Visual Terrain Mapping for Mars Exploration[1,2]

Clark F. Olson
Computing and Software Systems
University of Washington, Bothell
18115 Campus Way NE, Box 358534
Bothell, WA 98011-8246
+1-425-352-5288
cfolson@u.washington.edu

Larry H. Matthies, John R. Wright
Jet Propulsion Laboratory
California Institute of Technology
4800 Oak Grove Drive
Pasadena, CA 91109-8099
lhm@robotics.jpl.nasa.gov
john@tone.jpl.nasa.gov

Ron Li, Kaichang Di
The Ohio State University
Department of Civil and Environmental Engineering and Geodetic Science
2070 Neil Avenue, 470 Hitchcock Hall
Columbus, OH 43210-1275
li.282@osu.edu, di.2@osu.edu

Abstract—One goal for future Mars missions is to navigate a rover to science targets not visible to the rover, but seen in orbital or descent images. In order to support and improve long-range navigation capabilities, we generate 3D terrain maps using all available images, including surface images from the lander and/or rover, descent images from the lander, and orbital images from current and future Mars orbiters. The techniques used include wide-baseline stereo mapping for terrain distant from the rover, bundle adjustment for high-accuracy mapping of surface images, and structure-from-motion techniques for mapping using descent and orbital images. The terrain maps are compiled using a system for unifying multi-resolution models and integrating three-dimensional terrains.

TABLE OF CONTENTS

1. INTRODUCTION

For a Mars rover capable of long-range mobility, it is highly desirable to travel to science targets observed in orbital or descent images. However, current rover technologies do not allow rovers to autonomously navigate to distant targets with a single command. Since communication with Mars rovers usually occurs only once per day, navigation errors can result in the loss of an entire day of scientific activity.

In order to improve long-range navigation capabilities, we generate 3D terrain maps using all available images, including surface images from the lander and/or rover, descent images from the lander, and orbital images from current and future Mars orbiters. Laser rangefinder data from rovers and/or landers can also be integrated with this methodology.

We apply bundle adjustment techniques to overlapping stereo images from panoramic image sets in order to create highly accurate maps of terrain nearby the rover. This method automatically determines corresponding features between multiple pairs of stereo images, even in cases where the overlap is small. These correspondences are used to update the camera positions precisely. Finally, seamless maps are constructed from the stereo data.

Distant terrain is mapped using a combination of rover, lander, and orbital images. Onboard the rover, maps of distant terrain can be created using wide-baseline stereo vision. In addition, we create maps using images collected as a lander descends to the surface or from orbital images.

While conventional stereo vision performed on the rover has limited accuracy for distant terrain owing to the small distance between the stereo cameras (known as the baseline distance), we can achieve accurate mapping for distant terrain using wide-baseline stereo vision. In this technique, images from the same camera, but at different rover positions, are used to generate a virtual pair of stereo images with a large baseline distance. In this work, we overcome two significant problems. First, the relative positioning between the cameras is not well known, unlike conventional stereo vision, where the cameras can be carefully calibrated. In addition, the problem of determining the matching locations between the two images is more difficult owing to the different viewpoints at which the images are captured.

Images captured during the descent of a lander to the surface can be mapped using similar techniques. This data faces the additional difficulty that the direction of movement is towards the terrain being imaged, complicating image rectification. We determine the terrain height by resampling one image several times according possible terrain heights

[1] 0-7803-8155-6/04/$17.00© 2004 IEEE
[2] IEEEAC paper #1176, Version 2, Updated December 3, 2003

and selecting the height at which the best match occurs for each image location.

Techniques similar to wide-baseline stereo vision are applied to the creation of terrain maps using correspondences between descent images and orbital images or between orbital images from different sensors. In this case, the different sensors have different responses to various terrain features. For this reason, we use a representation of the images that is based on local entropy to perform matching.

The terrain maps computed from the various images are compiled using SUMMITT (System for Unifying Multi-resolution Models and Integrating Three-dimensional Terrains), which is designed to merge disparate data sets from multiple missions into a single multi-resolution data set.

A considerable amount of previous work has been done on robotic mapping. However, much of it concerns indoor robots, while we are concerned with mapping natural terrain with rovers and spacecraft. We also concentrate on the use of cameras for mapping, but other sensors have also been used for mapping Mars, including laser altimeter [1] and delay-Doppler radar [2]. Aside from our own work, much of the early and recent work on terrain mapping for rovers has been performed at CMU [3, 4, 5, 6]. Also of note is work at CNRS, where stereo mapping is performed using an autonomous blimp [7, 8].

In Section 2, we discuss the use of wide-baseline stereo mapping for terrain that cannot be mapped effectively using conventional stereo vision. Section 3 describes our approach to mapping surface terrain close to the rover using locally captured stereo panoramas. Methods by which images captured during a spacecraft descent to the planetary surface are given in Section 4. Section 5 discusses techniques by which multi-model image pairs (such as from different orbiters or orbital/descent image pairs) can be used to create three-dimensional maps. The SUMMITT system for integrating multi-resolution data sets is described in Section 6. We give our conclusions in Section 7.

2. WIDE-BASELINE STEREO VISION

The accuracy of stereo vision decreases with the square of the distance to the terrain position. For this reason, conventional stereo vision cannot accurate map terrain many meters away. One solution to this problem is to use a larger baseline distance (the distance between the cameras), since the depth accuracy is inversely proportional to this distance. This is problematic, however, since a rover with limited size cannot have two cameras with a large baseline distance.

We achieve an arbitrarily large baseline distance using two images captured by the rover at different positions. While this improves the accuracy of the range estimation, it introduces two new problems. First, stereo systems typically use a pair of cameras that are carefully calibrated so that the relative position and orientation between the cameras are known to high precision. This is not possible with wide-baseline stereo, owing to rover odometry errors. Second, the change in the viewpoint between the images makes stereo matching more difficult, since the terrain no longer has the same appearance in both images. Our algorithm addresses these problems using a motion refinement step based on the structure-from-motion problem of computer vision [9] and robust matching between the images [10].

Motion Refinement

We assume that an initial estimate of the motion between the positions at which the images were captured is available from the rover odometry (or other sensors). We refine this estimate by determining matching points between the images and updating the motion to enforce geometrical constraints that must be satisfied if the points are in correspondence.

Once the corresponding matches have been found (using, for example, normalized correlation), we seek precise values for the translation T and rotation R relating the positions of the rover camera at the two locations. In this procedure, we include in the state vector not only the six parameters describing the relative camera positions (only five are recoverable, since the problem can be scaled to an arbitrary size), but also the depth estimates of the features for which we have found correspondences. With this augmented state vector, the objective function that we minimize becomes the sum of the squared distances between the detected feature position in one image and the estimated feature position calculated using backprojection from the other image according to the current motion estimate. We use the Levenberg-Marquardt optimization technique [11] to adjust the motion parameters in order to minimize this function.

After we have determined the new motion estimate, we apply a rectification process that forces corresponding points between the images to lie on the same row in the images [12]. Figure 1 shows an example of matching features that were detected in an image pair with a baseline distance of 20 meters. Figure 2 shows the images after motion refinement and rectification has been performed. In this example, the camera pointing angles converged by 20 degrees from perpendicular so that the same rock was located at the center of both images.

It can be observed that there is relatively little movement between the features matched in the middle of the image. Features on the mountains move considerably to the right and foreground features move in the opposite direction. Recent work allows robust matching to performed in these cases also. After rectification, all of the corresponding features lie on the same image row, facilitating dense matching.

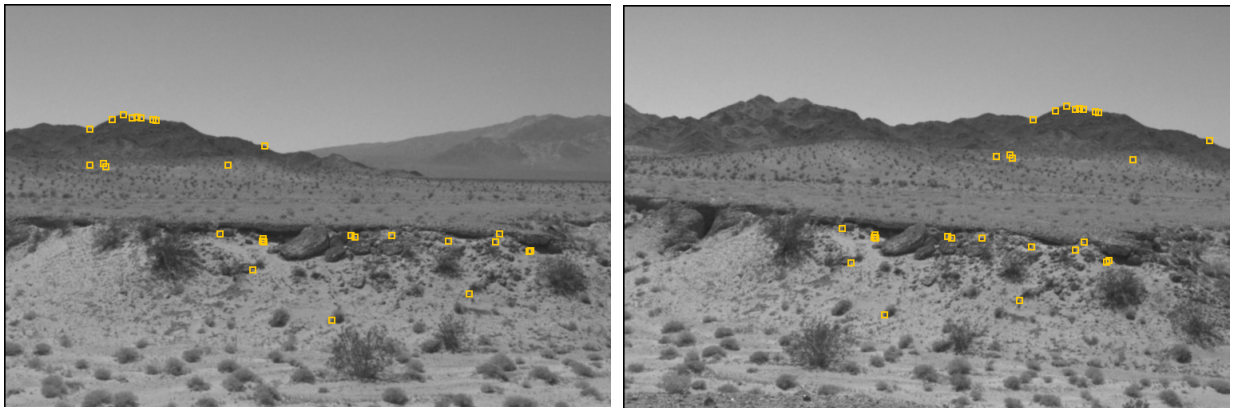

Figure 1: Matching features detected in a wide-baseline stereo pair.

Figure 2: Wide-baseline stereo pair after rectification. Corresponding features now line on the same image row.

Disparity Estimation

Image disparity is a measure of the difference in x-coordinate for corresponding points between two images. It has an inverse relationship to the depth of the terrain point represented by the corresponding image points. Given a pair of rectified images, we can compute the disparity for any point that is present in both images by searching along the corresponding row of the other image. We combine robust template matching [10] with an efficient stereo search algorithm in order to compute a dense disparity map. Every position in one image is given a disparity estimate, unless no corresponding match can be found in the other image.

In order to determine which points are in correspondence between the two images, we could use a measure such as normalized correlation or the sum-of-squared differences (SSD) applied to a small window (or box) around each point. However, owing to the difference in viewpoints, these measures do not produce good results [9]. Instead, we use a maximum-likelihood measurement that improves upon normalized correlation (and SSD) in two ways. First, normalized correlation compares only the pixels between the two images that are directly overlapping at some disparity of the image window with respect to the other image. If camera motion or perspective distortion causes pixels to move by different amounts between the two images, it will not be possible to find a window position where all of the pixels are correctly overlapped. Our distance measure allows pixels that are not directly overlapping to be matching by linearly combining the distance in the image with the difference in intensity. Computing the best distance for a pixel is no longer trivial with the formulation, since the best match may not be the overlapping pixel from the other image. However, efficient computation of the distances can be performed by precomputing a three-dimensional distance transform of the input data [10].

The second improvement over normalized correlation is that the possibility of an outlier is explicitly represented. In this application, any terrain feature that is visible in one image, but not in the other is an outlier for the matching process. Such outliers occur frequently for images taken from different viewpoints. In order to model such outliers, we use a probability density function for each pixel distance that has two terms, one for inliers and one for outliers, where each is weighted by an estimate of the probability of an outlier.

In order to perform dense matching between the rectified images using the measure described above, we use an efficient search strategy common in stereo vision. This strategy makes use of the observation that a brute-force implementation performs many redundant computations for adjacent positions of the template at the same displacement. We eliminate the redundant computation by storing the information for reuse as necessary for fast matching. Figure

3 shows a result of performing dense matching using the images from Figures 1 and 2. The results are displayed on the positions at which the points lie in the right image. High quality results are achieved on the right side of the image, since this is the area in which the terrain features are also present in the left image. Towards the left side (and bottom) of the image, the results degrade, since there is no matching feature (or the match is difficult to find).

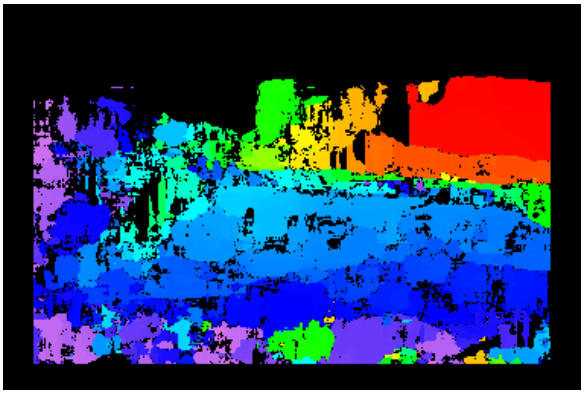

Figure 3: Color-coded disparity map computed from wide-baseline stereo pair. The largest (positive) disparities are red, while the smallest (negative) disparities are purple. Black values indicate that no match was found.

3. Mapping Surface Panoramas

We map surface panoramas using an automatic procedure that first selects candidate tie points in stereo images. Matches are detected both within stereo pairs and between adjacent stereo pairs to create an image network. Bundle adjustment is applied to the image network in order to improve the estimates of the camera and landmark positions. Finally, elevations maps and orthophotos are generated using dense matches detected between image pairs.

Automatic Selection of Tie Points

We have developed a systematic method for automatic tie point selection [13, 14, 15]. The procedure for selecting tie points within one stereo pair (intra-stereo tie points) includes interest point extraction using the Förstner operator, interest point matching using normalized cross-correlation coefficients, verification based on the consistency of parallaxes, and final selection by gridding. Figure 4 shows an example of automatically selected intra-stereo tie points from IMP (Imager for Mars Pathfinder) images. Tie points between adjacent stereo images (inter-stereo tie points) are extracted and selected with the help of a coarse elevation map generated from individual stereo pairs using the approximate orientation parameters. Figure 5 shows an example of automatically selected inter-stereo tie points from IMP images. Figure 6 shows an example of tie points automatically selected from rover images. Ultimately, the selected intra- and inter-stereo tie points build an image network.

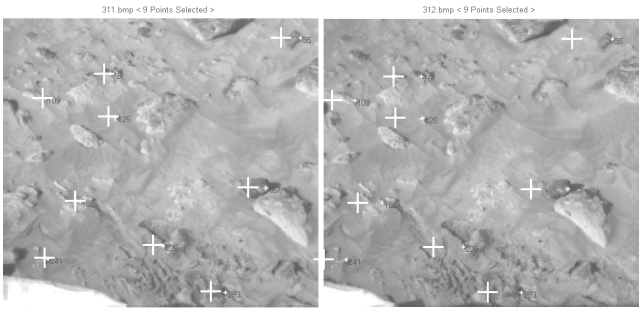

Figure 4: Intra-stereo tie points in lander IMP (Imager for Mars Pathfinder) images.

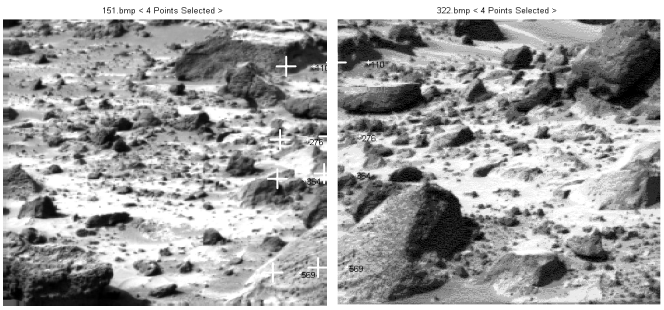

Figure 5: Inter-stereo tie points in lander IMP (Imager for Mars Pathfinder) images.

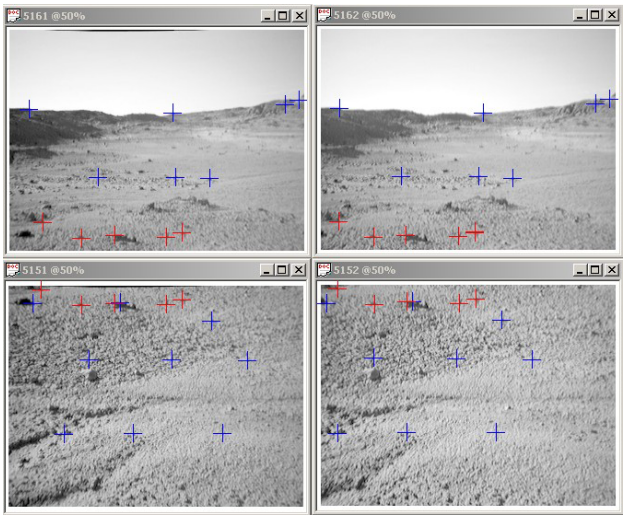

Figure 6: Automatically selected tie points from rover images (Blue crosses are intra-stereo tie points and red crosses are inter-stereo tie points.)

Bundle Adjustment

A bundle adjustment is applied to the image network to improve the accuracy of image orientation parameters as well as the 3D ground positions of the tie points. To ensure a high accuracy, we model the correlation between the position and attitude of the stereo camera and use this correlation as constraints in the least squares adjustment.

In an experiment on IMP data, the entire panorama consists of 129 images that form either an upper panorama and a

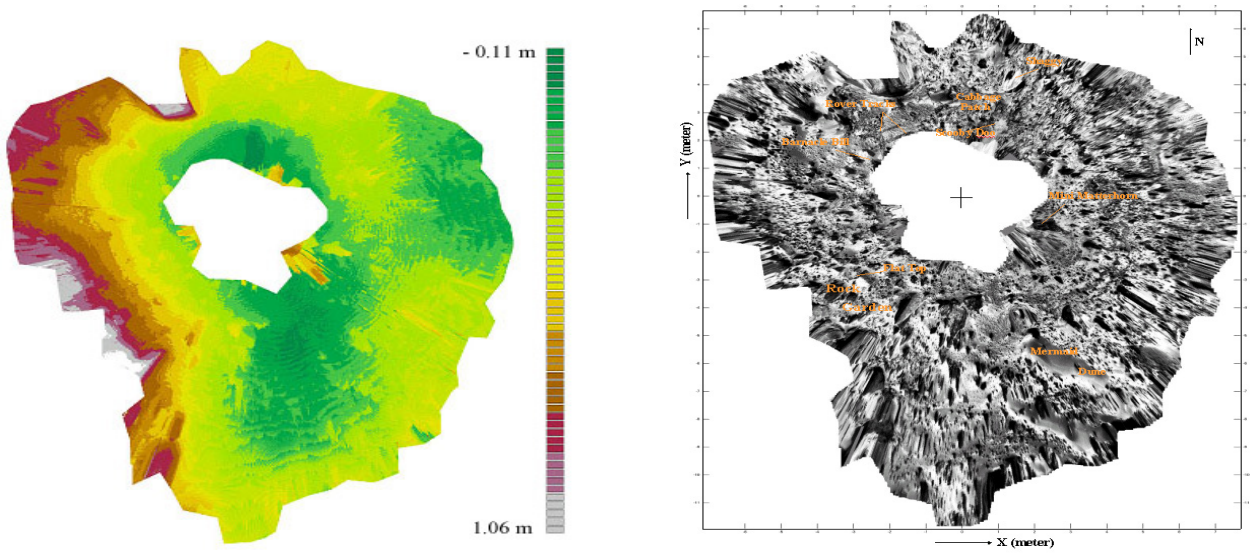

Figure 7: Elevation map and orthophoto of the Mars Pathfinder landing site.

lower panorama with horizontal links, or an entire panorama with both horizontal and vertical links. In the image network, there are 655 tie points, 633 of which are automatically selected and 22 that have been manually selected. Before adjustment, the precision is 4.61 pixels in the image space (distance between measured and back-projected image points) and 0.067 meter in the object space (distance between 3D positions triangulated from adjacent stereo pairs). After bundle adjustment, the precision is 0.80 pixels in the image space and 0.05 meter in the object space.

In an experiment processing terrestrial rover data, a panoramic image network was built by linking 36 mast images (18 pairs) with 220 automatically selected intra- and inter-stereo tie points. Before adjustment, the precision is 2.61 pixels in the image space and 1.62 meter in the object space. After bundle adjustment, the precision is 0.55 pixels in the image space and 0.08 meter in the object space. We can see that precision is improved in both the image space and the object space by the bundle adjustment. The improvement in image space (generally at a subpixel level) is larger than the improvement in the object space.

Elevation Map and Orthophoto Generation

After bundle adjustment, image matching is performed to find dense conjugate points for elevation map generation. The epipolar geometry and a coarse-to-fine matching strategy are applied to achieve both high speed and high reliability. The 3D ground coordinates of the matched points are then calculated by photogrammetric triangulation using the adjusted image orientation parameters. Finally, based on the 3D points, the elevation is generated using the Krigging interpolation method.

The orthophoto is generated by back-projecting the grid points of the refined elevation map onto the left (or right) image. A corresponding grayscale value is found and assigned to the grid point. In the area of overlap for adjacent images, the grid points will be projected to two or more images. The grayscale value is picked from the image in which the projected point is closest to its image center. The resultant elevation map and orthophoto of the Mars Pathfinder landing site are shown in Figure 7. The 3D visualization of the rover data is shown in Figure 8. Through visual checking of the orthophotos, we find no seams between image patches. This demonstrates the effectiveness of the bundle adjustment.

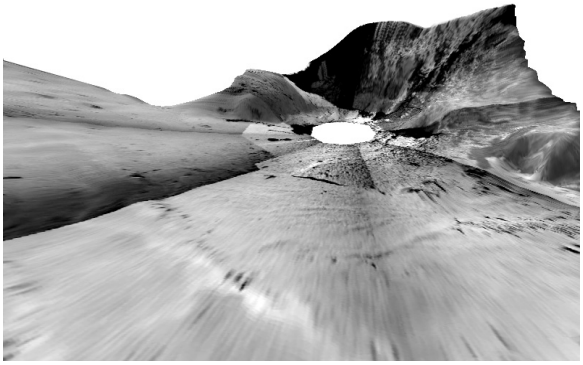

Figure 8: 3D visualization of terrestrial rover test site.

4. MAPPING DESCENT IMAGES

In addition to images captured by a rover or lander on the planetary surface, we can perform mapping using images captured above the surface. This is not limited to orbital images; it can also include the sequence of images captured by a spacecraft as it descends to the surface. An example of descent imagery of the moon from the Ranger 9 mission can be seen in Figure 9. Figure 10 shows a descent image sequence captured using a helicopter in the Mojave Desert for rover testing. Such descent images provide a link between the lower resolution images captured from orbit and the higher resolution images captured on the surface.

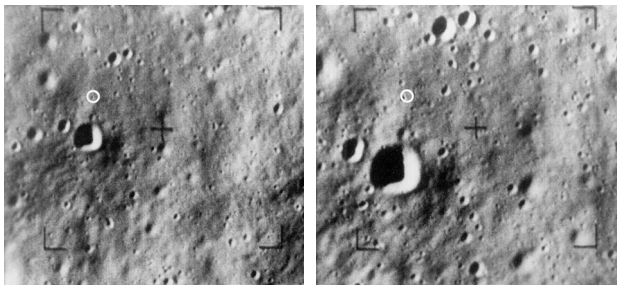

Figure 9: Images of the moon captured by the Ranger 9 spacecraft during descent to the surface.

Figure 10: Three images from a descent sequence of the Mojave Desert captured from a helicopter.

For matching between pairs of consecutive descent images, we can use techniques similar to those described for wide-baseline stereo, including motion refinement and dense disparity estimation [16]. For this case, the use of robust matching techniques for disparity estimation is less important, but a different problem arises. The images are captured in the direction in which the spacecraft is traveling (down towards the surface). In this case, the focus-of-expansion (the fixed point at which a terrain location appears at the same location in both images) is within the image boundaries. (For conventional stereo imaging, this point is at infinity.) This has two consequences. First, we cannot rectify the images along epipolar lines, since they pass through the focus-of-expansion, and this would resample the image highly unevenly. In addition, the depth cannot be accurately recovered near this point, since the computation is numerically unstable.

For the above reasons, our approach to recovering an elevation map from this data is somewhat different from wide-baseline stereo vision. We examine a discrete set of virtual planar surfaces perpendicular to the camera pointing axis that slice through the three-dimensional terrain. For each of these surfaces, one image is warped to appear as it would from the second camera position if all of the terrain was at this elevation. The warped image is compared to the second image at each position and the elevation the produces the closest match at each position is stored as the elevation estimate for that position. Increased accuracy is obtained through parabolic fitting of the scores achieved near the peak for each location.

Figure 11 shows an example of a map created from the descent images in Figure 10. The map was generated by rendering an image according to the elevations computed by the algorithm. A channel can clearly be seen running through the bottom of the image. However, the location of the focus-of-expansion can be seen in the center where the elevation estimates become somewhat unstable.

Figure 11: Map extracted from Mojave Desert images.

5. MULTI-MODAL DATA

After considering surface images and descent images, one might wonder whether similar techniques might be used for orbital images. In general, the answer is yes. We can even generate depth maps by performing matching between a descent image and an orbital image. If the same terrain location is imaged by the same sensor from a different location, then mapping can be performed using one of the previously described techniques. (If it is not imaged from a different location, we cannot extract any depth information.) On the other hand, if we want to perform matching between images captured from different sensors, then the problem becomes more complex, since the terrain will not necessarily have the same appearance in images from different sensors. Matching between images captured with different sensors is often called multi-modal matching.

When we perform matching between images captured with different sensors, we must use a measure for matching small image windows that is insensitive to the difference between the sensors, since different sensors have different responses to the various wavelengths of light. Our approach is to transform the images into a new representation that is less sensitive to sensor characteristics. The new representation replaces each pixel in the image with a local measure of the entropy in the image near the pixel [17]. This is based on the idea that each neighborhood in the image can be treated as a set of samples from a random variable. The entropy of the random variable is estimated from the sample values. Figure 12 shows an image and the same image after each pixel has been replaced with the local entropy computed using this method.

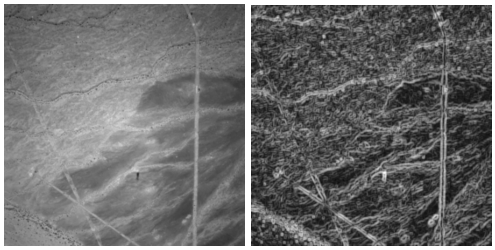

Figure 12: An aerial image and its representation using a local entropy measure.

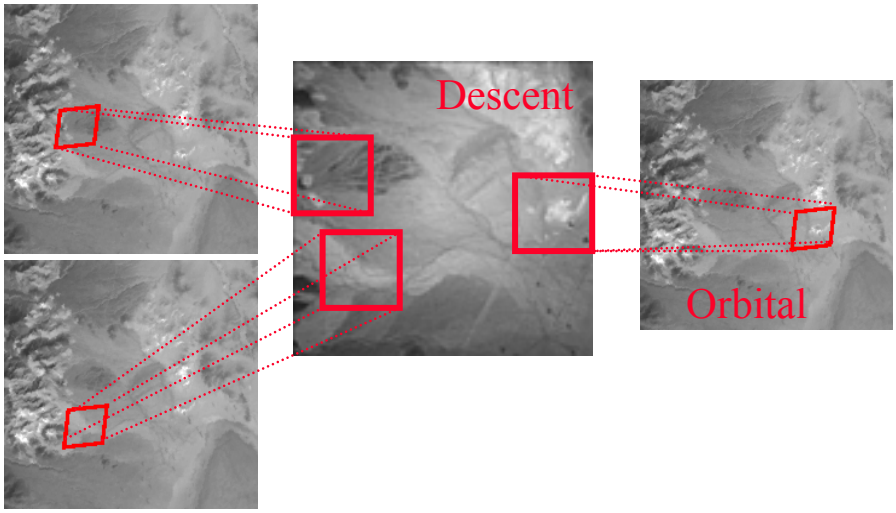

Figure 13: Matches found between an orbital image (left and right sides) and a descent image (center).

This entropy representation is very insensitive to sensor characteristics. In fact, it has been used to perform matching between infrared images and conventional images that have a very different appearance [17]. Once the images have been transformed into the entropy representation, the images can be compared in a straightforward manner using normalized correlation. An example of matches detected between an aerial image and an orbital image can be seen in Figure 13.

6. Integrating Data Sets

In this work, we have used SUMMITT for integrating maps from various sensors. SUMMITT (the System for Unifying Multi-resolution Models and Integrating Three-dimensional Terrains) is a suite of software tools for registering and merging terrain models from a variety of sources. The terrain models are then output in a variety of formats to support activities requiring visualization of the terrain.

The fundamental data structure for the terrain data is the octree. Octrees were chosen due to their inherent support of multi-resolution data, their ability to support rapid searches for nearest neighbors, and their small memory footprint. Multi-resolution support is necessary due to the disparate sources of terrain information that must be merged. Orbital cameras may be used to collect imagery and produce terrain models through stereo processing. In this case, the resolution of the terrain models could range from one meter up to multiple kilometers based on the quality and resolution of the optics, the orbital dynamics, and atmospheric constraints. Typically the resolution of terrain models produced from orbital cameras will be relatively constant as the range to the terrain is generally fixed by roughly circular orbits. Descent cameras, on the other hand, are constantly moving during data collection and their range to the terrain is varied. Thus, a camera in the early stages of data collection, at a higher altitude, might be capable of producing models with a resolution near one meter. The same camera might produce models at one centimeter resolution at a very low altitude. Finally, rover and lander cameras may produce models with sub-centimeter resolution. For example, on the Mars Exploration Rovers (MER) mission, the rovers carry three sets of stereo cameras ranging from nearly 180 degrees FOV down to about 15 degrees FOV. Given that the stereo baseline for each pair of cameras is comparable, the resolution of the terrain models produced will vary considerably among the cameras on a given platform. In addition, the models produced by a single camera will vary considerably in resolution due to the wide range of distances of the terrain from the camera. Thus, near-field objects may have up to 30 times finer resolution than far-field objects.

The terrain models are inserted into the octree by assuming that each sample is a volume with a cross-section equivalent to the model resolution. The coarser model data remains near the root of the tree while the finer data traverses nearer the leaves of the octree.

Each individual model, except the first, must be registered to the overall terrain model contained within the octree, prior to being merged into the octree. The registration process uses the Iterative Closest Points (ICP) algorithm, as described by Zhang [18] and applied by Nishino and Katsushi [19] and Miyazaki et al. [20], to compute an alignment transform. This transform is then applied to each sample in the model prior to inserting it into the octree. Since the ICP algorithm is an iterative process that requires finding nearest neighbors, the octree's support for rapid searches is very important to the success of the ICP algorithm.

Once models have been merged into the octree, terrain models can be generated. For many applications, multi-resolution triangle meshes are used. By treating the octree as a cloud of points, a variety of algorithms are available for this process, including those of Hoppe et al. [21] and Hilton et al. [22]. For rover missions, corresponding height maps are also required. The height maps are produced by binning

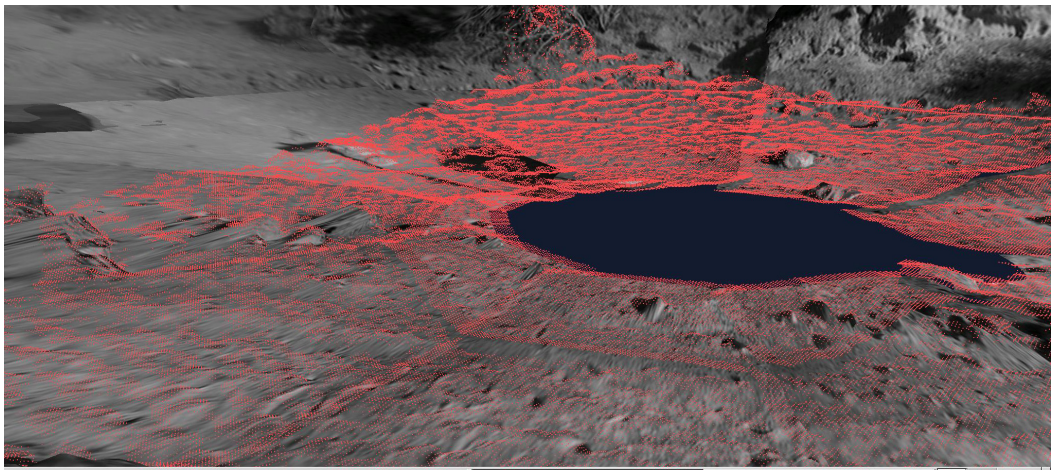

Figure 14: Elevation map generated from registered rover panorama data.

the octree data in the (x, y) plane and then selecting the largest z value within each bin. The bins can be of any desired resolution to support the visualization requirements. Repeated extraction of height maps at different resolutions produces multi-resolution maps for applications requiring such support.

Figure 14 shows an example of an elevation map generated using SUMMITT. In this case, panoramic stereo images were registered, integrated and, finally, rendered with the original image data.

7. CONCLUSIONS

Obtaining accurate maps of the terrain is critical for long-distance traverses on Mars. Without such maps, a rover may spend much time and energy venturing along what turns out to be a dead end. Maps are also critical for rover localization in long-distance traverses, since estimates from odometry and other sensors will grow in error without bound unless corrected using a global map.

We have described techniques for generating three-dimensional terrain maps spanning many resolutions, from the high-resolution maps generated using conventional stereo on-board a rover, through medium resolution maps generated using wide-baseline stereo or descent images, to the lower resolution (but wider area) maps generated from high-altitude descent images and/or orbital images.

The data sets are combined using a system (SUMMITT) that efficiently registers and integrates the data sets in a multi-resolution context. SUMMITT also provides tools for map visualization that are useful for planning.

8. ACKNOWLEDGEMENTS

We gratefully acknowledge funding of this work by the NASA Mars Technology Program.

9. REFERENCES

[1] M. T. Zuber et al., "Observations of the North Polar Region of Mars from the Mars Orbiter Laser Altimeter," *Science*, 283(5396): 2053-2060, December, 1998.

[2] J. K Harmon, R.E. Arvidson, E. A. Guinness, B. A. Campbell, M. A. Slade, "Mars Mapping with Delay-Doppler Radar," *Journal of Geophysical Research*, 104(E6): 14065-14089, 1999.

[3] M. Hebert, C. Caillas, E. Krotkov, I. S. Kweon, T. Kanade, "Terrain Mapping for a Roving Planetary Explorer," *Proceedings of the International Conference on Robotics and Automation*, vol. 2, pages 997-1002, 1989.

[4] I. S. Kweon, T. Kanada, "High-Resolution Terrain Map from Multiple Sensor Data," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 4(2): 278-292, February, 1992.

[5] E. Krotkov, "Terrain Mapping for a Walking Planetary Rover," *IEEE Transactions on Robotics and Automation*, 10(6): 728-739, December 1994.

[6] D. Huber, O. Carmichael, M. Hebert, "3-D Map Reconstruction from Range Data," *Proceedings of the IEEE International Conference on Robotics and Automation*, pages 891-897, 2000.

[7] S. Lacroix, I.-K. Jung, A. Mallet, "Digital Elevation Map Building from Low Altitude Stereo Imagery," *Robotics and Autonomous Systems*, 41(2-3): 119-127, November 2002.

[8] I.-K. Jung and S. Lacroix, "High Resolution Terrain Mapping Using Low Altitude Aerial Stereo Imagery," *Proceedings of the International Conference on Computer Vision*, pages 946-951, 2003.

[9] C. F. Olson, H. Abi-Rached, M. Ye, J. P. Hendrich, "Wide-Baseline Stereo Vision for Mars Rovers," *Proceedings*

of the IEEE/RSJ International Conference on Intelligent Robots and Systems, pages 1302-1307, 2003.

[10] C. F. Olson, "Maximum-Likelihood Image Matching," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 24(6): 853-857, June 2002.

[11] W. H. Press, S. A. Teukolsky, W. T. Vetterling, B. P. Plannery, *Numerical Recipes in C*, Cambridge University Press, 1988.

[12] A. Fusiello, E. Trucco, A. Verri, "A Compact Algorithm for Rectification of Stereo Pairs," *Machine Vision and Applications*, 12: 16-22, 2000.

[13] K. Di, R. Li, F. Xu, L. H. Matthies, C. F. Olson, "High Precision Landing Site Mapping and Rover Localization by Integrated Bundle Adjustment of MPF Surface Images," *International Archives of Photogrammetry and Remote Sensing*, vol. 34, part 4, pages 733-737, 2002.

[14] R. Li, K. Di, F. Xu, "Automatic Mars Landing Site Mapping Using Surface-Based Images," *ISPRS WB IV/9: Extraterrestrial Mapping Workshop – Advances in Planetary Mapping,* 2003.

[15] F. Xu, K. Di, R. Li, L. Matthies, C. F. Olson, "Automatic Feature Registration and DEM Generation for Martian Surface Mapping," *International Archives of Photogrammetry and Remote Sensing*, vol. 34, part 4, pages 549-554, 2002.

[16] Y. Xiong, C. F. Olson, L. H. Matthies, "Computing Depth Maps from Descent Imagery," *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, vol. 1, pages 392-397, 2001.

[17] C. F. Olson, "Image Registration by Aligning Entropies," *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, vol. 2, pages 331-336, 2001.

[18] Z. Zhang, "Iterative Point Matching for Registration of Free-Form Curves and Surfaces," *International Journal of Computer Vision*, 13(2): 119-152, 1994.

[19] K. Nishino, I. Katsushi, "Robust Simultaneous Registration of Multiple Range Images," *Proceedings of the 5th Asian Conference on Computer Vision*, pages 454-461, 2002.

[20] D. Miyazaki, T. Ooishi, T. Nishikawa, R. Sagawa, K. Nishino, T. Tomomatsu, Y. Takase, K. Ikeuchi, "The Great Buddha Project: Modelling Cultural Heritage through Observation," *Proceedings of the Sixth International Conference on Virtual Systems and Multimedia*, pages 138-145, 2000.

[21] H. Hoppe, T. DeRose, T. Duchamp, J McDonald, W. Stuetzle, "Mesh Optimization," *Proceedings of SIGGRAPH '93*, pages 19-26, 1993.

[22] A. Hilton, A. J. Stoddard, J. Illingworth, T. Windeatt, "Marching Triangles: Range Image Fusion for Complex Object Modelling," *Proceedings of the International Conference on Image Processing*, pages 381-384, 1996.

10. BIOGRAPHIES

Clark F. Olson received the B.S. degree in computer engineering in 1989 and the M.S. degree in electrical engineering in 1990, both from the University of Washington, Seattle. He received the Ph.D. degree in computer science in 1994 from the University of California, Berkeley. After spending two years doing research at Cornell University, he moved to the Jet Propulsion Laboratory, where he spent five years working on computer vision techniques for Mars rovers and other applications. Dr. Olson joined the faculty at the University of Washington, Bothell in 2001. His research interests include computer vision and mobile robotics. He teaches classes on the mathematical principles of computing, database systems, and computer vision and he continues to work with NASA/JPL on mapping and localization techniques for Mars rovers.

Larry H. Matthies PhD, computer science, Carnegie Mellon University, 1989; Supervisor, Machine Vision Group, Jet Propulsion Laboratory (JPL). Dr. Matthies has 21 years experience in developing perception systems for autonomous navigation of robotic ground and air vehicles. He pioneered the development of real-time stereo vision algorithms for range imaging and accurate visual odometry in the 1980's and early 1990's. These algorithms will be used for obstacle detection onboard the Mars Exploration Rovers (MER). The GESTALT system for obstacle avoidance for the MER rovers was developed in his group; this system is also currently the baseline for onboard obstacle detection for the 2009 Mars Science Laboratory (MSL) mission. His stereo vision-based range imaging algorithms were used for ground operations by the Mars Pathfinder mission in 1997 to map terrain with stereo imagery from the lander for planning daily operations for the Sojourner rover. He initiated the development at JPL of computer vision techniques for autonomous safe landing and landing hazard avoidance for missions to Mars, asteroids, and comets. His group developed the Descent Image Motion Estimation System (DIMES) that will be used to estimate horizontal velocity of the MER landers during terminal descent. Algorithms developed in his group for onboard crater recognition have

been selected as the backup approach to regional hazard avoidance for the MSL mission. He also conducts research on terrain classification with a wide variety of sensors for off-road navigation on Earth. He was awarded the NASA Exceptional Achievement Medal in 2001 for his contributions to computer vision for space missions. He is an Adjunct Assistant Professor of Computer Science at the University of Southern California.

John R. Wright is a member of the Visualization and Science Applications group at JPL. He is the task leader for the task of producing terrain models from various mission data as well as a software developer for the rover sequence generation tool for the upcoming Mars missions. He has also been involved in a variety of other visualization tasks at JPL including Cognizant Engineer for visualization for the Tropospheric Emissions Spectrometer (TES) instrument. John has been at JPL for nine years. Prior to joining JPL, he worked on terrain visualization tools at Hughes Aircraft Co. for flight simulator applications utilizing voxels.

Rongxing (Ron) Li is a professor at the Department of Civil & Environmental Engineering and Geodetic Science of The Ohio State University. He holds a B.S. and a M.S with honors in surveying and mapping from Tongji University, China, and a Ph.D. in photogrammetry and remote sensing from Technical University of Berlin, Germany. Dr. Li is a Participating Scientist of the Mars Exploration Rover mission, a member of the NASA Geodesy and Cartography Working Group, and a member of the Committee on National Needs for Coastal Mapping and Charting in the US National Academies. His research interests include digital mapping, Mars rover localization and landing site mapping, coastal and marine GIS, and spatial data structures.

Kaichang Di is a research associate at the Department of Civil & Environmental Engineering and Geodetic Science of The Ohio State University. He holds a B.S. with honor, a M.S. and a Ph.D. in photogrammetry and remote sensing, all from Wuhan Technical University of Surveying and Mapping, China. Dr. Di is a collaborator of the Mars Exploration Rover mission, and a member of the NASA Geodesy and Cartography Working Group. His research interests include digital mapping, Mars rover localization and landing site mapping, high-resolution satellite image processing, and spatial data mining.