, The Ohio State University
470 Hitchcock Hall, 2070 Neil Avenue, Columbus, OH 43210-1275
hwangbo.2@osu.edu

ABSTRACT

Rover localization is essential to the exploration of space. The availability of sub-meter resolution satellite imagery, especially HiRISE imagery, has opened the possibility of computing rover locations at higher accuracy by making use of detailed features seen in the satellite orbital images. This paper describes a new development towards automation of the rover localization process using orbital and ground images. Using a rigorous sensor model and bundle adjustment (BA) of HiRISE stereo imagery, high-resolution orthophotos and DEMs can be generated as soon as the orbital stereo images are acquired. A ground image network is also constructed using intra- and inter-stereo matching. From both types of imagery, a few landmarks are identified to be used as ground control points for the integration of the orbital and ground image networks. Rocks detected from both orbital and ground imagery serve as tie points for rover localization. From orbital images, rocks are extracted based on brightness values and the shape of dark spots. Rocks in ground images are extracted through dense stereo matching, rock peak and surface point extraction, and rock modeling. To narrow down a precise rover position, terrain match is performed using DEMs generated from orbital and ground imagery. Finally, distribution pattern matching is implemented for rocks detected from orbital and ground imagery. The rover position is adjusted based on 2D affine transformation obtained from rock pattern matching. The proposed method has been tested for the Spirit rover traverse. Experimental results show that the orbital/ground rock matching approach has performed successfully for MER rover localization.

INTRODUCTION

In space exploration missions, the most critical information that needs to be updated constantly is the location of mission components. Depending on the nature of each mission, these components can be unmanned vehicles (rovers) or astronauts. Without the availability of Global Positioning Systems (GPS) in the outer planets, the localization process is a challenging task. For that reason, orbiters are usually deployed before sending any lander to planets. One of the major objectives of orbiters is to gather a global dataset able to provide topographic information for the next lander mission. The current Mars Exploration Rover (MER) mission, started in January 2004, has far exceeded its initial goals in terms of distance traveled and operational lifetime (Arvidson et al., 2004; Li et al., 2004, 2005). In future planetary surface robotic operations, a rover would be expected to drive more than the 9 km that has been achieved by the Opportunity rover (Li et al., 2007). To achieve the longest possible exploration within the lifetime of a rover, it is desirable to travel longer distances within each command cycle. The MER vehicles have performed over 60,000 coordinated motions (the powering of either steering or drive motors continuously), activities that demand nearly constant attention (Maimone et al., 2007). Without highly detailed information about the terrain to be explored, autonomous navigation is jeopardized and little drive progress can be made. With the advent of sub-meter resolution HiRISE imagery, it is now possible to obtain highly detailed topography of the Martian surface. The crucial question is how to locate the rover in the context of topographic products generated by orbital datasets. This is also a key to the success of autonomous mobile robot navigation in the future.

Various methodologies have been researched for mobile robot localization and navigation. In the Mars Pathfinder mission, the rover Sojourner achieved an overall localization error of about 10% of the distance from the lander within an area of about 10x10 m using the dead-reckoning method (Matthies et al., 1995). The Jet Propulsion Laboratory (JPL) worked on several advanced methods including probabilistic self-localization (Olson, 2000), and visual odometry (VO) algorithms (Maimone et al., 2007). The Robotics Institute at Carnegie Mellon University (CMU) designed and developed robotic systems and vehicles. Field experiments gave a localization accuracy of 3-5% of the distance traveled using dead-reckoning technology that integrated wheel encoders, roll and pitch inclinometers and yaw gyro (Wettergreen et al., 2005). During surface operations of the MER mission, onboard positions of both rovers

are estimated within each sol by dead reckoning using the wheel encoders and IMU. The heading is also updated occasionally by sun-finding techniques using the Pancam images. A designed accuracy of 10% has been achieved (Li et al., 2004). Without external references from GPS systems, the dead-reckoning approach is subject to accumulative errors due to the measurement of relative pose. In the MER mission, the dead-reckoning error has accumulated mainly due to slippage between the rovers' wheels and the ground. Visual odometry has been used onboard the rover for precision instrument placement and the correction of slippage (Maimone et al., 2007). For rover localization over relatively long traverses, incremental BA technology is used to refine the rover positions (Li et al., 2004, 2007). Incremental BA provides accurate rover positions by building a strong image network along the traverse to maintain consistent overall traverse information. Although VO is fully automated onboard the rover, BA is conducted manually because of the difficulties in selecting cross-site tie points automatically. Li et al. (2007) proposed a new approach for automatic cross-site tie point selection using rock extraction, rock modeling, and rock matching. Both VO and BA are limited to the ground image network, leaving unanswered the question of finding the exact position of the rover in the context of topographic products from orbital imagery. Due to the inconsistencies between the orbital and ground image networks, updating rover positions in a HiRISE orthophoto is performed manually by identifying common features from both datasets.

This paper introduces a new method for automation of the rover localization process having an integrated orbital and ground image network using rock extraction, terrain matching and rock distribution pattern matching.

INTEGRATED ORBITAL AND GROUND IMAGE NETWORK

Typically, orbital image networks are built based on stereo tie points and ground control points. In the case of Mars, a series of orbiter missions before Mars Reconnaissance Orbiter (MRO) provided a lower resolution image network as well as vertical control. The widely used Mars Global Digital Image Mosaic (MDIM) was produced by the U.S. Geological Survey (USGS) using Viking orbiter images. With a ground resolution of 213 m, the latest version of MDIM (version 2.1) has an accuracy (RMS error of the control network) of about 250 m (Archinal et al., 2003). Currently, the best global Mars terrain model was acquired by MOLA (Mars Orbiter Laser Altimeter). The MOLA terrain model with a spatial resolution of about 1 degree (or 59 km at the equator) has an absolute vertical accuracy of 13 m with respect to the Mars center of mass (Smith et al., 1999). The higher-resolution MOLA DTM known as Mission Experiment Gridded Data Record (MEDGR) is available. Most importantly, from the mapping perspective, MEDGR provides a global control for registering orbital images from current and future missions. In recent years, regional maps of certain important sites have been generated using MOC (Mars Orbiter Camera) images from the MGS mission and THEMIS (Thermal Emission Imaging System) images from the Mars Odyssey mission (Kirk et al., 2003). The topographic maps generated from MOC NA (Narrow Angle) stereo images have a horizontal resolution of 5 m and an expected vertical precision of 1 m (Kirk et al., 2003). Shan et al. (2005) reported their efforts in photogrammetric registration of MOC NA imagery to MOLA profiles and automatic digital elevation model (DEM) generation from MOC NA stereo images. The Mars Express HRSC (High Resolution Stereo Camera) images have been used to map the Martian surface at high resolution (10 meters/pixel) and at super-high resolution (2 meters/pixel) in selected areas (Scholten et al., 2005).

When processing HiRISE imagery, the MEDGR gridded product having a resolution of 463 m per pixel provides the best source of ground control. However, a big difference in planimetric resolution prevents using it as horizontal control. Kirk et al. (2008) used coordinates measured from the 1.5 m/pixel MOC images previously controlled to MOLA for horizontal control. We adopted MOLA altimetry for vertical control but use landmarks derived from the MER dataset for horizontal control as well. In fact, the Spirit rover landing position is the origin of Local Coordinate System (LCS) and is centered at 175.47848 º in longitude and −14.571892 º in latitude. Figure 1 shows the process of integrating orbital and ground BA. First, orbital and ground imagery is separately processed to generate stereo tie points. The MER image network is built based on intra- and inter-stereo matching followed by BA of panoramic images in one site. A few landmarks identified from both orbital and ground images are selected as horizontal control for the adjustment of orbital exterior orientation (EO) parameters. Vertical control is obtained from the MOLA altimetry dataset. After the BA of orbital images is conducted, dense matching is performed to generate 3D point clouds and, eventually, a digital elevation model (DEM). A ground DEM is also produced after adjustment within one site. As mentioned earlier, the current MER mission incorporates dead reckoning and visual odometry as well as incremental BA using cross-site stereo matching. However, all of these efforts do not guarantee a perfect match between the topographic products from orbital imagery and an actual rover position. Therefore, common features are identified to connect the two image networks. The most recognizable objects from both dataset are rocks, especially

small ones (less than 1 m in length). Unlike rock matching between cross-site ground images, rocks seen from orbit could not be modeled or classified into different types due to the limited resolution of the HiRISE sensor. If a rock is 30 cm in diameter, it will occupy only one or two pixels in an orbital image. Currently, modeling small rocks from an orbital dataset does not yield accurate results due to the difficulty in distinguising the actual rock from its shadow. Rocks generated from both datasets are matched based on the distribution pattern. To prevent mismatches due to a lack of model parameters, the rover position is preliminarily estimated by comparing the DEMs from the two datasets. After terrain matching, the search area is narrowed down and common rocks are identified using rock pattern matching. Finally, the matched rocks serve as tie points for the integrated image network and the rover position is updated within the context of the orbital topography.

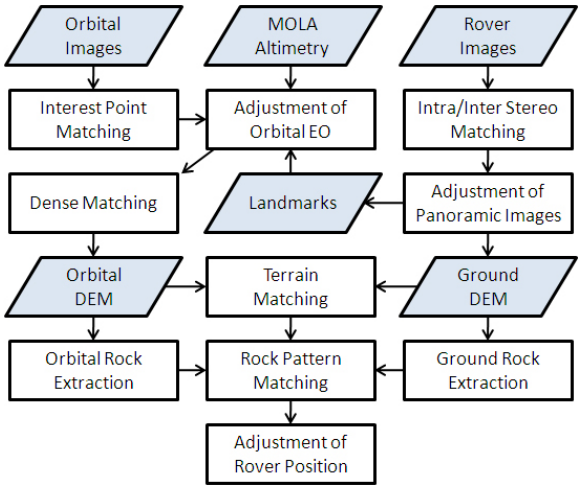

Figure 1. Conceptual flowchart of the integration of orbital and ground data.

Rigorous Modeling of HiRISE Geometry

HiRISE is a push-broom imaging sensor with 14 CCDs (10 red, 2 blue-green and 2 NIR). Each CCD consists of a block of 2048 pixels in the across-track direction and 128 pixels in the along-track direction. Ten CCDs covering the red spectrum (700 nm) are located in the middle. In the across-track direction, average overlap width between adjacent CCDs is about 48 pixels. However, the alignment of CCDs involves small shifts and rotations with regarding to the HiRISE optical axis. After excluding overlapping pixels, HiRISE can generate images with a swath of up to 20,264 pixels (across-track) and a resolution of 30 cm/pixel at an altitude of 300 km (McEwen et al, 2007). At such a high resolution, the IFOV (instantaneous field-of-view) is extremely small and, as result, the ground track speed becomes very fast. To improve the signal strength of "fast-moving" objects and to increase the exposure time, Time Delay Integration (TDI) technology has been incorporated in the instrument. As the MRO (Mars Reconnaissance Orbiter) spacecraft moves above the surface of Mars, TDI integrates the signal as it passes across the CCD detector by shifting the accumulated signal into the next row (line) of the CCD at the same rate as the image moves (line rate of 13,000 lines/sec = 1 line every 76 microsecond). Signals in each TDI block are transferred from line to line at ground track speed. A single pixel is formed by accumulating signals from the TDI block. HiRISE can use 8, 32, 64 or 128 TDI stages to match scene radiance to the CCD full well capacity. According to the HiRISE instrument information from SPICE kernels (available at ftp://naif.jpl.nasa.gov/pub/naif/MRO/kernels/ik), the observation time of a single pixel is defined as the Ephemeris Time (ET) when the center of the TDI block is exposed. The HiRISE instrument kernel provides the calibrated interior orientation parameters needed to calculate the pixel view direction with respect to the HiRISE frame (MRO_HIRISE_OPTICAL_AXIS). In the raw image, the row position of each pixel is related to the Ephemeris Time, which then determines the position and orientation of the HiRISE frame.

Bundle Adjustment of Orbital EO

Exterior orientation (EO) parameters, the positions of the camera perspective center and pointing angles at a specific time are provided in the SPICE kernels. The EO parameters of each image line can be retrieved by interpolating the spacecraft's trajectory and pointing vectors. In order to retrieve EO parameters from the SPICE kernel, the number of image lines along with the starting and ending times for a particular CCD image are needed. Ten sets of EO parameters are needed for a full red band HiRISE mosaic. However, since all CCDs are fixed to the HiRISE focal plane, they all share the same EO parameter at a particular time. Therefore it is possible to model EO parameters

from different CCDs by incorporating the simple offset between them. Thus, images simultaneously acquired by multiple CCD arrays can be processed together under a uniform rigorous sensor model in the BA instead of being processed strip by strip. The difficulty with the raw EO parameters is the discontinuities at CCD boundaries. The discontinuities can be removed using offsets calculated from the tie points in the overlap area between CCDs. Significant research has shown that the change in EO parameters over short trajectories can be well modeled using polynomials (Shan et al., 2005; Li et al., 2007, 2008). In this study, using the line (image row) index as a variable, third-order polynomials are selected to model the EO parameters. The best-fitting polynomial coefficients are calculated based on the least-squares method. Positions of the spacecraft can be modeled precisely with very small residuals. However, a large residual from the image-pointing parameters reveals an erratic signal that cannot be modeled by a polynomial. This signal represents jitter, or small motions of the spacecraft around its nominal pointing. The jitter effect was originally identified in the MOC images, but was found to be more severe for HiRISE because of HiRISE's higher resolution (Kirk et al., 2008). As it is caused by unpredictable forces acting on the camera, jitter shows a random pattern and is difficult to incorporate into a mathematical model. Jitter can be removed simply by adjusting the image coordinates. For each image coordinate, the pointing vector from the focal point is calculated and then projected onto the modeled focal plane based on the best-fitting polynomials. The intersection of the modeled plane and the pointing vector can be used as the new image coordinates free of the jitter effect.

The variability of the atmosphere of Mars requires a comparable adjustment of the spacecraft position. Kirk et al. (2008) reported more that 1 km of movement in the along-track direction and 260 m in the across-track direction. With a looser radial weighting, the vertical movement of the spacecraft was 500 m. This magnitude of adjustment, large RMS residuals (more than a few meters), and considerable scale difference with regarding to MOC DTM left absolute accuracy of the BA questionable. Obviously, further research is required to solve this problem. Our solution is to conduct free adjustment with tie points and to fit the HiRISE DTM with horizontal control points from the MER traverse and vertical control points from the MOLA altimetry.

HiRISE Stereo Matching

To achieve the full potential of the very high resolution data without jeopardizing quality, the geometry of the HiRISE sensor is used in stereo matching in addition to the correlation of image brightness values. This geometry is incorporated into the hierarchical matching process to provide a solid structure for stereo matching in higher-resolution levels. Matching starts with the images of lowest resolution; results are then transferred to the next higher level (higher resolution) with more interest points being extracted and matched. A hierarchical matching strategy enables utilization of the topological and geometrical relationships among the features at subsequent levels. This strategy is based on image pyramids generated from the original images. We started from an image scaled down to $1/16^{th}$ of the original; then stereo matching is continued at the $1/8^{th}$, $1/4^{th}$, $1/2^{nd}$, and original scale images. At the lowest level, the geographic locations of interest points are estimated by assuming a flat terrain. This enables automatic pairing of interest points in stereo images. The search radius is confined to the neighborhood of the point of interest. The cross-correlation coefficients are calculated for interest point pairs from each image within the search range. Interest points from both images that mark each other as the matching point with the highest correlation are chosen as a matching pair. The parameters to determine which matching pair is acceptable can be customized at each level. The search window size, the correlation window size, and the correlation coefficient limit all affect the automatic image matching result. At each subsequent level, points from the previous level are matched again to achieve higher matching precision. Then a TIN (Triangulated Irregular Network) parallax surface is generated from these matched points using the Delaunay triangulation method. This TIN is used to estimate the corresponding tie points. If a point is outside of the Delaunay triangles, the parallaxes of the nearest points are used to estimate the tie point coordinates. To improve matching performance for points located around the boundary of each CCD, HiRISE imaging geometry is fully utilized. Instead of mosaicking images from separate CCDs, our approach links the TIN surface using the best-fitting alignment derived from interest point matching between adjacent CCDs.

Automatic blunder (mismatch) detection is performed at each level by eliminating outliers based on the distribution of the elevation of neighboring points. For each point, a small local DEM surface is constructed from the matched points and modeled as a plane. Then the standard deviation of the plane estimation, σ, is calculated. If the residual of a point exceeds a certain threshold in terms of σ, it is considered an error and eliminated.

After matching the interest points generated from the highest-resolution images, 50-pixel grid points are defined, followed by 10-pixel grid points, to form a basis for further matching. To generate a 1-m-resolution DEM of the terrain, 2-pixel grid points are matched. For one image stereo, the overall processing time needed for the 3D point cloud generation at a 2-pixel grid level is a few hours. Evenly distributed tie points between the stereo images are selected from the set of matched interest points to be used in the subsequent BA. The final DEM is generated after BA and elimination of matching errors. Interpolation of a 1-m-resolution DEM takes about 12 hours for 80,000-line image

stereos. For generation of sub-meter level DEMs, dense matching is performed for every pixel in the images of highest resolution.

Incremental Bundle Adjustment of the Ground Image Network

The MER rover is equipped with four pairs of stereo camera systems: panoramic cameras (Pancam), navigation cameras (Navcam), and forward- and rear-looking hazard cameras (Hazcam). Hazcam, Navcam, and Pancam cameras are used for imaging the terrain at close-, middle-, and far-range, respectively. They all are 1024×1024-pixel frame-transfer digital cameras (Maki et al., 2003). Navcams and Pancams are installed on the Pancam Mast Assembly, which can rotate 360° in azimuth and ±90° in elevation. A panoramic image with 10% overlap requires 10 pairs of Navcam images or 27 pairs of Pancam images. A Navcam panorama consists of 20 images configured as one vertical row (tier); it can be generated within 30 minutes. A Pancam panorama usually consists of three tiers and takes several days to complete. Hazcams are installed both front and rear just under the solar panel. They can perform real-time feature matching and build up the terrain slope in the close range, to help the rover to select a safe route.

MER uses Pancam as a sun sensor finder. MER is also equipped with an IMU and an odometer to record its movements. These instruments work well most of the time, except when climbing a hill or crossing surfaces covered with soft soil, when slippage can occur and the traverse calculated by IMU and odometer will no longer be accurate. Therefore, the pointing parameters of each image in the network (camera center position and three rotation angles) are adjusted to their optimal values using the least-squares method, with landmarks serving as tie points. This approach depends on the availability of distinguishable landmarks detected by the rover and shown on the rover images.

The BA for improving rover localization results consists of three steps: within-site bundle adjustment (to remove the within-site inconsistency), cross-site rigid transformation (to improve initial parameters), and cross-site bundle adjustment (to refine parameters iteratively). From the observation equation (Equation 1a) and the error equations (Equations 1b and 1c), the BA in Equation 1d is derived. Although the solution cannot be an absolute answer without ground control points, it is an optimal one. Equation 1e represents the rigid transformation, in which the transformation matrix B (representing rotation and translation) is calculated from cross-site tie points (Xu, 2004):

$$(x, y) = f(X, Y, Z, X_s, Y_s, Z_s, \omega, \varphi, \kappa) \quad (1a)$$

$$v = a_1dX + a_2dY + a_3dZ + a_4dXs + a_5dY_s + a_6dZ_s + a_7d\omega + a_8d\varphi + a_9d\kappa - l \quad (1b)$$

$$V = AX - L \quad (1c)$$

$$X = (A^TPA)^{-1}A^TPV \quad (1d)$$

$$P_2 = BP_1 \quad (1e)$$

where (x, y) is the image coordinate of the tie points, (X, Y, Z) is the object coordinate of the tie points, (X_S, Y_S, Z_S) is the camera coordinate, (ω, φ, κ) is the camera attitude, l is an observation, L is observation vector, v is the residual of the observation, V is the residual vector, $a_{1\text{-}9}$ are the partial derivatives of the collinear equations, A is a matrix containing the partial derivatives from each observation, X is the unknown vector (X, Y, Z, X_S, Y_S, Z_S, ω, φ, κ), P is the weight matrix, B is the matrix of rigid transformation, P_1 represents the homogeneous coordinates of tie points before adjustment, and P_2 represents the homogeneous coordinates after adjustment. After the BA, the new camera position is calculated by applying a rigid transformation to renew the rover's location (Equation 1e).

ORBITAL AND GROUND DATA INTEGRATION

In MER mission, the correspondence of the rover position to the orbital image is crucial. From locating the target to finding the safest route, when ground data is not available the information needed for rover navigation comes from orbital orthophotos and DEMs. According to statistical analysis of the MER Spirit rover traverse conducted by The Ohio State University, without using ground control the maximum difference between a bundle-adjusted rover track and the HiRISE orthophoto without correction was about 90 m for a total traverse length of 6 km. After applying manual adjustment, the maximum difference was reduced to 9 m. Since orbital images do not have sufficient resolution to match model parameters with ground rocks, the distribution pattern is the only criterion for rover matching. If slippage causes major errors in odometer readings, the worst-case scenario would be to incorrectly locating the rover position solely because the rock distribution pattern matches. To prevent such a case, 3D terrain matching is proposed to narrow down possible rover positions based on two terrain signatures, elevation and slope.

Terrain Matching

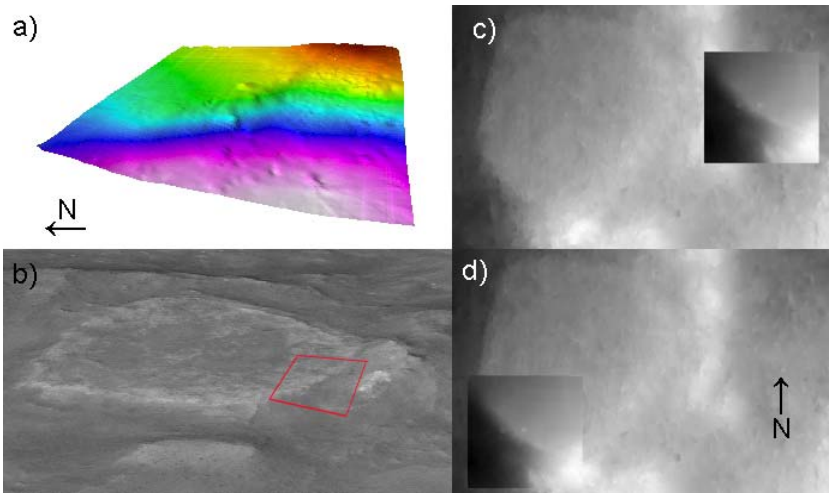

Figure 2. DEMs: a) ground DEM of site 13136 in 3D view, b) orbital DEM of Home Plate in 3D view (red rectangle is identified coverage of site 13136 DEM), c) ground DEM overlayed on orbital DEM before terrain match, and d) ground DEM overlayed on orbital DEM after terrain match.

The common 3D terrain matching method treats the DEM as an image and matches DEMs according to image correlation. To compare relative variations in elevation, the mean elevation is removed from each DEM. The correlation of DEMs is calculated as stated as:

$$DEM_o'(x, y) = DEM_o(x, y) - (\Sigma_x\Sigma_y\, DEM_o(x, y)) / (n_x \cdot n_y) \quad (2a)$$
$$DEM_g'(x, y) = DEM_g(x, y) - (\Sigma_x\Sigma_y\, DEM_g(x, y)) / (n_x \cdot n_y) \quad (2b)$$

where (x, y) are the ground coordinates, DEM_o is the orbital DEM, DEM_g is the ground DEM, n_x is the number of grid pixels in the x direction, and n_x is the number of grid pixels in the y direction. The correlation of DEMs measures the similarity of the overall pattern of the terrain.

The slope information (defined as inclination of the terrain in 3D space) is derived from the DEMs. Incorporating slope helps to match distinctive features such as ridges and hills, showing abrupt changes in elevation. The correlation of slope is obtained in the same manner as that of DEMs, except for using absolute values:

$$Cor_{DEM}(\Delta x, \Delta y) = [\Sigma_x\Sigma_y(DEM_o'(x, y) - DEM_g'(x+\Delta x, y+\Delta y))^2 / (n_x \cdot n_y)]^{1/2} \quad (3a)$$
$$Cor_{Slope}(\Delta x, \Delta y) = [\Sigma_x\Sigma_y(Slope_o(x, y) - Slope_g(x+\Delta x, y+\Delta y))^2 / (n_x \cdot n_y)]^{1/2} \quad (3b)$$

where $(\Delta x, \Delta y)$ is the shift of the coordinates, Cor_{DEM} is the correlation of the DEMs, Cor_{Slope} is the correlation of slope, $Slope_o$ is slope generated from the orbital DEM, and $Slope_g$ is slope from the ground DEM. Finally, terrain matching is done by identifying the position with the maximum weighted sum of correlation coefficients. The weights for each type of correlation coefficient are decided based on the ratio of the standard deviation of the elevation over that of the slope:

$$Cor(\Delta x, \Delta y) = w_1 \cdot Cor_{DEM}(\Delta x, \Delta y) - w_2 \cdot Cor_{DEM}(\Delta x, \Delta y) \quad (4)$$

where Cor is the correlation of the terrain, and w_1 and w_2 are the weight parameters. The position of the rover is shifted to where the correlation of the terrain is the highest. Figure 2 shows the terrain matching result from site 13136. Figure 2a is the ground DEM generated by Navcam panoramas. The area covered by ground DEM is identified as the southern part of Home Plate (Figure 2b). However, without integrated BA of the orbital and ground image networks, the distance between the rover position from telemetry data and the manually identified rover position on the orbital orthophoto was 90 m (Figure 2c). This telemetry data is available at NASA's Planetary Data Systems (PDS). After terrain matching, this difference decreased significantly, to 0.22 m (Figure 2d). Considering that the resolutions of the DEMs, is different, it is neither possible nor necessary to find the exact position. The objective of terrain matching is to

obtain a reasonable estimation so that the search area can be reduced for the rock distribution pattern matching. However, if the terrain has little variation, as in a flat region, the matching result could be erroneous. In that case, use of terrain matching result as an estimation should be limited.

Rock Extraction from Ground Imagery

Rocks are one of the major features found in rover imagery of the Martian terrain. Gor et al. (2001) used image intensity information to detect small rocks and range information to detect large rocks from Mars rover images. Thompson et al. (2005) developed a rock detection method based on segmentation, detection, and classification using texture, color, shape, shading, and stereo data from the Zoë rover. Li et al. (2007) proposed new methods of rock modeling from multiple rover sites for autonomous cross-site tie point selection. Their strategy is to identify rock peaks and model the surface points as hemispheroid, semi-ellipsoid, cone, or tetrahedron. Rocks extracted from cross-site Navcam images then are matched based on the distribution pattern and model parameters.

In orbital-ground data integration, the objective for rock extraction is to identify rocks that are visible in both orbital and ground imagery. Since rocks are to be used as tie points, size is also an important factor. As the resolution of HiRISE images is about 30 cm, rocks with diameters less than 30 cm would not be visible from orbit. On the other hand, if the rock is too large, it would not be accurate enough to be used as a tie point. The ideal size of rock for this purpose is one that represents a few pixels in the orbital image. Typically, rocks are identifiable from orbit due to their shadow and the difference in image intensity as distinct from the terrain. Therefore a rock extraction strategy for orbital-ground integration should incorporate image intensity as well.

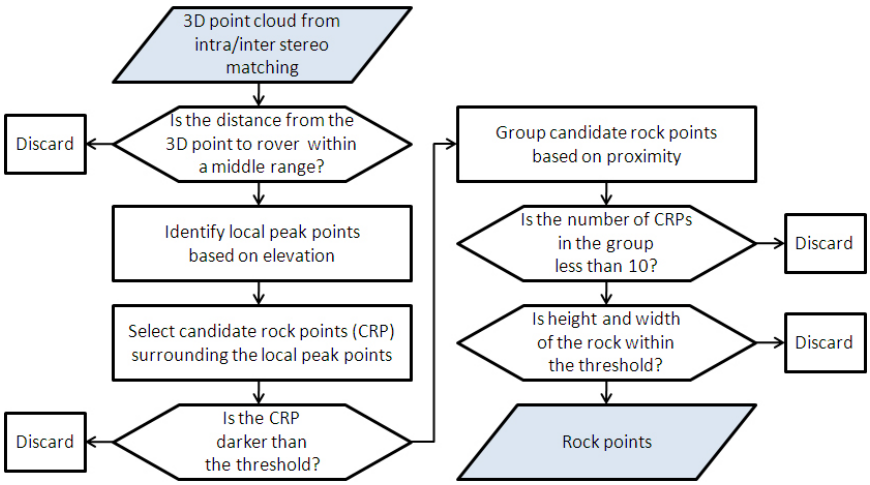

Figure 3. Conceptual flowcharf of the ground rock extraction process.

Figure 3 illustrates the ground rock extraction process.

1. From the 3D point cloud generated by Navcam intra-stereo matching, points that are too close or too far are eliminated. Navcam was originally designed with the intention of imaging the middle range (3~25 m). Points beyond that range have lower positional accuracy.
2. A rock is composed of a continuous surface and a peak that is higher than the terrain surface. Therefore a distinction between the terrain and the rock comes from differences in elevation. To find rock peaks, local peaks with maximum elevation from candidates within a close range (1 m) are identified. The height of a point is defined as the elevation subtracted by that of the lowest point within 1 m. To be considered as a rock peak, the height of local peaks should exceed 22 cm.
3. Candidate rock points (CRPs) are points close (within 1 m) to these local peaks. CRPs could include points that belong to the terrain as well as rocks.
4. To eliminate terrain points from CRPs, statistics of image intensity are calculated. Points with image intensity less than a certain threshold are accepted as CRPs.
5. Segmentation of CRPs is performed based on proximity in image space. Since local peaks are selected within a range of 1 m, smaller rocks within that range are separated in this step. Each segmented group of CRPs is now called a rock points group (RPG).

6. If the number of CRPs in the RPG is less than ten, it is likely to be very small rock or partial shadow, so these smaller RPGs are eliminated.
7. Each RPG is modeled as a rock based on height and width. Height is defined as the elevation difference between highest CRP and lowest CRP. Width is the distance between leftmost CRP and rightmost CRP in the image space. The reason for using image space is that object space coordinates behind the visible side of the rock surface are less accurate. If the height is over 15 cm and the width is between 18 cm and 1 m, the RPG is selected as a rock.

Figure 4 shows the results of ground rock extraction from a Navcam image at site 13136. Among the 3D points, only those within a range of 3 ~ 25 m from the rover were selected (Figure 4a). Local peaks within the 1-m range became candidate rock points (Figure 4b). Points surrounding local peaks are also CRPs (Figure 4c). Figure 4d shows CRPs with darker brightness values. After segmenting the CRPs, rock points groups with less than 10 points are eliminated (Figure 4e). Finally, RPGs that satisfy the height and width constraints are selected as rocks (Figures 4f and 4g). These rocks from ground images are later matched with orbital rocks based on distribution patterns.

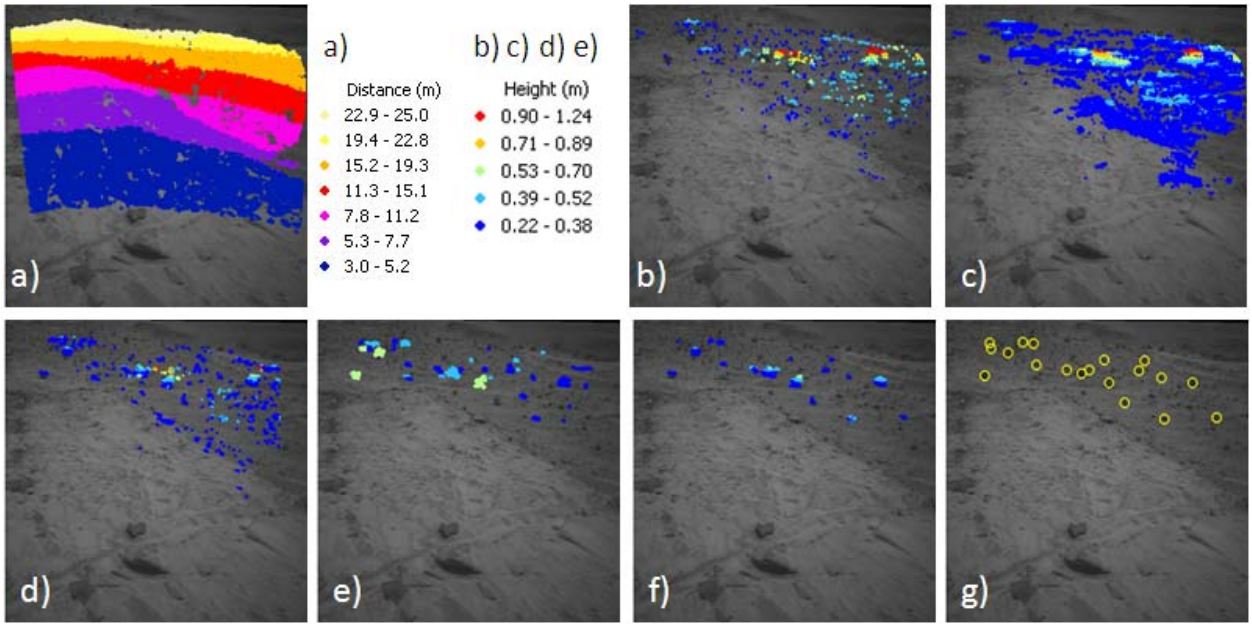

Figure 4. Rock extraction from site 13136: a) is the result from step 1, b) is from step 2, c) is from step 3, d) is from step 4, e) is from step 6, f) is from step 7, and g) shows the final selected ground rocks circled on the ground image.

Rock Extraction from Orbital Imagery

In orbital imagery, rocks are distinguished from the terrain mainly based on pixel brightness. Since the illumination of the Martian terrain can change based on the geology and the topography, locally dark pixels are identified by comparing pixel brightness values with those of the neighborhood average. For the orbital image I, the following procedures were used to extract orbital rocks.

1. Local average I_{ave} is the convolution of the average filter M and the original image I, so that $I_{ave} = M * I$.
2. Local darkness of a pixel is obtained by subtracting the average pixel brightness value of its neighboring pixels from the image intensity of the pixel. The map of local darkness is therefore $I_{dark} = I - I_{ave}$.
3. If the subtracted value of I_{dark} for a pixel (i, j) is less than $threshold_1$, the pixel belongs to the DarkPixel set:

$$DarkPixel = \{(i, j) | I_{dark}(i, j) < threshold_1\}. \quad (5a)$$

4. However, not all dark pixels are considered as rocks. Rock candidates should meet a certain criterion as a group of rock pixels. For that reason, the elements of the DarkPixel are then grouped based on proximity, forming connected components CC such that:

$$CC_k = \{(i_c, j_c) | (i_c, j_c) \in \text{DarkPixel}, i_c' - 1 \le i_c \le i_c' + 1 \text{ and } j_c' - 1 \le j_c \le j_c' + 1 \text{ for } (i_c', j_c') \in CC_k\} \quad (5b)$$

5. If a CC contains at least one pixel whose I_{dark} value is less than a $threshold_2$, it becomes a dark connected component (DarkCC). Reasons for not using a single threshold value for extracting rock pixels are twofold: 1) using a less strict $threshold_1$ in the first place eliminates disconnection CCs from dark pixels that actually belong to the same rock, and 2) by applying a stricter $threshold_2$ in a later step, a CC with a very dark pixel is distinguished as a rock:

$$DarkCC_k = \{(i_d, j_d) | (i_d, j_d) \in CC_k, \in I_{dark}(i_d, j_d) < threshold_2\} . \quad (5c)$$

6. The size and shape of a DarkCC are important factors to determining if it can be used as a tie point. By conducting space intersection with corresponding tie points from the other stereo image, object space coordinates (x_{ij}, y_{ij}, z_{ij}) are obtained for the image space coordinates (i, j).
7. For each DarkCC, major and minor axes are defined. Major axis V_{major} is the vector of longest distance in the DarkCC. Minor axis V_{minor} is orthogonal vector to the major axis. For positional accuracy, the maximum width of a DarkCC should not exceed 2 m in object space even considering the shadow. The ratio of the length of the minor axis to that of the major axis should be larger than 0.2. Using Equation 5d, those DarkCC meeting the criteria are selected as orbital rocks (OrbRock).

$$OrbRock_k = \{(i, j) | (i, j) \in DarkCC_k, (i', j') \in DarkCC_k, \mathbf{V} = (x_{ij} - x_{i'j'}, y_{ij} - y_{i'j'}), |\mathbf{V}_{major}| \ge |\mathbf{V}|, \mathbf{V}_{minor} \perp \mathbf{V}_{major}, |\mathbf{V}_{major}| \le 2, |\mathbf{V}_{minor}| / |\mathbf{V}_{major}| > 0.2\} \quad (5d)$$

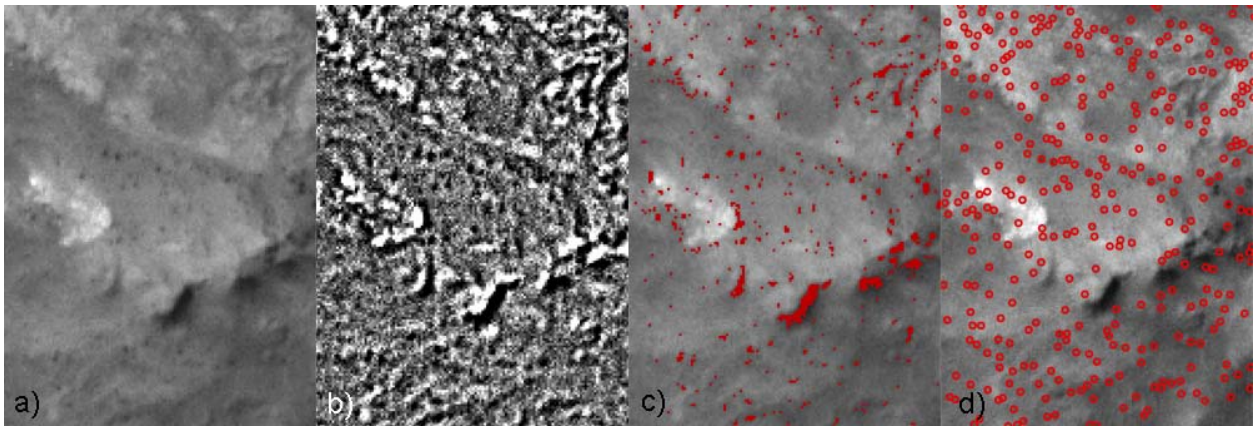

Figure 5. Rocks extracted from HiRISE orbital image of the southern part of Home Plate: a) original image I, b) local brightness I_{dark}, c) the DarkPixel (red pixels) overlaid on the original image, and d) the orbital rocks circled on the original image.

Figure 5 shows the result of orbital rock extraction from the image PSP_001777_1650. The area displayed is the southern part of Home Plate, including the Spirit rover site 13136. From the original image (Figure 5a), the local brightness is calculated (Figure 5b). The size of the average filter used for this example is 9 by 9 pixels. In Figure 5c, the red pixels belong to the DarkPixel set, which contains relatively dark pixels in the neighborhood. After segmenting the dark pixels into connected components based on their proximity and applying additional constraints about the size and brightness values of the members, the final rocks are obtained (Figure 5d).

Rock Pattern Matching

Once orbital and ground rocks are obtained and terrain matching is completed, rock pattern matching is conducted. It is assumed that the rover position adjusted by terrain matching is closer than a certain range to the real location. Depending on the effectiveness of the terrain match, the estimated distance tends to vary. If the terrain shows a distinctive shape such as a crater or a ridge, the accuracy is better than in the case of a flat terrain. Accordingly, pairs of orbital/ground rocks within the range are identified. These pairs are used to define translation vectors from estimated rover positions to the actual ones. MER rovers correct their attitudes by finding the sun. Therefore the rotation angles are accurate enough to be considered zero at this stage (Li et al., 2007). Therefore all the ground rocks

are shifted solely based on the translation vector from each orbital/ground rock pair. Then orbital/ground rock pairs that are within 0.5 m are selected as matched rocks. Based on these matched rock pairs, 2D affine transformation is derived as shown in Equation 1e. The translation vector with the maximum number of paired rocks is chosen. If there are multiple vectors with the same number of pairs, the mean residual between matched rocks after affine transformation should be the minimum. Finally, the rover position is adjusted using the 2D affine transformation with the maximum number of pairs and the minimum mean residuals.

Figure 6 shows the rock pattern matching results for site 13136. In Figure 6a, the rover position and the coordinates of 55 ground rocks resulting from terrain matching are shown. After implementing the rock pattern matching procedure, the optimal solution was derived (Figure 6b). In total, 10 ground rocks were matched with orbital rocks. When the extracted rocks are too close to each other, two or more ground rocks could be paired with the same orbital rock, and vice versa. In that case, the orbital/ground rock pair with the least residual is selected and the other ones, which are duplicates, are eliminated from the matched rock pairs.

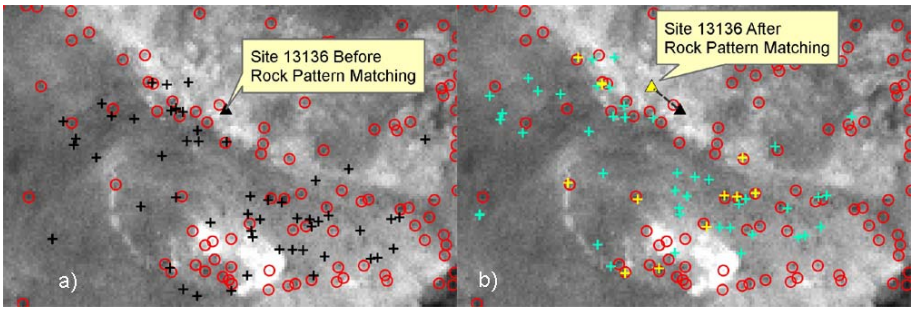

Figure 6. Ground rock distribution: a) after terrain matching (black crosses are ground rocks, black triangles are rover position), and b) after rock pattern matching (yellow crosses are matched ground rocks, blue crosses are unmatched ground rocks, and yellow triangles are adjusted rover position after 2D affine transformation).

EXPERIMENTAL RESULTS AND DISCUSSION

We generated topographic mapping products from a stereo pair of HiRISE images that cover the Columbia Hills area of the MER Spirit rover landing site PSP_001513_1655 (27.1 cm/pixel, 80,000 rows) and PSP_001777_1650 (26.3 cm/pixel, 40,000 rows). The extent of PSP_001777, about 10.5 × 5.5 km on the ground, is entirely covered by PSP_001513. The two images have a convergence angle of 19.8 degrees. The objective of this test was to examine the effectiveness of the rock matching results for locating the rover in the orbital topographic product. For that reason, the orbital products were obtained from the integrated image networks using MOLA altimetry data and MER traverse. For this test, seven sites were selected along the Spirit rover traverse (Figure 7). Terrain matching, orbital and ground rock

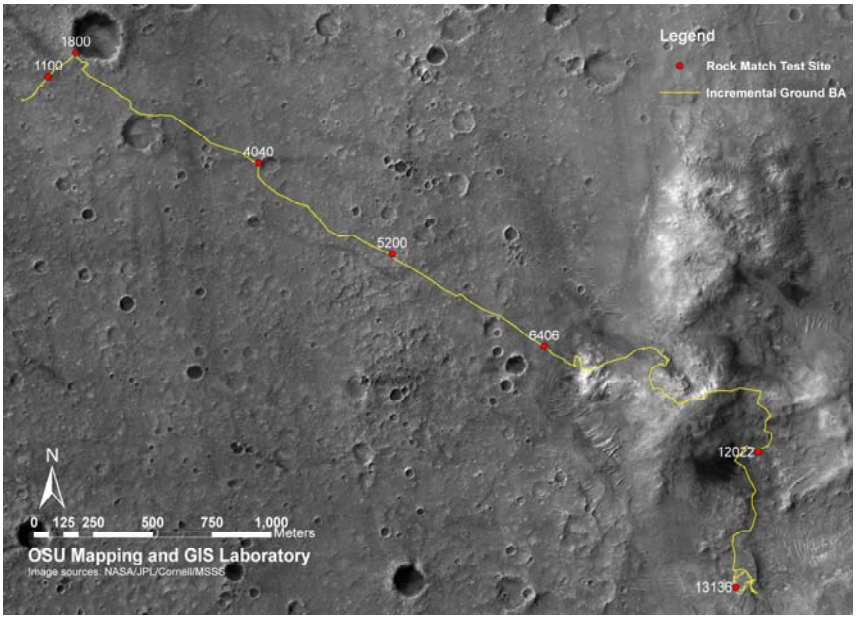

Figure 7. Rock matching test sites along the Spirit rover traverse (red dots).

extraction and rock pattern matching were conducted at each site. The adjusted rover position was compared with the rover position manually identified on HiRISE orthophotos. The current MER mission utilizes incremental BA processing using cross-site stereo matching.

Table 1 shows the rock matching results for the seven test sites. Telemetry data from the PDS was used as the initial rover position. Terrain matching was performed based on the DEMs derived from orbital and ground datasets, and the rover position was updated. The number of ground rocks extracted from Navcam stereos varied from 12 (site 5200) to 54 (site 13136) depending on the quantity of rocks in that area meeting the criteria of size, height and image intensity. After rock pattern matching was performed, the number of ground rocks matched with orbital rocks ranged from 5 to 13, depending on the availability of the corresponding orbital rocks. The ratio of matched ground rocks to extracted rocks ranged from 18.5 % up to 50.0 %.The wide range of this result is influenced by the fact that rocks seen from the ground are often not visible in the orbital imagery due to a lack of distinctive variation in image intensity when compared to the surrounding area. When compared to the manually identified rock pairs and rover position, the rock matching results for the seven test sites were found to be correct.

Table 1. Rock matching results from ground imagery

Site	1100	1800	4040	5200	6406	12022	13136
Distance from the origin	159.7 m	309.6 m	1299.5 m	2037.3 m	2802.4 m	4660.2 m	5832.7 m
Number of ground rocks	20	36	23	12	14	26	54
Number of matched ground rocks	10	13	5	6	6	8	10
Percent matched	50.0 %	36.1 %	21.7 %	50.0 %	42.9 %	30.8 %	18.5 %

In this paper, a new method for automated rover localization based on the integration of orbital and ground image networks is presented. This method consists of three major steps: (a) construction of an orbital image network from HiRISE stereo images based on the rigorous sensor model, (b) construction of a ground image network based on MER telemetry data and visual odometry, and (c) integration of these orbital and ground image networks using a newly developed methodology. The first step of this integration consists of using the MER traverse as horizontal control for the orbital image network. As the rover proceeds, its positions are adjusted using orbital/ground rock matching. The capability for automatic rover localization in the context of orbital image network is examined. Test results show that this method can effectively extract rocks and adjust rover positions comparable to the labor-intensive incremented BA currently used for the MER mission. This technique opens a new possibility for autonomous rover localization performed onboard unmanned vehicles for future space exploration. Further research is needed to combine incremental BA with the proposed orbital/ground data integration. Based on the performance of the rover localization method, it also will be important to develop measures of quality control for different scenarios.

ACKNOWLEDGEMENTS

Funding for this research by the NASA Mars Applied Information Systems Research Program is acknowledged.

REFERENCES

Archinal, B.A., R.L. Kirk, T.C. Duxbury, E.M. Lee, R. Sucharski, and D. Cook, 2003. Mars Digital Image Model (MDIM) 2.1 Control Network, *ISPRS Working Group IV/9 Workshop "Advances in Planetary Mapping 2003"*, Houston, March 2003.

Arvidson, R.E., R.C. Anderson, P. Bartlett, J.F. Bell III, D. Blaney, P.R. Christensen, P. Chu, L. Crumpler, K. Davis, B.L. Ehlmann, R. Fergason, M.P. Golombek, S. Gorevan, J.A. Grant, R. Greeley, E.A. Guinness, A.F.C. Haldermann, K. Herkenhoff, J. Johnson, G. Landis, R. Li, R. Lindermann, H. McSween, D.W. Ming, T. Myrick, L. Richter, F.P. Seelos IV, S.W. Squyres, R.J. Sullivan, A. Wang, and J. Wilson, 2004. Initial localization and physical properties experiments conducted by Spirit at Gusev Crater, *Science*, 305(5685): 821-824.

Gor, V., R. Castano, R. Manduchi, R.C. Anderson, and E. Mjolsness, 2001. Autonomous rock detection for Mars terrain, *AIAA Space 2001—Conference and Exposition*, Albuquerque, NM, August 2001.

Kirk, R.L., E. Howington-Kraus, B. Redding, D. Galuszka, T.M. Hare, B.A. Archinal, L.A. Soderblom, and J.M. Barrett, 2003. High-resolution topomapping of candidate MER landing sites with Mars orbiter camera narrow-angle images, *J. Geophys. Res.*, 108(E12): 8088, doi:10.1029/2003JE002131.

Kirk R.L., E. Howington-Kraus, M.R. Rosiek, J.A. Anderson, B.A. Archinal, K.J. Becker, D.A. Cook, D.M. Galuszka, P.E. Geissler, T.M. Hare, I.M. Holmberg, L.P. Keszthelyi, B.L. Redding, W.A. Delamere, D. Gallagher, J.D. Chapel, E.M. Eliason, R. King, and A.S. McEwen, 2008. Ultrahigh resolution topographic mapping of Mars with MRO HiRISE stereo images: Meter-scale slopes of candidate Phoenix landing sites, *J. Geophys. Re*s., *113*, E00A24, doi:10.1029/2007JE003000.

Li, R., K. Di, L.H. Matthies, R.E. Arvidson, W.M. Folkner, and B.A. Archinal, 2004. Rover localization and landing site mapping technology for 2003 Mars Exploration Rover mission, *Photogrammetric Engineering & Remote Sensing*, 70(1): 77-90.

Li, R., S.W. Squyres, R.E. Arvidson, B.A. Archinal, J. Bell, Y. Cheng, L. Crumpler, D.J. Des Marals, K. Di, T.A. Ely, M. Golombek, E. Graat, J. Grant, J. Gulnn, A. Johnson, R. Greeley, R.L. Kirk, M. Maimone, L.H. Matthies, M. Malin, T. Parker, M. Sims, L.A. Soderblom, S. Thompson, J. Wang, P. Whelley, and F. Xu, 2005. Initial results of rover localization and topographic mapping for the 2003 Mars Exploration Rover mission, *Photogrammetric Engineering & Remote Sensing*, 71(10):1129–1142.

Li, R., K. Di, A.B. Howard, L.H. Matthies, J. Wang, and S. Agarwal, 2007. Rock modeling and matching for autonomous long-range Mars rover localization, *Journal of Field Robotics: Special Issue on Space Robotics*, 24(3): 187-203.

Li, R., J.W. Hwangbo, Y. Chen, and K. Di, 2008. Rigorous Photogrammetric Processing of HiRISE Stereo Images for Mars Topographic Mapping, *XXI ISPRS Congress*, Beijing, China, July 2008.

Maki, J.N., J.F. Bell, K.E. Herkenhoff, S.W. Squyers, A. Kiely, M. Klimesh, M. Schwochert, T. Litwin, A. Johnson, M. Maimone, E. Baumgartner, A. Collins, M. Wadsworth, S.T. Elliot, A. Dingizian, D. Brown, E.C. Hagerott, L. Scherr, R. Deen, D. Alexander, and J. Lorre, 2003. Mars exploration rover engineering cameras, *J. Geophys. Res.*, 108(E12), 8071, doi:10.1029/2003JE002077.

Maimone, M.W., P.C. Leger, and J.J. Biesiadecki, 2007. Overview of the Mars exploration rovers' autonomous mobility and vision capabilities, *IEEE International Conference on Robotics and Automation (ICRA) Space Robotics Workshop*, Roma, Italy, 14 April 2007.

Matthies, L., E. Gat, R. Harrison, B. Wilcox, R. Volpe, and T. Litwin, 1995. Mars microrover navigation: performance evaluation and enhancement, *Autonomous Robots*, 2(4): 291-311.

McEwen, A.S., E.M. Eliason, J.W. Bergstrom, N. T. Bridges, C. J. Hansen, W. A. Delamere, J. A. Grant, V. C. Gulick, K. E. Herkenhoff, L. Keszthelyi, R. L. Kirk, M.T. Mellon, S. W. Squyres, N. Thomas, and C. M. Weitz, 2007. Mars reconnaissance orbiter's High Resolution Imaging Science Experiment (HiRISE), *J. Geophys. Res.*, *112*, E05S02, doi:10.1029/2005JE002605.

Olson, C.F., 2000. Probabilistic self-localization for mobile robots, *IEEE Transactions on Robotics and Automation*, 16(1): 55–66.

Scholten, F., K. Gwinner, T. Roatsch, K. Matz, M. Wählisch, B. Giese, J. Oberst, R. Jaumann, G. Neukum, and the HRSC Co-Investigator Team, 2005. Mars express HRSC data processing – Methods and operational aspects, *Photogrammetric Engineering & Remote Sensing*, 71(10): 1143–1152.

Shan, J., J-S.Yoon, S. Lee, R. Kirk, G. Neumann, and C. Acton, 2005. Photogrammetric analysis of the Mars global surveyor mapping data, *Photogrammetric Engineering & Remote Sensing*, 71(1): 97-108.

Smith, D. E., M.T. Zuber, S.C. Solomon, R.J. Phillips, J.W. Head, J.B. Garvin, W.B. Banerdt, D.O. Muhleman, G.H. Pettengill, G.A. Neumann, F.G. Lemoine, J.B. Abshire, O. Aharonson, C.D. Brown, S.A. Hauck, A.B. Ivanov, P.J. McGovern, H.J. Zwally, and T.C. Duxbury, 1999. The global topography of Mars and implications for surface evolution, *Science*, 284(28): 1495-1503.

Thompson, D., S. Niekum, T. Smith, and D. Wettergreen, 2005. Automatic detection and classification of geological features of interest, *Proceedings of the IEEE Aerospace Conference*, Big Sky Montana, March, 2005.

Wettergreen, D., N. Cabrol, J. Teza, P. Tompkins, C. Urmson, V. Verma, M.D. Wagner, and W.L. Whittaker, 2005. First experiments in the robotic investigation of life in the Atacama Desert of Chile, *IEEE International Conference on Robotics and Automation*, Barcelona, Spain, April 2005.

Xu, F., 2004. Mapping and localization for extraterrestrial robotic explorations, Ph.D. Dissertation, The Ohio State University, Columbus, OH.