

TOWARDS AUTONOMOUS MARS ROVER LOCALIZATION: OPERATIONS IN 2003 MER MISSION AND NEW DEVELOPMENTS FOR FUTURE MISSIONS

K. Di^{a,*}, J. Wang^a, S. He^a, B. Wu^a, W. Chen^a, R. Li^a, L. H. Matthies^b, A. B. Howard^b

^a Mapping and GIS Laboratory, CEEGS, The Ohio State University, 470 Hitchcock Hall, 2070 Neil Avenue, Columbus, Ohio 43210, U.S.A. {di.2, wang.813, he.119, wu.573, chen.1284, li.282}@osu.edu

^b Jet Propulsion Laboratory, California Institute of Technology, Pasadena, California 91109, U.S.A. {lhm, abhoward}@robotics.jpl.nasa.gov

Commission I, ICWG I/V

KEY WORDS: Absolute Orientation, Bundle Adjustment, Close Range Photogrammetry, Computer Vision, Matching, Navigation

ABSTRACT:

During the 2003 Mars Exploration Rover (MER) mission, onboard rover localization has been performed primarily by IMU, wheel-odometry, and sun-sensing technologies. In cases where the rover experiences slippage caused by traversing loose soil or steep slopes, particularly in a crater, the onboard visual odometry (VO) technique is applied. A bundle adjustment (BA) method has been performed on Earth to achieve a high-accuracy solution of rover positions by building and adjusting an image network containing all panoramas and traversing images along the entire traverse. An innovative method has been developed to automate cross-site tie-point selection so that BA-based rover localization can be performed autonomously onboard the rover. Recent results of MER mission operations and field test results are reviewed to demonstrate the effectiveness of this autonomous rover localization technology.

1. INTRODUCTION

In a planetary rover mission, localization of the rover with a high degree of accuracy is of fundamental importance both for safe rover navigation and for the achievement of science and engineering goals (Arvidson et al., 2004). During the 2003 Mars Exploration Rover (MER) mission, IMU, wheel odometry, and sun-sensing technologies are all being used to estimate rover positions and attitudes within a nominal accuracy of 10 percent. As of April 16, 2008 (Sol 1524 for Spirit; Sol 1503 for Opportunity), Spirit has traveled 6.67 km while Opportunity has traveled 11.09 km (actual distances traveled, not odometry measures). Onboard visual odometry (VO) is being used to track terrain features appearing in sequential images in order to correct errors caused by wheel slippage (Maimone et al., 2007). Due to limitations in computational speed, VO has only been applied to relatively short distances where the rovers have traveled on steep slopes or across loose soils, for example.

In support of MER mission operations, researchers at the Mapping & GIS Lab of The Ohio State University (OSU) have been collaborating with JPL and other mission teams in performing bundle adjustment (BA)-based rover localization and topographic mapping since the landing of the two rovers in January 2004 (Li et al., 2005; Di et al., 2008). This BA technology uses tie points to link images taken at different rover locations, thereby forming an image network and allowing adjustment of the image orientation parameters to improve localization accuracy. Topographic maps, rover traverse maps, and updated rover locations have been produced and distributed to the science and engineering team members through a WebGIS for science analysis, long term planning and mission operations (Li et al., 2007a).

The key to the success of BA-based rover localization is selection of a sufficient number of well-distributed tie points for linking the images along the rover traverse. Autonomous rover localization requires full automation of tie-point selection. From the beginning of MER operations, tie points linking a stereo pair (intra-stereo tie points) and tie points linking adjacent stereo pairs within one panorama taken at one rover location (inter-stereo tie points) were selected automatically during MER operations. However, cross-site tie points (ones that link panoramas taken at different rover locations) were selected manually during MER operations for the first three years. Recently, we developed an innovative approach to automatic cross-site tie-point selection so that BA-based rover localization can be autonomously performed onboard the rover (Li et al., 2007b). The new approach has been verified using actual Spirit rover data as well as field test data acquired at Silver Lake, California. This new autonomous BA software has been applied in MER operations since August 2007.

2. MARS ROVER LOCALIZATION DURING MER MISSION OPERATIONS

At the Gusev Crater landing site, localization of the Spirit rover has been performed sol by sol based on incremental bundle adjustment using full or partial Navcam/Pancam panoramic images along with, occasionally, forward- and backward-looking Navcam/Pancam middle-point survey images. The achievable localization accuracy has been evaluated based on a consistency check of the BA results. Overall, after BA, 2D accuracy generally ranged from less than 1 up to 1.5 pixels while 3D accuracy was at a centimeter to sub-meter level (Li et al., 2005; Di et al., 2008). Figure 1 shows the Spirit bundle-adjusted traverse map as of Sol 1524 in the area of Home Plate.

* Corresponding author.

3. NEW DEVELOPMENTS FOR AUTONOMOUS ROVER LOCALIZATION

In our new autonomous rover localization approach, VO and BA methods are integrated with the expectation of achieving high efficiency and full automation. As illustrated in Figure 5, BA is performed at waypoints (panoramic sites and mid-point survey positions), while VO is performed between waypoints. The BA obtains the following data from VO: tracked features, refined image-orientation parameters as an input, and first and last stereo pairs. After BA, rover positions are updated at subsequent waypoints. The overall flowchart of this autonomous BA-based rover localization process is shown in Figure 6. This process includes initialization of image parameters (including inputs from VO), extraction and matching of interest points, selection of tie points, and bundle adjustment.




Figure 5. Configuration of the onboard image network

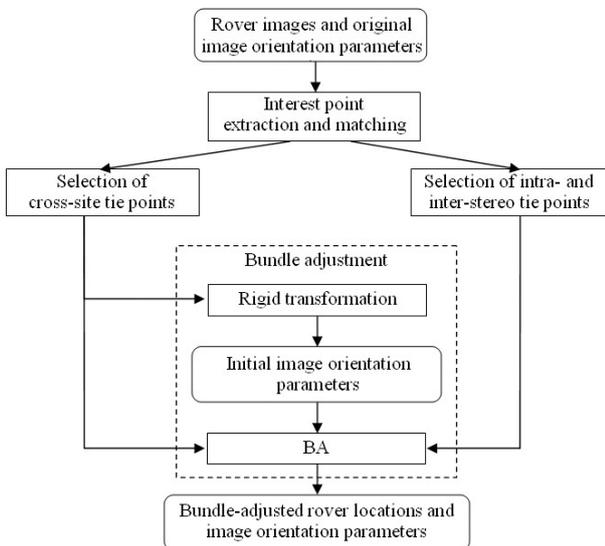


Figure 6. Flowchart of autonomous BA-based rover localization

The key to the success of autonomous BA is selection of tie points, in particular, cross-site tie points. A great challenge is that the cross-site tie points can look significantly different from different viewpoints, especially from forward- and backward-looking images. We have developed a new approach to automatic selection of rocks as cross-site tie points through rock extraction, rock modeling and rock matching (Figure 7).

Pre-screening and fault detection algorithms were also developed to ensure there is no mismatch in the final tie-point selection results.

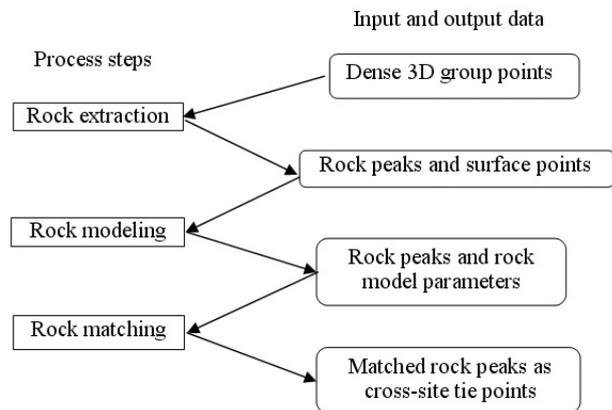


Figure 7. Diagram of automatic cross-site tie-point selection

Rocks are the major landmarks that can be easily identified in most of the ground images. Usually, rocks are composed of distinct rock peaks and surface points. A rock peak is extracted as the local maxima within a predetermined window from set of 3D ground points generated by dense matching of stereo images. Starting from this extracted rock peak, a plane is estimated using those terrain points within an area of 70×70 cm from the rock peak. The initial rock height H is calculated as the perpendicular distance from the peak to the fitted plane. Surface points are searched for iteratively among the candidate points above the fitted plane using a dynamic search range of kH , where k varies from 0.3 to 1.7 based on a ground truth experiment in which manual measurements of rocks at the Spirit site were made and the coefficient k was calculated. Figure 8 shows examples of rock peaks and rock surface points extracted from Spirit rover images. The green dots are the rock peaks, while the red dots are the extracted surface points.

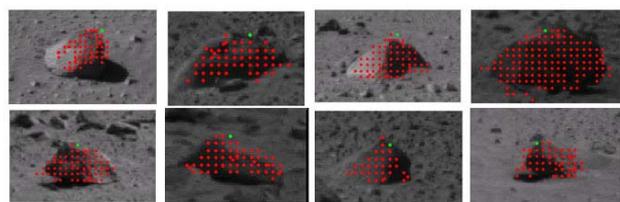


Figure 8. Examples of extracted rocks showing peaks (green dots) and surface points (red dots)

Each rock is then modeled using one of a number of analytical surface models such as hemispheroid, semi-ellipsoid, cone and tetrahedron. The parameters of each individual rock model for a rock are estimated by a least-squares fitting using the surface points on the rock. The model with the minimum root-mean-square error is considered the best model for that rock.

Rock matching was used to find corresponding rocks in the two sets of rocks extracted from two different sites. The rock matching technique we have developed uses rock pattern matching to describe global rock-distribution patterns and rock model matching to depict individual rock similarities (Li et al.,

2007b). The peaks of the matched rocks that pass both rock pattern matching and rock model matching are taken as cross-site tie points. To ensure the effectiveness of the matching, and an even distribution of the tie points, we define a 4×4 grid in the area of overlap between the two sites. Within each grid cell, we select a limited number of significant rocks (e.g., up to 3), which are usually the highest rocks in the grid cell. Only the selected significant rocks at the current site are used to find their corresponding rocks at the adjacent site. Figure 9 shows an example of automatically selected cross-site tie points between Sites 1200 and 1300 of Spirit rover. These two sites are 23 m apart and the image data were acquired on Sols 61 and 62, respectively.

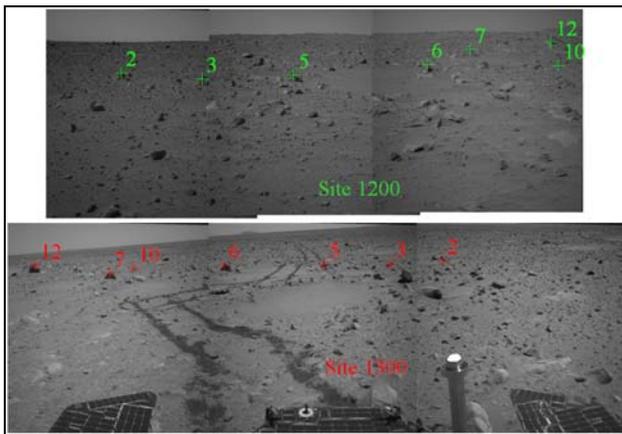


Figure 9. Automatically selected cross-site tie points between Sites 1200 and 1300 of Spirit rover

Pre-screening and fault detection were based on extensive tests and statistical analysis. We found that traverse distance, distance ratio, and the number of used peaks were the most important factors for fault detection. The distance ratio compares two distances: a rock to the camera position at on site versus the same rock to the camera at the adjacent site. At the pre-screening step, pairs with the following conditions were excluded: 1) traverse leg length being less than 30 m, or 2) number of rock peaks extracted being less than 20, which is not sufficient for significant peak selection. In fault detection, we excluded rocks with distance ratios less than 0.3, rocks with unreliable modeling parameters, rocks with unmatched local terrain at both sites, and sites whose number of matched rocks in the final result were less than 3. These pre-screening and fault detection strategies ensure that the successfully selected cross-site tie points are of high quality.

4. VERIFICATION OF AUTONOMOUS ROVER LOCALIZATION TECHNOLOGE USING SPIRIT ROVER DATA AND FIELD TEST DATA

4.1 Verification using Spirit Rover Data

We have tested our new software using a 318 m traverse (19 pairs of sites) taken by Spirit from Sols 574 to 648 in the Husband Hill summit area. The test results are shown in Figure 10. Black dots show sites where Navcam or Pancam panoramic images were taken. Green segments delineate traverse legs outside of the test area. Red segments designate those traverse segments that have passed fault detection and successfully

completed BA. Both yellow and orange segments show those segments that failed the BA. Yellow segments designate those excluded by pre-screening, while orange segments designate those excluded by fault detection. A success rate of 68 percent (13 out of 19 pairs) was achieved, which is a very successful result considering that the MER-A traverse was not designed originally for autonomous BA.

The performance of the autonomous BA is shown in Figure 11, where differences between the blue (telemetry-based) and red (BA-based) lines represent the differences between the telemetry and BA positions. Traverse segments that passed fault detection and successfully finished BA are indicated by solid red lines while the dashed red lines represent those pairs of sites that failed BA, whether excluded by pre-screening or fault detection. Test results using MER data have shown that the proposed method is effective for medium-range traverse segments (up to 26 m). As an example, in the first segment (Sites 11304 to 11308), BA corrected the rover's position by 5.6 percent (0.95 m out of a total segment length of 16.96 m).

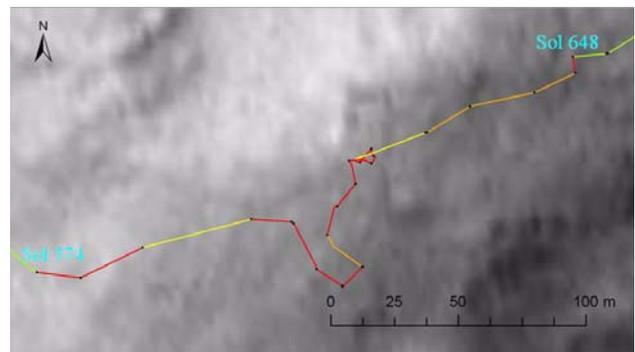


Figure 10. Map of the bundle-adjusted rover traverse of the Spirit rover in the Husband Hill summit area

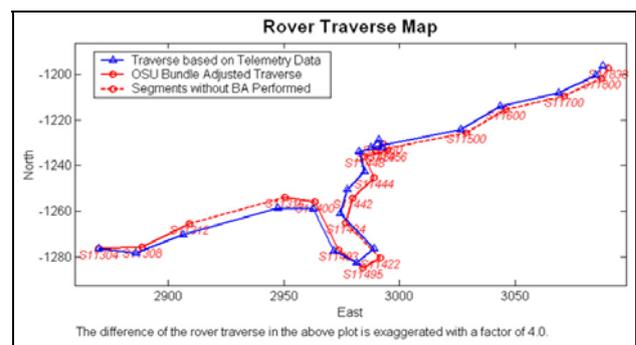


Figure 11. Comparison of Spirit rover traverses in the Husband Hill summit area. Blue line is the telemetry-based traverse and the red line is the traverse computed from the autonomous BA.

Since August 2007, this newly developed software has been employed to perform automatic rover localization for the Spirit rover in the Home Plate area (Figure 1) in the Earth-based data processing environment for ongoing MER mission operations. The developed software has been able to automatically select cross-site tie points for 71 percent of the total number of 38 traverse segments. Over a traverse of 270.92 m, it has corrected the rover's position by 11.03 m (4.07 percent). For the remaining 29 percent of the traverse segments, despite being

unable to select enough cross-site tie points (whether due to lack of significant rocks or long traverse lengths greater than 30 m) the software was still very helpful in assisting operators to rapidly select cross-site tie points for segments within 30 m. The process was reduced to just several minutes. In the past, it would take tens of minutes or even hours to manually select only one cross tie point. This demonstrates that the software is being effectively used in the ongoing MER mission for daily operations.

4.2 Verification using Silver Lake Field Test Data

In order to test the performance of the autonomous BA algorithm and the integration of BA and VO, a field test was conducted at Silver Lake, California, on January 14, 15 and 16, 2007. A radio-controlled LAGR rover (Matthies et al., 2005), capable of capturing panorama and VO images, traversed about 5.5 kilometers (as shown in Figure 12). VO images were taken continuously at a rate of 2 fps; BA panoramic images were taken at the ends of traverse segments with a typical segment length of 25 m. The positions of the rover were obtained from the DGPS (Differential GPS) at a data acquisition rate of 2 Hz, which matched the VO image acquisition rate. The DGPS-determined rover positions were used as ground truth to evaluate the localization accuracy of BA and the integration of BA and VO.



Figure 12. Silver Lake traverse (base map from Google Map)

Along the entire 5.5 km traverse, the rover acquired about 20,000 frames of VO images with a step length of 30 cm and 186 sites of panoramic images (3534 stereo pairs). Both the VO and panoramic images are 1024×1024 pixels in size and 8 bits in grayscale. The terrain captured in these images falls into three categories: rocky outcrops, bushes, and dry lakebed (see Figure 12). The stretch of rocky outcrops ran for about 210 m and was imaged by panoramas from 14 sites. Panorama images from 80 sites covered the bushy area along a traverse about 2.2 km. Although the shape of bushes is different from that of rocks, which are the main features on Martian surface, our software achieved a correct percentage of about 76. The remaining images were mainly obtained on the dry lakebed. Cracks running across the dry lakebed were very convenient for image matching in VO, but make it impossible to pick cross-site tie points for BA.

In order to evaluate the performance of both the BA and the integration of BA and VO, we tested our software in two different ways: without VO data, and with the integration of

VO data. Before the VO processing result was provided, the software was tested only with the panorama data. The positions of the panorama images were initially obtained by DGPS. They were added with an artificially set noise of 10 percent of the distance between two consecutive panorama sites. This 10 percent figure was used to make this positioning error equivalent to the maximum positioning error from wheel odometry. After the cross-site tie points were selected, the image position and attitude (both having errors) are refined by bundle adjustment. This test was conducted in an area of rocky outcrops with a traverse of 14 segments (206 m). For these 14 segments, the software was able to automatically achieve correct cross-site tie points within 11 segments. One segment of 31 m was excluded by pre-screening and 2 segments were excluded by fault detection. The success rate, therefore, is 79 percent (85 percent after pre-screening).

Computation of the VO and BA integration was performed in the local coordinate frames (east-north-up) of three stretches of traverses. Each local frame origin is at the center of the first panorama of that stretch. The orientation (azimuth and tilt) of the first panorama in each local frame was calculated manually based on the first panorama, an adjacent panorama for this stretch of traverse, and DGPS data. We matched the first VO pair to the first panorama (already in the local frame) and then transformed all the VO poses to the local frame. After BA, we evaluated the localization error at the end point by comparing the BA-derived position and DGPS position (as shifted to the local frame). Figure 13 shows the BA and VO integration results in the rocky outcrop area. The blue, red and black lines represent rover traverses from VO, from integration of VO and BA, and from ground truth, respectively. We can observe that the integrated BA and VO were significantly better than the initial VO result. This indicates that the BA panoramas improved the geometric strength of the image network and provided better localization accuracy than VO alone. Relative accuracy improved from 27.1 to 3.9 percent. As for automatic cross-site tie-point selection, we obtained the same success rate as results without VO data in the rocky outcrop area.

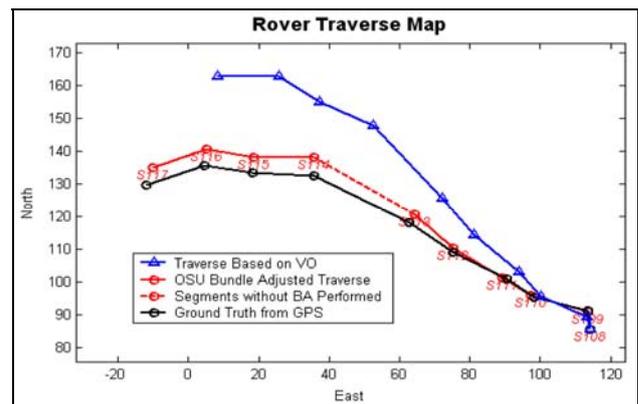


Figure 13. Result of integration for rocky outcrop area near Silver Lake, CA (units: m)

We also tested the software in a bushy area around Silver Lake using a total of 81 pairs and a traverse length of 2.2 km. Among the 81 sites were 10 consecutive pairs (about 0.7 km) whose traverse lengths were approximately 50 m each, which is beyond our software's ability to reliably extract features. The fault detection module excluded these 10 pairs, thus they are

not included in the statistical analysis. For the rest of the 71 pairs, covering a 1.6 km traverse, the success rate was 59 percent after pre-screening. This result is not as good as results for the rocky outcrop area because the characteristics of bushes are sufficiently different from the rock shapes for which our software was designed. Despite the lower success rate, this software still achieved much better localization results than VO only. The 1.6 km traverse was evaluated as two separate parts because the corresponding VO data were processed separately. For the first part of the traverse (1.2 km), the relative localization accuracy was reduced from 19.7 to 4.1 percent (see Figure 14). The relative localization accuracy for the second part of the traverse (0.4 km) was refined from 9.9 to 8.7 percent. It is to note that in the traverse shown in Figure 14, VO failed at 7 positions due to a lack of sufficient image overlap. The large localization error (19.7 percent) is predominately wheel odometry error. Integration of BA and VO significantly reduced this error.




Figure 14. Test results for traverse of bushy area near Silver Lake, CA (units: m)

5. CONCLUSIONS

During the MER mission, the rovers are primarily localized onboard by IMU, wheel-odometry, and sun-sensing technologies. VO technique has been effectively applied onboard over short distances to correct slippage errors. The BA method has been performed on Earth to achieve a high-accuracy solution of the entire rover traverse. Localization based on integrated BA and VO based has greatly supported mission operations for safe navigation and for achieving scientific and engineering goals.

We have developed a new approach to autonomous localization for long-range rover traverses for future rover missions. This new approach integrates VO and BA with the expectation of achieving high efficiency and full automation. In particular, an automatic cross-site tie-point selection method has been developed to enable the BA to be autonomous. Test results using MER's Spirit rover data as well as field test data acquired at Silver Lake, California, have verified the effectiveness of our autonomous BA software. This software could be used for

onboard autonomous Mars rover localization in a rock-abundant landing site (like MER's Spirit landing site) in future Mars rover missions.

ACKNOWLEDGEMENTS

This work was partially performed at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration (NASA). Funding for this research by the NASA Mars Exploration Program and the NASA Mars Technology Program is acknowledged.

REFERENCES

- 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. Haldemann, K. Herkenhoff, J. Johnson, G. Landis, R. Li, R. Lindemann, 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. Localization and physical properties experiments conducted by Spirit at Gusev Crater. *Science, Special Issue on MER 2003 Mission*, 305(5685), pp.821-824, doi:10.1126/science.1099922.
- Di, K., F. Xu, J. Wang, S. Agarwal, E. Brodayagina, R. Li, L. Matthies, 2008. Photogrammetric Processing of Rover Imagery of the 2003 Mars Exploration Rover Mission. *ISPRS Journal of Photogrammetry and Remote Sensing*, 63, pp. 181-201, doi:10.1016/j.isprsjprs.2007.07.007.
- Li, R., K. Di, A.B. Howard, L.H. Matthies, J. Wang, and S. Agarwal, 2007b. Rock Modeling and Matching for Autonomous Long-Range Mars Rover Localization. *Journal of Field Robotics*, 24(3), pp. 187-203.
- Li, R., K. Di, J. Wang, X. Niu, S. Agarwal, E. Brodayagina, E. Oberg, and J.W. Hwangbo, 2007a. A WebGIS for Spatial Data Processing, Analysis, and Distribution for the MER 2003 Mission. *Photogrammetric Engineering and Remote Sensing, Special Issue on Web and Wireless GIS*, 73(6), pp. 671-680.
- Li, R., S. W. Squyres, R. E. Arvidson, B. A. Archinal, J. Bell, Y. Cheng, L. Crumpler, D. J. Des Marais, K. Di, T. A. Ely, M. Golombek, E. Graat, J. Grant, J. Guinn, 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 and Remote Sensing, Special issue on Mapping Mars*, 71(10), pp. 1129-1142.
- Maimone, M., Y. Cheng, and L. Matthies, 2007. Two years of Visual Odometry on the Mars Exploration Rovers. *Journal of Field Robotics*, 24(3), pp. 169-186.
- Matthies, L., M. Turmon, A. Howard, A. Angelova, B. Tang, 2005. *Journal of Machine Learning Research*. 1, pp.1-48.