

9-2012

Position Estimation by Registration to Planetary Terrain

Aashish Sheshadri
Carnegie Mellon University

Kevin Peterson
Carnegie Mellon University

Heather L. Jones
Carnegie Mellon University

William Whittaker
Carnegie Mellon University

Follow this and additional works at: <http://repository.cmu.edu/robotics>

 Part of the [Robotics Commons](#)

Published In

International Conference on Multisensor Fusion and Information Integration (MFI), 432-438.

This Conference Proceeding is brought to you for free and open access by the School of Computer Science at Research Showcase @ CMU. It has been accepted for inclusion in Robotics Institute by an authorized administrator of Research Showcase @ CMU. For more information, please contact research-showcase@andrew.cmu.edu.

Position Estimation by Registration to Planetary Terrain*

Aashish Sheshadri, Kevin M. Peterson, Heather L. Jones and William L. “Red” Whittaker

Abstract— LIDAR-only and camera-only approaches to global localization in planetary environments have relied heavily on availability of elevation data. The low-resolution nature of available DEMs limits the accuracy of these methods. Availability of new high-resolution planetary imagery motivates the rover localization method presented here. The method correlates terrain appearance with orthographic imagery. A rover generates a colorized 3D model of the local terrain using a panorama of camera and LIDAR data. This model is orthographically projected onto the ground plane to create a template image. The template is then correlated with available satellite imagery to determine rover location. No prior elevation data is necessary. Experiments in simulation demonstrate 2m accuracy. This method is robust to 30° differences in lighting angle between satellite and rover imagery.

I. INTRODUCTION

Global localization for planetary rovers has bearing on many applications of exploratory and scientific interest. Rovers will be expected to maintain precise pose estimates over long traverses, build detailed models of interesting regions, and accurately correlate in-situ measurements with satellite imagery to achieve high science return. Autonomous global localization results in more efficient use of human operators, enabling them to focus on scientific and exploration objectives instead of rover positioning. It also makes possible driving without communication for operations such as forays into radio-dark craters at the poles of the Moon. Despite its impact, precise, autonomous, global localization has not yet been achieved.

Current means for localization include absolute methods like radio tracking from a satellite and matching to elevation data, and relative methods like wheel odometry, inertial navigation, and visual odometry. Absolute methods using radio rely on infrastructure and are limited to operate within line-of-sight of the radio beacon. Elevation data is often much lower resolution than available imagery, limiting localization precision. Relative methods drift over time and become inaccurate without human input. These approaches are insufficient for long-range traverse where regular updates from Earth are infeasible and line of sight is not

guaranteed. The need is for precise, autonomous, globally consistent localization.

This paper demonstrates the feasibility of precise global rover localization by registration to terrain using image correlation. The availability of high-resolution satellite imagery [1] [2] is exploited to provide orthographic terrain maps. Local terrain geometry and appearance are modeled using camera and LIDAR (LIght Detection And Ranging) panoramas. The panoramas are transformed into an orthographic view and correlated with satellite imagery to determine position.

Accuracy of the technique is evaluated in simulation. Sensitivity to resolution, camera height, pose uncertainty, and lighting is evaluated. Results demonstrate 94% probability of better than 2m global accuracy under typical rover conditions with graceful degradation in off-nominal conditions. The technique is also evaluated in the context of a traverse to a site of interest. A 5km rover traverse to an interest site is simulated with the rover localized after every 100m traverse. Results demonstrate a successful traverse with localization error less than 2m for 94% of positions.

II. RELATED WORK

Terrestrial outdoor robots depend heavily on GPS for localization. Typical approaches use an integrated GPS/INS to provide highly accurate, low latency updates to position. In particular the successful DARPA Grand Challenge vehicles all used GPS as primary means for localization [3] [4]. In environments where GPS is periodically unavailable, radio beacons are commonly used to bridge GPS outages [5]. In environments where GPS is only occasionally available, visual odometry and bundle adjustment provide stable pose for medium range traverses [6]. The infrastructure to perform GPS or radio-based localization is unavailable on planets other than the Earth.

Non-GPS sensor-based localization is used terrestrially when motion is repetitive. Notable examples are shipping dock automation [7] and autonomous urban driving [8]. In these environments, sensor-centric maps are built ahead of time using SLAM. These approaches assume that the route can be driven to build the map before the map is needed for autonomous operations. This is not the case in planetary exploration.

Planetary rovers have historically relied on a combination of relative localization and human input to provide reliable localization. During the Pathfinder mission, human operators triangulated the Sojourner rover position by periodically viewing the rover [9] with an on-board camera.

* Research supported by NASA contract NNX11CB55C.

A. Sheshadri was with the Robotics Institute, Carnegie Mellon University, 5000 Forbes Ave, Pittsburgh, PA 15213 USA. He is now with the Department of Computer Science, The University of Texas at Austin, 1616 Guadalupe, Suite 2.408 Austin, TX 78701 (e-mail: aashish.sheshadri@gmail.com).

K. M. Peterson, H. L. Jones, and W. L. Whittaker are with the Robotics Institute, Carnegie Mellon University, 5000 Forbes Ave, Pittsburgh, PA 15213 USA (e-mail: {kp | hlj | red}@cs.cmu.edu).

Coverage was limited to areas with line of sight to the lander.

The Mars Exploration Rovers (MERs) with semi-autonomous traverse capabilities far exceeded their expected traverse distance, travelling a combined 41.7 km, this was achieved over 2947 sols (nearly 8 years) [10]. The MERs were first localized at their landers using radiometric tracking from the Mars Odyssey orbiter. Cartographic triangulations combining imagery from lander descent, ground panoramas, and orbital imagery served as a resource for human updates to optimally determine absolute lander position [11].

MER localization beyond lander line of sight is achieved primarily by wheel odometry and IMU (Inertial Measurement Unit) measurements. Infrequent heading updates from sun sensing reduce angular drift. Visual odometry is used intermittently to detect slip and provide accurate position estimation. Bundle Adjustment is performed incrementally, sol-by-sol, on downlinked rover imagery to further improve position estimates [11].

Visual odometry on the MERs is used sparingly due to high processing times (up to 3 minutes per frame on its 20MHz processor) and the requirement for sufficient overlap between frames to extract dependable features [12]. As a result, the longest traverse that utilized visual odometry was < 10 meters [11].

Experimentation has shown that velocity measurement using optical flow and visual odometry produce pose estimates with two to three percent error over several kilometers of travel in Moon- and Mars-relevant terrain [13]. While still relative, these measurements represent acceptable error for missions with periodic human oversight.

Many approaches to improving relative pose estimation exist. A brief introduction to individual localization methods and issues with integration to reduce position uncertainty is discussed in [14]. [15] discusses a method to register locally generated elevation maps to previously generated elevation maps, both generated using stereo vision techniques. Prior elevation maps are generated by raising the rover stereo mast to gather a partial panorama view of the intended traverse region before making the traverse. [16] investigates a solution to enhancing accuracy of visual localization by tracking selective pixels frame to frame. The accumulated error is minimized by tracking a few these pixels over multiple frames. Xu presents a method for accurate rover

localization by autonomously extracting tie points between subsequent pairs of stereo images, however the localization method is semi-autonomous as there is need for human verification of the tie points [17]. Olson and Abi-Rached introduced wide baseline stereo to build accurate models of distant terrain and hence enable localization and navigation over larger areas [18]. Kubota et al. propose a scheme that builds local elevation maps using stereo or LIDAR, extracts feature points in individual local maps and matches these points between local maps under geometric constraints [19].

The relative localization methods, discussed above, drift over time. Consequently, over long traverses pose becomes inaccurate and external input is required. Accurate real-time global localization will enable kilometer-per-command traverses and fully autonomous operation.

The VIPER algorithm provides a global position solution by identifying unique structure on the horizon (e.g., mountain peaks and local variation) in image panoramas [20]. The localization errors are on the order of 100 meters in terrain with appreciable local variation in elevation. Barfoot, et al. proposed global localization by correlating digital elevation maps (DEMs) from prior orbital data with LIDAR-based DEMs produced on-board the rover [21] [22]. Because DEMs built from orbital data are typically low resolution, this requires long-range 3D mapping on-board the rover. Both of these methods require significant variation in terrain elevation to provide pose estimates.

Recent availability of high-resolution (on order of 0.25m to 1.0m per pixel) imagery of the Moon and Mars provides new means for localization. Variation in terrain appearance is independent of variation in elevation and serves as an alternative to the approaches discussed above. The approach proposed here correlates local visual appearance with overhead imagery of terrain to produce accurate global localization. Experiments show that pose estimates with better than two-meter accuracy can be achieved under nominal conditions. The resulting position estimates are independent of local terrain geometry and do not drift.

III. GLOBAL ROVER LOCALIZATION

The approach is to correlate terrain appearance near the rover with high resolution satellite imagery. Terrain appearance is measured using a panorama of camera and LIDAR images. To accomplish this, the rover collects pairs of camera and LIDAR images of the surrounding terrain. Each camera image is projected onto the corresponding

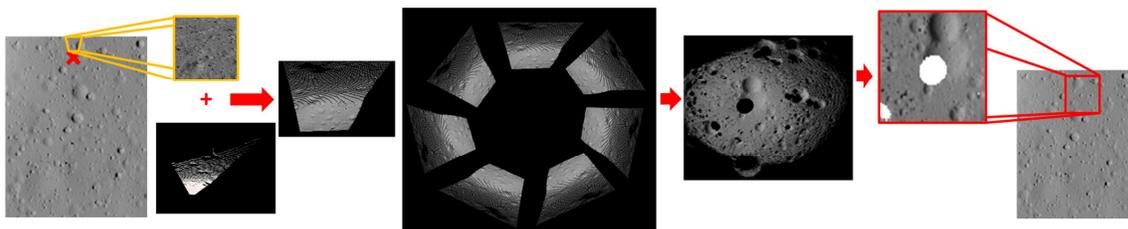


Figure 1 Global Localization Method Overview. From left to right: 1) Camera and LIDAR image pairs are collected around the rover. 2) Imagery is projected onto LIDAR geometry to produce colorized LIDAR. 3) Colorized LIDAR frames are stitched together to produce a panorama. 4) The panorama is projected into an overhead orthographic view and correlated with terrain to determine location.

LIDAR geometry to produce a colorized point cloud. The colorized clouds are stitched together into a cohesive model projected into an overhead orthographic view and cropped into a small template. The template is correlated with prior imagery (from e.g., satellite or flyover). Figure 1 depicts this process.

A. Building the Panorama

Panoramas can be built by rotating an encoded sensor head to produce a 360-degree view of surrounding terrain. The motion of the camera and LIDAR between frames is then known to the precision of the sensor head encoders. This is the approach evaluated in the results section of this paper.

If the cameras cannot be rotated independently of the rover chassis, panoramas can be generated by rotating the rover chassis. In this case the motion of the camera and LIDAR between frames is not known precisely and must be reconstructed. The motion of the rover is estimated by stitching consecutive LIDAR-camera frames with visual odometry and the iterative closest point algorithm. The initial transformation between frames is computed using stereo visual odometry [23]. Initial transformations are refined using iterative closest point on the LIDAR frames [24]. Localization accuracy is largely dependent on template accuracy, which in turn depends on the accuracy with which individual frames are stitched to form a single model. This was found to require 75% overlap between frames when the rover base rotates, whereas little to no overlap is required when using encoded sensor head rotations, so fewer frames are needed to build the panorama.

Both methods of generating panoramas were evaluated. Representative templates are shown in Figure 2. The two techniques produce equivalent localization performance.

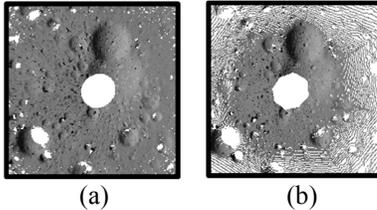


Figure 2 (a) Orthographic view of the panoramic model stitched using the Rotating Rover Base approach. The higher point density in this model is due to the larger number of frames required. The templates produce equivalent localization results. (b) Orthographic view of the panoramic model generated using the Rotating Camera Head approach. White pixels indicate locations where no information is known about the terrain.

B. Generating the Template

Once the transformations between frames are determined, each LIDAR point cloud is projected into its corresponding camera image to recover camera color and texture associated with the local geometry. The colorized frames are stitched together using the frame-to-frame transformations to build a cohesive 3D panorama in the reference frame of the first image of the sequence.

The panorama is rotated from the rover-relative orientation into world-frame orientation using known rover attitude. It is assumed that rover yaw is known from a star tracker or sun sensor [11] and that local vertical is known from accelerometer measurements of the gravity vector.

The now locally-level colorized geometry is projected into an overhead orthographic image to use as a template for correlation. This is done by binning colorized LIDAR points into a 2D raster with the same resolution as the map imagery. When multiple LIDAR points lie in the same bin, the average color of the points is used.

Because terrain is uneven and self-occluding, geometry for some regions of the model is not known. These locations are marked as blank.

C. Estimating Rover Position

Rover position is determined by correlating the orthographic template with satellite imagery. Since satellite imagery is available at varied resolutions, and the panorama is always at a higher resolution than the satellite imagery, the template is constructed to match the resolution of the satellite imagery. A search area is selected using the best guess of the rover position and the uncertainty associated with that position. Position is determined by running a normalized cross correlation between the template and the map interest region. Normalized cross correlation was modified to correlate only measured terrain in the template i.e. regions marked blank were ignored.

The objective is to determine position of the template t in a two dimensional overhead image f . Let $f(x, y)$ denote the intensity value of image f of size $M_x \times M_y$ at the point (x, y) , $x \in \{1, \dots, M_x\}$, $y \in \{1, \dots, M_y\}$. Similarly, let $t(x, y)$ denote the intensity value of template t of size $N_x \times N_y$ at the point (x, y) and $g(x, y)$ denote in binary validity of measurements at the point (x, y) with $x \in \{1, \dots, N_x\}$ and $y \in \{1, \dots, N_y\}$. The position (u_{pos}, v_{pos}) of template t in image f is evaluated as the maximum value of the weighted normalized cross correlation value γ at each point (u, v) for f and the template t , which has been shifted by u steps in x and by v steps in the y . Equation (1) gives the definition of the weighted normalized cross correlation coefficient.

$$\gamma = \frac{\sum_{x,y} w_{x-u,y-v} (f(x,y) - \bar{f}_{u,v}) (t(x-u,y-v) - \bar{t})}{\sqrt{\sum_{x,y} w_{x-u,y-v} (f(x,y) - \bar{f}_{u,v})^2 \sum_{x,y} w_{x-u,y-v} (t(x-u,y-v) - \bar{t})^2}} \quad (1)$$

$$\bar{f}_{u,v} = \frac{1}{n_t} \sum_{x=u}^{u+N_x-1} \sum_{y=v}^{v+N_y-1} w_{x-u,y-v} f(x, y) \quad (2)$$

$$\bar{t} = \frac{1}{n_t} \sum_{x=1}^{N_x} \sum_{y=1}^{N_y} w_{x,y} t(x, y) \quad (3)$$

$$w_{x,y} = \begin{cases} 1, & \text{if } g(x, y) = 1 \\ 0, & \text{Otherwise} \end{cases} \quad (4)$$

In (1) $\bar{f}_{u,v}$ denotes the mean value of $f(x, y)$ within the area of template t shifted to (u, v) calculated using (2) and \bar{t} is the mean value of template t calculated using (3). In (2) and (3) n_t denotes the number of valid pixels in template t . Weights $w_{x,y}$ are assigned using (4).

IV. EXPERIMENTS

A. Overview

Localization accuracy and robustness was evaluated, on datasets generated in simulation, under a variety of conditions. Accuracy was evaluated as a function of search area (Section IV.D) and rover height (Section IV.E). Robustness was evaluated under varied lighting conditions (Section IV.F) and multiple map resolutions (Section IV.G). Section IV.H demonstrates a long rover traverse to an interest region.

B. Simulation

Test images and maps were generated using a photo-realistic Lunar terrain renderer. Terrain was derived from real data as follows. First, a base terrain mesh was generated from a 2m/post digital terrain map of the Apollo 11 region. This map lacks sufficient detail for rover views, so small rocks and craters were added to the terrain following a statistical distribution derived from Surveyor data [25]. Procedural modeling was used to allow arbitrary detail. Terrain components were combined, detail textures were added, and simulated color and depth data were generated using an off-the-shelf raytracer [26].

Rover images were synthesized using a pinhole camera model with the appropriate field of view and resolution for each sensor. LIDAR data was generated in a similar fashion, but uses the depth channel of the rendering in place of color. Range data was transformed from the orthographic projection of the depth buffer to a perspective projection with rays from the LIDAR center. Satellite imagery was synthesized using an orthographic camera positioned above the terrain looking down.

C. Experimental Setup

The experimental setup is based on a Lunar mission scenario in which a rover carries scientific instruments to prospect for ice and map the distribution of volatiles [27]. Localization precision of 2m is specified for this mission, so successful rover localization is defined in the following experiments as a position estimate within 2m of the actual rover location. A rover designed for this mission is shown in Figure 3.



Figure 3 The rover, Polaris, modeled in these experiments. Polaris is a water prospecting rover designed to carry the RESOLVE payload to the Moon. It carries two cameras and a flash LIDAR for hazard detection and navigation. RESOLVE requires two-meter position accuracy to correlate in-situ measurements with satellite imagery.

The rover has a rotating pan-tilt sensor head with a camera and a flash LIDAR. The sensor head is at a height of

1.5m from the ground. The flash LIDAR has a 45deg Field of View (FOV) with a resolution of 128x128 pixels. The camera has a 45deg FOV and 1000x1000 pixel resolution. For the purpose of this work the camera is assumed to be collocated with the LIDAR. The rover is assumed to have an accurate orientation estimate at all times with the help of a sun sensor and onboard IMU. For panoramas, the camera is pointed down 32.1° from horizontal.

An 800x800m section of the DEM(Section IV.B) was chosen randomly as a test region. Orthographic views of the chosen region were captured at a resolution of 0.25m/pixel for use as maps. Rover panoramas were generated using 8 LIDAR and camera image pairs evenly spaced over one 360° rotation. The localization method was tested (Section IV.D – IV.G) at fifty randomly chosen rover locations over the 800x800m area. The angle of the sun relative to the terrain is adjustable in simulation. For most experiments, the sun angle in the overhead map image was the same as the sun angle in the rover images, though section IV.F investigates the impact of differences in lighting angle between the overhead map and the rover images

D. Localization under Variation in Search Window Size

This experiment determines the sensitivity of localization accuracy to variation in search window size. This experiment is set up as detailed in Section IV.C.

An orthographic view of the region was generated at 0.25m resolution. Localization templates were generated corresponding to each rover panorama as detailed in Section III.B. The templates cover a 15x15m square region centered at the rover. The template resolution was set at 0.25m/pixel to match resolution of the overhead map. Search windows were chosen for each of rover location as a square region of the overhead containing that location.

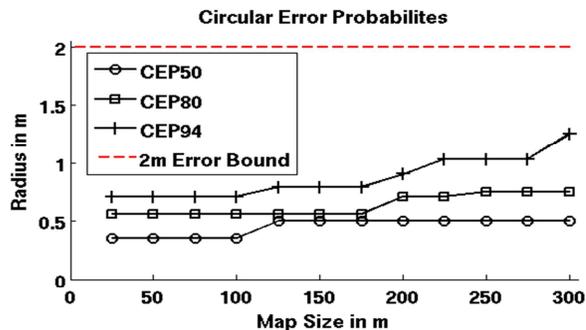


Figure 4 Circular error probabilities for 94%, 80%, and 50% confidence levels plotted as a function of map search window size for a 1.5m camera height. As the search window increases, error increases but stays below the two-meter error threshold.

The search window size was varied from 25x25m to 300x300m with 25m increments in edge length. Results were analyzed using Circular Error Probability (CEP) as the error metric. The CEP94, for example, is the radius of the circle, centered on the true position, which contains 94% of the results (even-numbered CEP values were used, since they correspond to whole numbers in a set of 50 templates). Results of the experiment placed CEP94 within the 2m error bound consistently through the variation in search window

size. Figure 4 shows plots of CEP94, CEP80 and CEP50 for a camera height of 1.5m.

E. Localization under Variation in Sensor Head Height

This experiment examines the dependence of localization accuracy on template size. This experiment is set up as detailed in Section IV.C, except that the rover sensor head height is varied from 0.5m to 3m in steps of 0.5m.

Since a higher sensor head enables a larger view of the local environment, larger image templates were generated with increasing sensor head height, see Figure 5. For the sensor head at 0.5m, the templates were orthographic views of a 5x5m square region centered at the rover. An increase in sensor head height of 0.5m increased the edge length of the template coverage region by 5m. Hence, the largest template was 30x30m. The template resolution was set at 0.25m/pixel to match the resolution of the overhead map. Figure 5 shows the templates generated for a single rover position with sensor head height varying from 0.5m (bottom) to 3m (top).

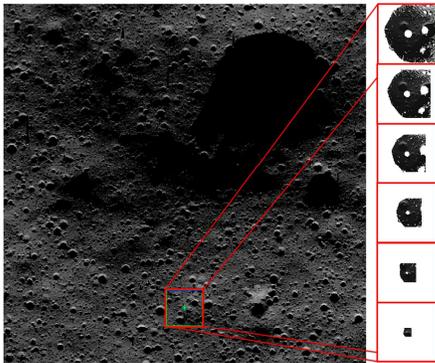


Figure 5 Exemplar templates (right) from one of the search windows (left) showing appearance change as a function of rover height. The bottom-most template is taken from a 0.5 meter high camera, the top-most template is taken from a 3 meter height. The 3m high camera exhibits very low error, while the 0.5m high camera exhibits high error.

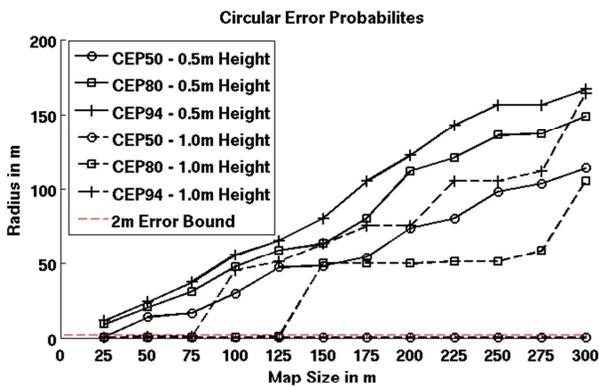


Figure 6 CEP94, CEP80 and CEP50 for rover heights of 0.5m and 1m as the search window is varied from 25m to 300m. Error increases with search window size. The 0.5m high camera exhibits high error rates in all cases, while the 1.0m high camera exhibits acceptable error for small search windows.

The templates for each sensor head height were tested with square search windows from 25m to 300m wide. The results were analyzed using CEP as the error metric. Figure 6 plots CEP50, CEP80 and CEP94 for two of the rover

configurations - sensor head at 0.5m and 1m. CEP94 results for the rest of the sensor head heights are well within the 2m error bound, so these results are not shown as plots. TABLE I lists the number of templates successfully localized.

TABLE I. NUMBER OF TEMPLATES SUCCESSFULLY LOCALIZED WITH VARIATION IN ROVER HEIGHT AND SEARCH WINDOW SIZE.

Search Window	Rover Height (m)					
	0.5m	1.0m	1.5m	2.0m	2.5m	3.0m
100m	12	48	50	50	50	50
200m	8	38	48	49	50	50
300m	7	32	47	48	49	50

F. Localization under Variation in Lighting

This experiment examines the robustness of the localization method to differences in lighting between the map and the rover images. This experiment is set up as detailed in Section IV.C, except that the lighting angle of the maps was varied to simulate a mismatch between the map and rover views.

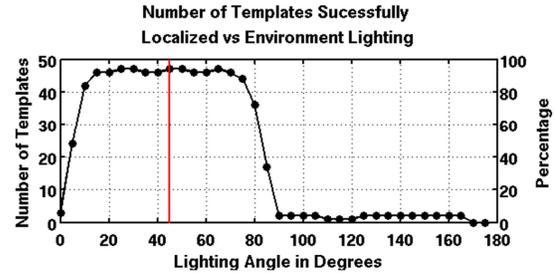


Figure 8 Number of templates successfully localized as the lighting in elevation is varied from 0° to 180° with azimuth at 0° simulating equatorial lighting conditions. Performance is level until the lighting phase angle reaches 40°. Red line indicates rover lighting angle.

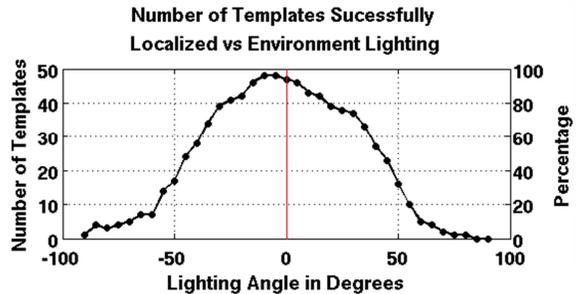


Figure 7 Number of templates successfully localized as the lighting in azimuth is varied from -90° to +90° with elevation fixed at 5° simulating polar lighting conditions. Performance drops off after about 30°. Red line indicates rover lighting angle.

Two sets of panoramas were generated: one representing nominal polar lighting and one representing nominal equatorial lighting. Image templates were generated corresponding to each rover panorama as detailed in Section III.B. The template resolution was set at 0.25m/pixel to match resolution of the generated overhead views.

Orthographic views of the 800x800m region were generated with varied lighting to represent polar and equatorial lighting conditions. In the polar case, lighting

azimuth was varied by $\pm 90^\circ$ in 5° increments with elevation fixed to match the rover views. In the equatorial case, lighting elevation angle was varied from 0° to 175° in 5° increments with azimuth fixed to match the rover views.

The rover was localized within a $300 \times 300\text{m}$ search window. Figure 7 shows a plot of number of successful results against variation in lighting for the polar case. Figure 8 shows this same information for the equatorial case.

G. Localization under Variation in Map Resolution

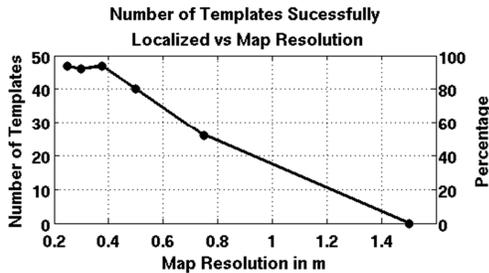


Figure 9 Number of templates successfully localized as map resolution is varied for a 1.5m high camera. Performance is stable until map resolution reaches 0.4m, then slowly drops off. Increasing rover camera height shifts this curve to the right.

This experiment examines localization performance as a function of overhead map resolution. This experiment is set up as detailed in Section IV.C, except that template resolution is varied from 0.25m to 1.5m.

Templates were generated from rover panoramas at 0.25m, 0.3m, 0.375m, 0.5m, 0.75m and 1.5m resolutions (these resolutions correspond to a 15×15 template that is 60, 50, 40, 30, 20, and 10 pixels across, respectively). Overhead maps at 0.25m per pixel were downsampled to match the resolution of the templates.

The rover was localized within a $300 \times 300\text{m}$ search window. Figure 9 shows the number of templates successfully localized for each map resolution.

H. 5km Traverse to an Interest Region

This experiment examines the validity of a long rover traverse under polar lighting using the localization method presented. Initial and final location (Interest Region), separated by 5km, were chosen in the world frame and a path was planned to connect the start and destination. The path was sampled every 100 meters to generate a total of 50 rover views.

To account for traverse uncertainty the sample positions were shifted randomly. The uncertainty in x and y was drawn from a normal distribution with zero mean and 8m

variance. Several rover positions were manually shifted to avoid a large shadow, a result of the polar lighting angle selected for the test, which was clearly visible in the overhead map. The rover was then localized within a $300 \times 300\text{m}$ search window.

The rover was successfully localized in 47 instances of 50 attempts over the traverse, with a mean error of 0.53m. Figure 10 shows a plot of the traverse.

V. DISCUSSION

Experimental results demonstrate 94% probability of better than 2m localization accuracy under nominal conditions. Probability of successful localization gracefully degrades with reduced template size, increased lighting phase angle, and decreased map resolution. Failed localization attempts result in position estimates randomly distributed within the search window.

Probability of successful localization increases with the number of pixels in the template (e.g., with increased map resolution or camera height). Under nominal conditions (0.25m resolution map with 1.5m camera height), the CEP94 was 1.25m – well below the 2m error bound. At 3m height, localization never failed and exhibited an CEP94 of 1.1m. Reducing the search window size results in an increase in the number of successful attempts. A $100 \times 100\text{m}$ search window resulted in zero failed attempts using the standard template.

Localization performs well under equatorial lighting conditions with up to $\pm 40^\circ$ difference between rover and map images. Performance rapidly degrades with further increase in the phase angle. Under polar lighting, localization with $\pm 30^\circ$ difference in azimuth performs well. This indicates that rover localization will function under a wide variety of lighting conditions with a single map. This is important because planetary maps are typically available for only a select set of lighting angles.

Unsurprisingly, accuracy decreases with map resolution and increases with rover height. Acceptable performance was achieved with 0.4m resolution maps at 1.5m camera height. If maps of this resolution are not available, a taller viewpoint would improve accuracy to acceptable levels. Additionally, the camera angle was fixed for all experimentation. Optimizing the camera angle might improve performance, particularly when the camera is low.

Localization failed in the traverse experiments on three occasions with errors of 44m, 92m and 173m, while the mean radius error of the rest of the attempts was 0.53m. The

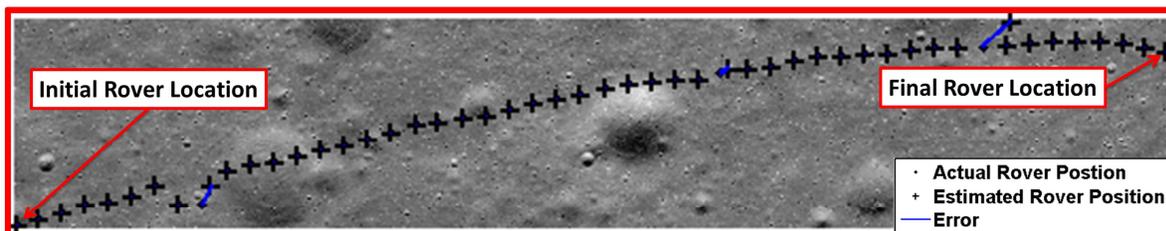


Figure 10 5km rover traverse to an interest region.

failed attempts indicated a localization result dependent on the search window size and well away from the actual position, hence monitoring the localization result over time using a filter will identify such cases.

VI. CONCLUSION

An approach to global localization for planetary rovers was presented and verified in simulation. The approach correlates observed terrain with prior imagery to produce localization solutions with better than 2m accuracy in nominal conditions. Localization performs well when lighting phase angle is within thirty degrees. Performance variation was evaluated as a function of search area size, map resolution, and rover height.

Several promising approaches to increasing localization robustness are apparent and will be investigated in future work. Template uniqueness and information content determine the likelihood of correct registration. Measurements of template information content could be used to determine the likelihood of successful localization with a given template.

Typical rover localization schemes incorporate a series of measurements over time, as well as measurements from many sensors, using filtering techniques. This method provides the precise, global position updates needed to prevent such a filter from drifting. Future work will filter measurements obtained using this method over time. Measurements from digital elevation maps, wheel and visual odometry, and inertial sensors will also be incorporated.

Field experimentation is required to validate these simulation results and will be performed in future work. Changing weather conditions influence both satellite and rover imagery, an issue which needs further investigation for successful implementation in Earth-like environments. Effects of sensor noise and calibration errors in rover models will be investigated in future work. Interactions between localization schemes and on-line planning and control are also topics for future work.

ACKNOWLEDGMENT

The authors acknowledge Jason Koenig for his help in modeling and simulation. This research is supported in part by donations from the NVIDIA Corporation.

REFERENCES

- [1] G. Chin et al., "Lunar Reconnaissance Orbiter Overview: The Instrument Suite and Mission," *Space Science Reviews*, vol. 129, no. 4, pp. 391-419, April 2007.
- [2] T. H. Ebben et al., "The HiRISE camera on-board MRO," in *Proc. of the SPIE Conference on Focal plane arrays for space telescopes III*, vol. 6690, Boulder, CO, 2007, pp. 1-22.
- [3] C. Urmson et al., "Autonomous driving in urban environments: Boss and the urban challenge," *J. of Field Robotics*, vol. 25, no. 8, pp. 425-466, 2008.
- [4] M. Montemerlo et al., "Junior: The Stanford entry in the Urban Challenge," *J. of Field Robotics*, vol. 25, no. 9, pp. 569-597, 2008.
- [5] S. Cobb, "GPS Pseudolites: Theory, Design, and Applications," Stanford University, Palo Alto, CA, PhD Thesis 1997.
- [6] S. Scherer et al., "River mapping from a flying robot: state estimation,

- river detection, and obstacle mapping," *Autonomous Robots*, vol. 33, no. 5, pp. 189-214, May 2012.
- [7] H. F. Durrant-Whyte, "An autonomous guided vehicle for cargo handling applications," *Int. J. of Robotics Research*, vol. 15, no. 5, pp. 407-440, October 1996.
- [8] S. Thrun and C. Urmson, "Self-Driving Cars," San Francisco, CA, Keynote talk at IEEE/RSJ Int. Conf. on Intelligent Robots and Systems 2011.
- [9] J. R. Matijevic et al., "Characterization of the Martian Surface Deposits by the Mars Pathfinder Rover, Sojourner," *Science*, vol. 5344, no. 278, pp. 1765-1768, December 1997.
- [10] NASA. (2012, May) Spirit and Opportunity. [Online]. http://www.nasa.gov/mission_pages/mer/opportunity-update.html
- [11] R. Li et al., "Initial Results of Rover Localization and Topographic Mapping for the 2003 Mars Exploration Rover Mission," *Photogrammetric Engineering & Remote Sensing*, vol. 71, no. 10, pp. 1129-1142, October 2005.
- [12] M. Maimone et al., "Two years of visual odometry on the mars exploration rovers: Field reports," *J. of Field Robotics*, vol. 24, no. 3, pp. 169-186, July 2006.
- [13] D. Wettergreen et al., "Design and field experimentation of a prototype Lunar prospector," *Int. J. of Robotics Research*, vol. 29, no. 12, pp. 1550-1564, October 2010.
- [14] S. Lacroix and A. Mallet, "Integration of concurrent localization algorithms for a planetary rover," in *Proc. of 6th Int. Symp. on Artificial Intelligence, Robotics and Automation in Space*, St-Hubert, Quebec, June 2001.
- [15] R. Volpe et al., "Enhanced Mars Rover Navigation Techniques," in *Proc. of the IEEE Int. Conf. on Robotics and Automation*, San Francisco, CA, 2000.
- [16] S. Lacroix et al., "Rover self localization in planetary-like environments," in *Proc. of 5th Int. Symp. on Artificial Intelligence, Robotics and Automation in Space*, Noordwijk, The Netherlands, June 1999, pp. 433-440.
- [17] F. Xu, "Automation in Mars Landing-Site Mapping and Rover Localization," in *International Society for Photogrammetry and Remote Sensing (ISPRS) Congress*, Istanbul, Turkey, 2004.
- [18] C. F. Olson and H. Abi-Rached, "Wide-Baseline Stereo Experiments in Natural Terrain," in *Proc. of 12th Int. Conf. on Advanced Robotics*, Seattle, WA, 2005, pp. 376-383.
- [19] T. Kubota et al., "Map Matching Scheme for Position Estimation of Planetary Explorer in Natural Terrain," in *Proc. of the IEEE Int. Conf. on Robotics and Automation*, Rome, Italy, April 2007, pp. 3520-3525.
- [20] F. Cozman et al., "Outdoor visual position estimation for planetary rovers," *Autonomous Robots*, vol. 9, no. 2, pp. 135-150, September 2000.
- [21] P. J. F. Carle and T. D. Barfoot, "Global rover localization by matching lidar and orbital 3D maps," in *Proc. of the IEEE Int. Conf. on Robotics and Automation*, Anchorage, AK, 2010, pp. 881-886.
- [22] T. Barfoot et al., "Field testing of a rover guidance, navigation, and control architecture to support a ground-ice prospecting mission to Mars," *Robotics and Autonomous Systems*, vol. 59, no. 6, pp. 472-488, June 2011.
- [23] B. Kitt et al., "Visual Odometry based on Stereo Image Sequences with RANSAC-based Outlier Rejection Scheme," in *Proc. of the IEEE Intelligent Vehicles Symposium*, San Diego, CA, June 2010, pp. 486 - 492.
- [24] R. B. Rusu and S. Cousins, "3D is here: Point Cloud Library (PCL)," in *Proc. of the IEEE Int. Conf. on Robotics and Automation*, Shanghai, China, 2011.
- [25] NASA Marshal Space Flight Center, "Lunar surface models," NASA Space Vehicle Design (Environment) NASA-SP-8023, 1969.
- [26] Blender Foundation. (2011) Blender 2.59. [Online]. <http://www.blender.org/>
- [27] W. E. Larson et al., "RESOLVE: Ground Truth for Polar Lunar Volatiles as a Resource," in *62nd International Astronautical Congress*, Cape Town, South Africa, October 2011.