Ray Tracing and Use of Shadows as Features for Determining Location in Lunar Polar Terrain

Eugene Fang and William "Red" Whittaker

Abstract Ice is the most valuable resource on the Moon. It exists only at the poles where shadows are extensive and drivable routes are short. Robot routes to reach this ice are tenuous. Sun-synchronous lunar polar routes offer order-of-magnitude greater duration and range, if such routes are achievable. Sun-synchrony is brittle in the sense that a rover must be at precisely scheduled time and place, so special localization techniques are warranted. Methods for terrain-based localization that work at equatorial regions are challenged at the lunar poles, where the grazing sunlight casts long shadows that obscure and change views over time. The shadows are shown here to accentuate craters as localization features. This paper presents a method that improves terrain registration at the poles of the Moon by probabilistically considering sensor and terrain uncertainty, and exploiting shadows as semantic features for localization. This method is validated and evaluated in simulated experiments.

1 Introduction

Solar powered rovers that explore the poles of the Moon for water ice and other volatiles will rely on sun-synchronous routes to extend a mission from weeks to months [15]. These spatio-temporal routes that chase sunlight and avoid shadows are highly brittle and require accurate localization to ensure rovers are in the right place at the right time. Compared to at equatorial regions, terrain relative navigation is challenged at the lunar poles. Long sweeping shadows preclude views of unlit regions and change the appearance at any location over time as the Sun clocks the diurnal cycle. This causes large errors in current methods for terrain registration, as

Eugene Fang · William "Red" Whittaker
Robotics Institute, Carnegie Mellon University, 5000 Forbes Ave. Pittsburgh, PA 15213
e-mail: {fangeugene, red}@cmu.edu

true rover positions may be interpreted as false negatives due to changing views and appearances.

This research develops techniques to improve accuracy of terrain relative localization at the lunar poles by considering both near-field and far-field geometry, accounting for sensor and terrain uncertainty, and exploiting shadows as semantic features for registration in an otherwise colorless and homogeneous environment. The true geometry and appearance of terrain influences both the observations from rover stereo imagery as well as orbital 3D models and imagery. Maximizing the likelihood of a rover observation given an expectation based on orbital data yields an estimate of position.

This paper presents a ray tracing method for registering rover stereo imagery to digital elevation models that considers near-field and far-field geometry and takes into account sensor and terrain uncertainty. This method is then augmented to incorporate shadows as features for localization. Section 2 discusses related work in geometric and appearance based terrain registration. Section 3 presents a ray tracing approach for geometric terrain registration. An augmentation of the method to exploit shadows as features is described in Section 4. Experiments and results are presented in Section 5. Section 6 discusses conclusions and directions for future research.

2 Related Work

Common approaches for registering rover imagery to DEMs utilize skyline matching [17, 4]. These methods extract the skyline – the curve where the horizon meets the sky – from a rover image (usually a panorama). Next, the skyline is compared to simulated skylines rendered at different positions within a digital elevation model. The similarity between the detected and simulated skyline provides a likelihood of the rover being at each position.

Accuracy of skyline based approaches is limited by the effects of parallax. When skyline features are distant – typical for planetary exploration rovers that operate on benign terrain far away from hills and other significantly 3-dimensional terrain features – positions hundreds of meters apart may have very similar skylines. Furthermore, these approaches are especially susceptible to the presence of shadows on the Moon. The black sky due to the lack of atmosphere can be difficult to differentiate from dark shadows (Figure 1), resulting in an incorrectly detected skyline. Current skyline matching methods do not account for this, and localization accuracy greatly degrades as a result.

Other approaches for terrain registration work by aligning local 3D models generated by onboard sensors to 3D orbital maps. Common variations of this method typically generate local rover point clouds from stereo imagery or LIDAR, downsample them to match the resolution of orbital DEMs, and then find their best alignment to orbital DEMs using variations of the Iterative Closest Point (ICP) algorithm [20, 12] or rock feature extraction [5]. A limitation of stereo methods, however,

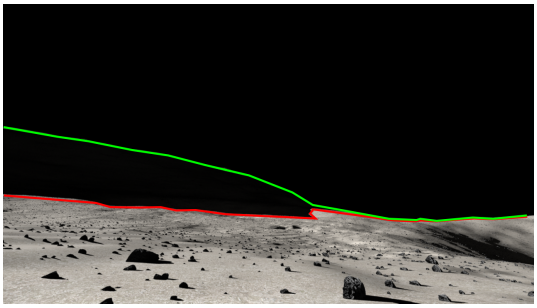

Fig. 1: A rover-perspective render near Shackleton crater on the Moon illustrates how skyline matching approaches for terrain registration are easily confounded by shadows on the Moon. The dark sky is difficult to differentiate from skyline features in shadow, resulting in a detected skyline (red) drastically different from the true skyline (green).

is the direct relationship between distance and uncertainty in distance. As a result, these methods discard portions of point clouds outside a given radius around the rover where the uncertainty in depth due to disparity error exceeds the resolution of the DEM [13]. However, discarding far-field information reduces the capability of distinguishing the difference between locations that look locally similar and also reduces the size of the basin of attraction for localization.

Methods that use LIDAR or other time-of-flight sensors do not suffer as much from the same issue of uncertainty increasing with distance, so far-field points can also be used. One such method, Multi-frame Odometry-compensated Global Alignment, extracts features from local and global 3D models such as topographic peaks [2] for registration. However, it relies on a high-power surveying-grade LIDAR for centimeter-accurate range measurements up to 1.5 km, which is impractical for non-terrestrial applications where power and mass budgets are much stricter.

Automated registration of photos to prior orbital imagery has been deeply explored for terrestrial aerial vehicles [19, 14] and precise Mars and lunar landings. [10, 11]. Many of these methods are all not well suited for surface rovers, where photos taken on the surface differ vastly in scale and have different spatial frequency content from those taken from orbit. Approaches that directly match vehicle images to orbital images require images of the same scale. Those that estimate position using SIFT features fail due to the different level of detail and spatial frequency content.

Some methods attempt to bridge the difference in scales by using a multi-resolution approach [7], but rely on a series of images taken over the length of a traverse in order to accurately localize. These suffer from issues when the terrain local to the rover is barren. Other methods overcome disparity in scale by only correlating salient features in images such rocks and craters [9, 3], but require that a sufficient number of features that are present and large enough to be visible in orbital imagery.

Even in optimal conditions, all aforementioned approaches suffer from dynamic illumination and shadows which change and obscure features. Synthetic images for registration based on orbital data and simulated lighting could be generated, but this idea has not been explored and the necessary accuracy and resolution of these synthetic images is unknown. Semantic features for localization offers benefits over low-level image features, which suffer from ambiguity and dependence on viewpoint [16]. However, these high-level features are only useful for terrain-relative localization if they can be reliably observed or predicted using orbital data.

This research is distinct from prior methods in two important ways. First, it incorporates both near-field and far-field geometry using a ray tracing approach for registering rover-centric stereo imagery to orbital terrain geometry that considers sensor noise and terrain uncertainty. Second, it exploits shadows as semantic features for localization, which handles differences in scale, frequency, and illumination. The developed method enables accurate map-based positioning at the lunar poles where current algorithms fall short.

3 Ray-Traced Stereo Scan Matching to Geometric Maps

This section presents a novel method for registering rover stereo images to 3D orbital maps that addresses many of the shortcomings of existing geometric terrain registration methods. Rays that correspond to pixels in stereo images are traced through a 3D map at a candidate pose to evaluate likelihood of observation based on occupancy. The sum of log likelihoods is then maximized through optimization to find the most likely pose within the map given the observations. The method accounts for anisotropic stereo noise to incorporate both near-field and far-field data, considers uncertainty in elevation maps, and explicitly models terrain free-space and occlusion.

3.1 Stereo Correspondence and Ray Construction

The method begins with a pair of undistorted and rectified stereo images captured by a rover. A block matching stereo correspondence algorithm [6] is run to obtain a disparity image (Figure 2).

The vectors $r_i = (x_i, y_i, z_i)$ that start at the camera origin and pass through each pixel (u_i, v_i) of the disparity image are computed using the camera principal point (c_x, c_y) and focal length f (Figure 3).

$$x_i = (u_i - c_x)/f$$

$$y_i = (v_i - c_y)/f$$

$$z_i = 1$$

(a) Left stereo image (b) Right stereo image (c) disparity image

Fig. 2: A disparity image (c) constructed from a left (a) and right (b) stereo image pair. Red indicates large disparity, blue indicates small disparity, white indicates zero disparity, and black indicates no stereo match.

These vectors, or *rays*, form the basis for individual observations that will be evaluated for the likelihood of occurrence.

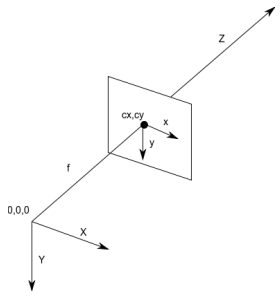

Fig. 3: Camera coordinate system

3.2 Camera Sensor Model Based on Disparity

Each traced ray from the disparity image (Figure 2c) comes from a pixel with either (a) non-zero disparity, (b) zero disparity, or (c) no stereo match due to poor features or lack of features. To calculate depth and uncertainty for each type of observation, they must be handled separately.

3.2.1 Non-Zero Disparity

Pixels with non-zero disparity typically make up the majority of a disparity image. For standard rectified stereo geometry, the inverse relationship between depth Z and disparity d is given by:

$$Z = \frac{fB}{d}$$

where f is the focal length and B is the stereo baseline [18]. Due to this non-linear inverse relationship, Gaussian error in pixel-space disparity results in non-Gaussian, asymmetric error in depth (Figure 4). A plot of the likelihood of occupancy at a depth for different observed stereo disparities is shown in Figure 5. Values of $f = 50px$, $B = 0.5m$, and $\sigma_{disparity} = 0.1px$ are chosen to exaggerate the effect of stereo disparity for illustrative purposes. Large disparities correspond to smaller depths with lower uncertainty. Small disparities correspond to larger depths with higher uncertainty.

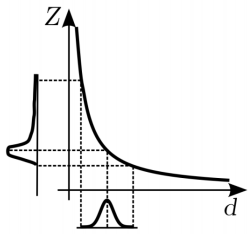

Fig. 4: Gaussian error in disparity d results in non-Gaussian, asymmetric error in depth Z. Image credit: [8]

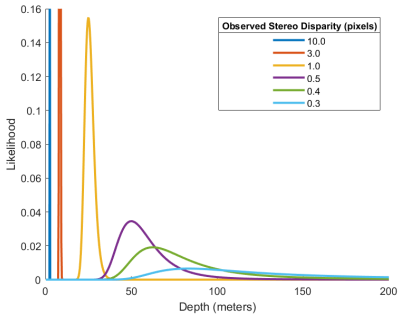

Fig. 5: As the observed stereo disparity decreases, the likely depth and uncertainty in depth both increase.

3.2.2 Zero Disparity

For many registration algorithms, such as those that rely on point clouds, pixels with zero stereo disparity cannot be used because they provide no information about depth. However, with ray tracing, they provide a lower bound for which the likelihood of occupancy at a depth – though possibly infinitesimally small – is nonzero. To overcome the issue of infinitesimally small likelihoods rounding to zero given the finite precision of computing, rays are only traced out to a finite range. Figure 6 shows an example of the likelihood of depth given an observation with zero stereo disparity up to a range of 2500 meters.

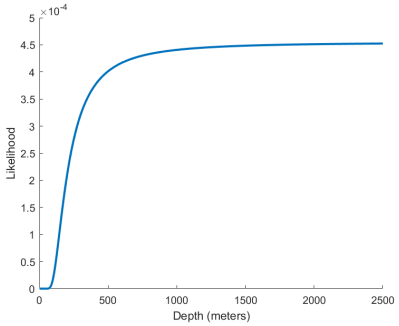

Fig. 6: The likely depth for an observation with zero stereo disparity. Range is limited to 2500 meters.

3.2.3 No Stereo Match

In cases where there is no stereo match due the lack of features, registration relies on semantic scene classification. This idea is explored in Section 4.

3.3 Observed Terrain Occupancy

In order for the evaluation of observation likelihood that accounts for occlusion, the raw likelihood of occupancy as a function of depth is not used. Instead, the cumulative likelihood of occupancy is computed (Figure 7). This gives the probability a observed ray with a particular stereo disparity is occupied before a given depth along that ray, which enables occlusion-aware observation likelihood maximization.

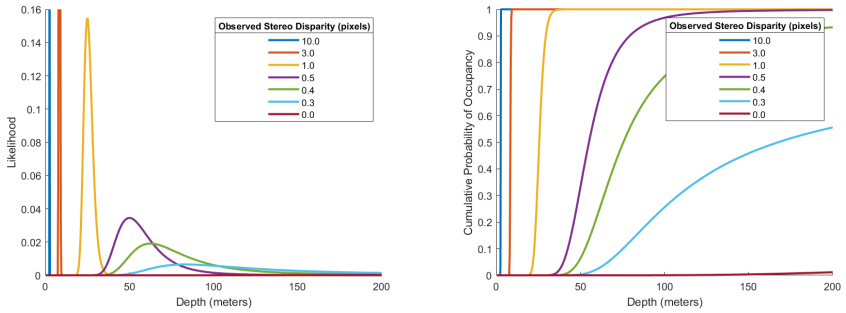

(a) Model of occupancy likelihood as a function of depth for several observed stereo disparities

(b) Model of cumulative likelihood of occupancy for several observed stereo disparities

Fig. 7: The probability that an observed stereo disparity is occupied before a given depth is equal to the cumulative likelihood of occupancy.

3.4 Expected Terrain Occupancy

To evaluate the likelihood of a ray observation, the observed and expected cumulative occupancy likelihood along the ray are compared. The expectation is computed by tracing the ray in an elevation map at a candidate pose (Figure 8). Due to the nature of ray tracing, computing the expected occupancy can incorporate uncertainty in the elevation map. Marching along a ray, the probability of occupancy is given by:

$$P(o_r|e_r) = cdf(o_r|\mu(e_r), \sigma(e_r))$$

where o_r is the elevation at range r, e_r is the cell in the elevation map that contains o_r, μ and σ are the mean and variance of the elevation measurement, and cdf is the cumulative probability distribution function.

3.5 Observation Likelihood Maximization

$P(o_i|e_i)$ is the likelihood of obtaining a single ray observation o_i given an expected observation e_i for ray i, and is calculated by computing the inverse absolute area between the observation and expectation curves (Figure 9).

Summing the log likelihood across all rays, yields the overall likelihood.

$$\mathscr{L} = \sum_i log(P(o_i|e_i))$$

Nelder-Mead optimization is used to maximize the likelihood over a candidate pose. If initialized close enough to the true pose, the estimate of the rover’s 3D pose

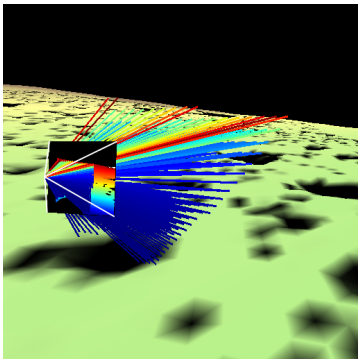

Fig. 8: Given a candidate pose in an elevation map, rays are traced within the elevation map to find the expected cumulative likelihood of occupancy along each ray.

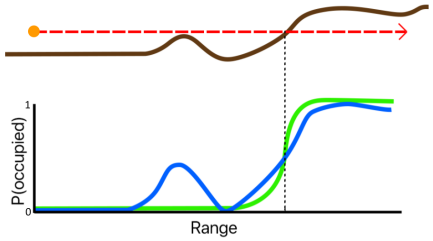

Fig. 9: A ray (dotted red) traced from a candidate pose (orange) intersects the terrain (brown). The observed probability of occupancy along the ray is shown in green, and the expected probability of occupancy along the ray is shown in blue (assuming some terrain uncertainty). The likelihood of the observation given the expectation is scored as the inverse absolute area between the two curves.

converges to the true pose. Experiments in Section 5 explores the performance of this optimization.

4 Detection & Registration of Shadow Features

Pure geometric methods for terrain registration fail when local topology is not distinct. Appearance-based terrain registration succeeds when geometric methods do not by correlating colors, textures, and other visual features to prior orbital imagery. However, appearance-based methods are severely limited at the lunar poles where there is no distinguishing color, textures depend on illumination, and visual features may be occluded by shadow. Furthermore, correlating photos taken by a surface rover to those taken by an orbiting satellite is generally challenging, since the pho-

tos differ vastly in vantage. These limitations are compounded by the constantly varying lighting and shadows.

Instead of being hindered by shadows, this research explores the idea of modeling them over time and exploiting them as features. This section presents a novel use of shadows as features for terrain registration that overcomes the vantage challenge of appearance-based registration by relying on semantic features rather than raw image features, seamlessly incorporates shadow registration into the ray traced stereo scan matching algorithm, and demonstrates its advantages in simulated experiments.

4.1 Shadow Detection

Semantic detection of shadows in an image provides features for terrain registration, and succeeds even in image areas with no features for stereo correspondence (Figure 10). The early approach for this research for automated detection of large shadows uses an adaptive threshold on intensity and then filters those by the size of connected components.

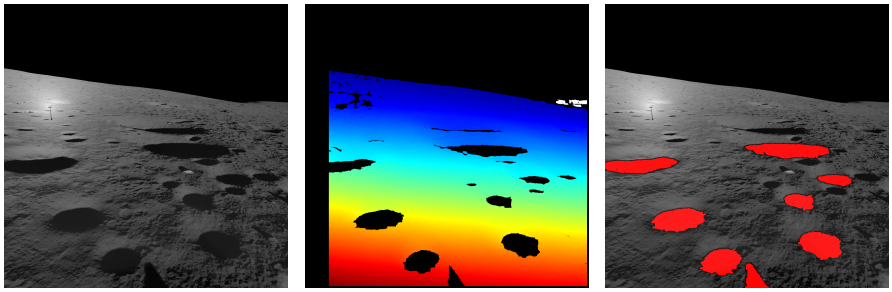

(a) Left stereo image

(b) Disparity image. Black areas have no stereo correspondence.

(c) Detected shadow features (red)

Fig. 10: Regions in stereo imagery (a) with no features for stereo correspondence (b) can still provide semantic features for registration (c).

4.2 Shadow Registration using Ray Tracing

Since the ray traced stereo scan matching algorithm described in Section 3 minimizes the log likelihood of of an observation o_i given an expected observation e_i, incorporation of shadow features into the registration process requires only a probability model $P(o_i|e_i)$ for each ray corresponding to a shadow. With the current naive

approach for shadow detection, o_i is binary. Thus, $P(o_i|e_i)$ is given by the probability that a shadow ray intersects a given patch of terrain multiplied by the probability of that patch of terrain being in shadow. Figure 11 shows an illustration of rays corresponding to shadows being traced within an elevation map.

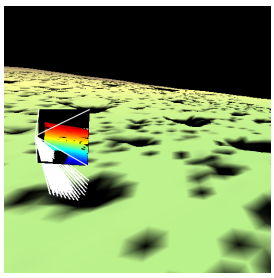

Fig. 11: Given a candidate pose in an elevation map, rays corresponding to semantically detected shadows (white) are traced within the map to find the likelihood of the ray intersecting an area of terrain in shadow.

5 Experiments & Results

Ray traced scan matching for terrain registration was evaluated against other methods in a photo-realistic lunar simulator developed by NASA Ames Research Center and Open Robotics [1] (Figure 12). The simulator covers an area of 1.2x1.2 km and uses a DEM by Nobile Crater near the south pole of the Moon. Fractal expansion was used to upsample the DEM to a resolution of 4cm/pixel in a small 160x160 meter region at the center, and gradually decreasing resolution outwards. The rover's stereo cameras have a resolution of 1024x1024 pixels, a 70 degree field of view, and a stereo baseline of 40 cm.

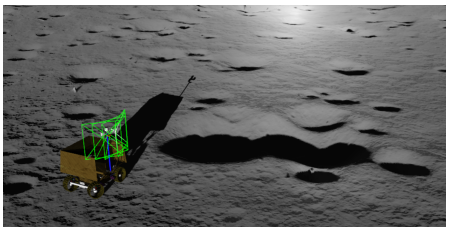

Fig. 12: A screenshot from NASA's planetary rover simulator for lunar missions showing a rover, its cameras' field of views, and the surrounding terrain.

The simulated rover was commanded to take an overlapping panorama of six pairs of stereo images centered 15 degrees below the horizon. Different terrain registration algorithms used these images as input and were randomly initialized at 10 locations up to 10 meters away from the true position. After attempting registration from each of these initialization conditions, the residual error with the best convergence score was then recorded. This process was repeated over a 10 meter grid within the high-resolution 160x160 meter portion of the simulated environment, and the results are shown in Figure 13 and Figure 14). On average, compared to ICP, the developed methodology decreases the average error from 0.5 m to 0.15 m. and decreases the maximum error from 3.10 m to 0.71 m.

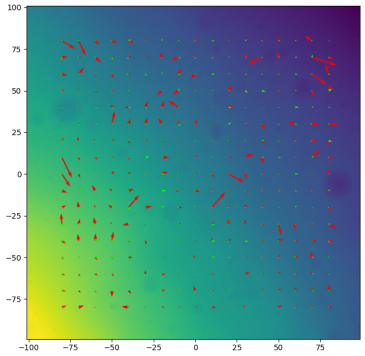

Fig. 13: The residual localization error (scaled 5x for clarity) at 289 locations within the simulated environment are compared. Point-to-plane ICP (red) has significantly larger errors than ray traced scan matching (green).

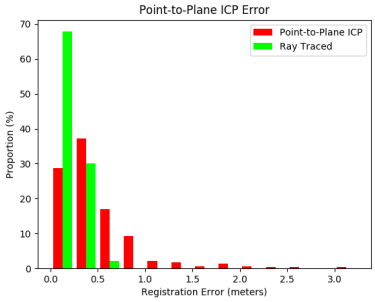

Fig. 14: In this particular dataset, ray traced scan matching (green) decreases both the mean and maximum registration error when compared to point-to-plane ICP (red).

Geometry-only and geometry-and-shadow ray traced stereo scan matching were compared at 289 locations within the same 160x160 meter region of the lunar simulator. For each method at each location, the likelihood field on a 20x20 meter grid around the true position in x-y are computed. For reliable optimization, the desired shape of these likelihood fields is a continuous, steep slope towards the true optima. Figure 15 compares the likelihood fields of geometry-only and geometry-and-shadow scan matching over the entire 160x160 meter region with a fixed color scale. Qualitatively, the addition of registering shadow features seems to decrease positional uncertainty overall. Quantitatively, by computing the volume under the

likelihood fields, the addition of shadows decreases the volume by 26% on average, and never increases it.

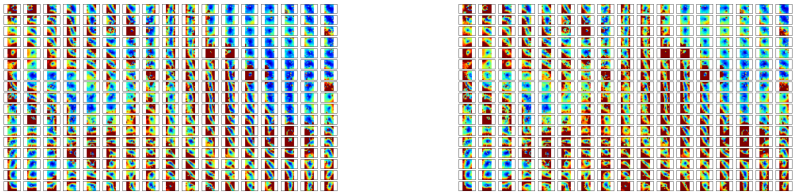

Fig. 15: Geometry-only (left) and geometry-and-shadow (right) likelihood fields around true positions for 289 locations in a 160x160 meter region of the lunar simulator. Color scales fixed, with blue as high-likelihood and red as low-likelihood. Overall, the addition of shadows decreases positional uncertainty and never increases it.

Using a single core on a multi-core Intel i7 laptop processor from 2011, ray tracing registration took on average 2.9x longer than ICP to converge (19.2 vs. 6.6 seconds) for each initialization in these experiments. Despite the comparatively worse performance, this is still reasonable for lunar polar sun-synchronous routes that can require maximum driving speeds as low as 1 centimeter per second [15].

6 Conclusion and Future Work

Sun-synchronous routes could be the key to lunar water if robots could reliably and autonomously navigate them. These routes are delicate, and requires certainty of location at given times. This research developed an approach to accurately localize to the terrain at the lunar poles. In contrast to priors, this method utilizes near-field and far-field geometry and effectively utilizes shadows as features for registration, yielding a reduction in maximum registration error by a factor of four compared to ICP in simulated experiments when initialized within 10 meters of the true position. Future work will analyze the efficacy of this method with real field data.

Acknowledgements This work was partially supported by the NASA Space Technology Research Fellowship under grant NNX16AM68H.

References

1. Mark Allan, Uland Wong, P Michael Furlong, Arno Rogg, Scott McMichael, Terry Welsh, Ian Chen, Steven Peters, Brian Gerkey, Moraan Quigley, et al. Planetary rover simulation for lunar exploration missions. In *2019 IEEE Aerospace Conference*, pages 1–19. IEEE, 2019.
2. Patrick JF Carle, Paul T Furgale, and Timothy D Barfoot. Long-range rover localization by matching lidar scans to orbital elevation maps. *Journal of Field Robotics*, 27(3):344–370, 2010.
3. Yang Cheng, Andrew E Johnson, Larry H Matthies, and Clark F Olson. Optical landmark detection for spacecraft navigation. 2003.
4. Fabio Cozman, Eric Krotkov, and Carlos Guestrin. Outdoor visual position estimation for planetary rovers. *Autonomous Robots*, 9(2):135–150, 2000.
5. Kaichang Di, Zhaoqin Liu, and Zongyu Yue. Mars rover localization based on feature matching between ground and orbital imagery. *Photogrammetric Engineering & Remote Sensing*, 77(8):781–791, 2011.
6. Pedro F Felzenszwalb and Daniel P Huttenlocher. Efficient belief propagation for early vision. *International journal of computer vision*, 70(1):41–54, 2006.
7. G Foil, C Cunningham, DS Wettergreen, and WL Whittaker. Onboard detection and correction of orbital registration errors using rover imagery. 2014.
8. Daniel Gutiérrez-Gómez, Walterio Mayol-Cuevas, and Josechu J Guerrero. Inverse depth for accurate photometric and geometric error minimisation in rgb-d dense visual odometry. In *2015 IEEE International Conference on Robotics and Automation (ICRA)*, pages 83–89. IEEE, 2015.
9. Ju Won Hwangbo, Kaichang Di, and Rongxing Li. Integration of orbital and ground image networks for the automation of rover localization. In *ASPRS 2009 Annual Conference*, 2009.
10. Andrew E Johnson, Yang Cheng, James F Montgomery, Nikolas Trawny, Brent Tweddle, and Jason X Zheng. Real-time terrain relative navigation test results from a relevant environment for mars landing. In *AIAA Guidance, Navigation, and Control Conference*, page 0851, 2015.
11. Andrew E Johnson and James F Montgomery. Overview of terrain relative navigation approaches for precise lunar landing. In *2008 IEEE Aerospace Conference*, pages 1–10. IEEE, 2008.
12. Ara V Nefian, Xavier Bouyssounouse, L Edwards, Taemin Kim, Emily Hand, Jared Rhizor, Matthew Deans, George Bebis, and Terrence Fong. Planetary rover localization within orbital maps. In *2014 IEEE International Conference on Image Processing (ICIP)*, pages 1628–1632. IEEE, 2014.
13. AV Nefian, LJ Edwards, D Lees, L Keely, TJ Parker, and M Malin. Automatic rover localization in orbital maps. In *Lunar and Planetary Science Conference*, volume 48, 2017.
14. Jaehong Oh, Charles K Toth, and Dorota A Grejner-Brzezinska. Automatic georeferencing of aerial images using stereo high-resolution satellite images. *Photogrammetric Engineering & Remote Sensing*, 77(11):1157–1168, 2011.
15. Nathan D Otten. *Planning for Sun-Synchronous Lunar Polar Roving*. PhD thesis, 2018.
16. John G Rogers, Alexander JB Trevor, Carlos Nieto-Granda, and Henrik I Christensen. Simultaneous localization and mapping with learned object recognition and semantic data association. In *2011 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 1264–1270. IEEE, 2011.
17. Fridtjof Stein and Gerard Medioni. Map-based localization using the panoramic horizon. *IEEE Transactions on Robotics and Automation*, 11(6):892–896, 1995.
18. Richard Szeliski. *Computer vision: algorithms and applications*. Springer Science & Business Media, 2010.
19. Alexander Wong and David A Clausi. Arrsi: Automatic registration of remote-sensing images. *IEEE Transactions on Geoscience and Remote Sensing*, 45(5):1483–1493, 2007.
20. Mark Woods, Andrew Shaw, Estelle Tidey, Bach Van Pham, Lacroix Simon, Raja Mukherji, Brian Maddison, Gary Cross, Aron Kisdi, Wayne Tubby, et al. Seekerautonomous long-range rover navigation for remote exploration. *Journal of Field Robotics*, 31(6):940–968, 2014.