

Complementary Flyover and Rover Sensing for Superior Modeling of Planetary Features

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

Abstract This paper presents complementary flyover and surface exploration for reconnaissance of planetary point destinations, like skylights and polar crater rims, where local 3D detail matters. Recent breakthroughs in precise, safe landing enable spacecraft to touch down within a few hundred meters of target destinations. These precision trajectories provide unprecedented access to bird’s-eye views of the target site and enable a paradigm shift in terrain modeling and path planning. High-angle flyover views penetrate deep into concave features while low-angle rover perspectives provide detailed views of areas that cannot be seen in flight. By combining flyover and rover sensing in a complementary manner, coverage is improved and rover trajectory length is reduced by 40%. Simulation results for modeling a Lunar skylight are presented.

1 Introduction

This paper presents complementary flyover and surface exploration for reconnaissance of point destinations, like skylights and polar crater rims, where local 3D detail matters (See Fig. 1). In contrast to past missions where regional characterization was the goal, missions to point destinations will detail local terrain geometry, composition, and appearance. Characterization of this type requires high density sampling and complete coverage. Standard rover-only approaches are inefficient and

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

Robotics Institute, Carnegie Mellon University, 5000 Forbes Ave, Pittsburgh, PA 15213,
e-mail: hlj@cs.cmu.edu, uyw@andrew.cmu.edu, kp@cs.cmu.edu, aashish.sheshadri@gmail.com, red@cs.cmu.edu

Jason Koenig

Computer Science Department, Carnegie Mellon University, 5000 Forbes Ave, Pittsburgh, PA 15213, e-mail: jrkoenig@andrew.cmu.edu

cannot generate the coverage required for complete 3D modeling. Complementary flyover and surface exploration meets the requirements for modeling point features with higher efficiency than alternative approaches.


Persistent light illuminates polar locations on the Moon and Mercury. These destinations could serve as bases of operations or power stations for exploitation of polar resources, but for polar destinations, even small rocks cast long shadows, and unexpected shadows can be mission-ending for small rovers. Precise knowledge of 3D structure on the meter-scale and smaller is needed to predict where shadows will fall.

Sub-surface caverns may harbor life on Mars. They may be the best hope for human habitation on the Moon. They can provide windows into a planet's past geology, climate, and even biology. Skylights, formed by partial cave ceiling collapse, provide access to sub-surface voids. They have been conclusively shown to exist on Mars [6] and the Moon [3], and evidence supports their existence on other planetary bodies throughout the solar system [2]. Surface robots can approach and scan skylight walls, but skylight geometry prevents viewing the hole floor from a surface perspective.

Orbiters currently in service around the Moon and Mars are generating higher resolution data than ever before, but there are limits to what can be done from orbital distances. Even with a very good laser, the Lunar Reconnaissance Orbiter (LRO) sees a 5m radius laser spot on the ground from its nominal 50km mapping orbit [17], limiting modeling precision. LRO's camera is higher resolution, at 0.5m per pixel for the 50km orbit [18]. Stereo processing can be used to create a 2m per post digital elevation map (DEM) from a pair of these images, but this only works for lit terrain. Skylights and polar craters contain terrain that is always in shadow. More detail, captured by flyover, is needed to see hazards on the scales that matter for robotic explorers.

New breakthroughs in terrain-relative navigation enable unprecedented precision in lander trajectory. This makes possible, for the first time, low-altitude lander flyover exploration of point targets. Precise, safeguarded landing can be achieved with real-time data from cameras and LIDAR (LIght Detection And Ranging), enabling a lander to identify a safe landing location and maneuver past hazards to safely touch down. Flyover data can further inform subsequent rover exploration for effectiveness, safety and coverage not possible in traditional missions with multi-

Fig. 1 Complementary flyover and surface modeling concept: a lander captures views of a terrain feature during final descent flyover. A rover carried by the lander returns to examine the feature in more detail.



kilometer landing ellipses. The combination of two perspectives, flyover birds-eye and rover on-the-ground, enables construction of the high-quality models needed to plan follow-on skylight exploration and science missions or develop detailed shadow prediction for crater rims. This paper presents a simulation of combined lander and rover modeling of a Lunar skylight. A comparison is made between a model built with lander data only, a model built with rover data only, and a model built from combining lander and rover data, in which the rover views are chosen based on holes in the lander model.

Section 2 discusses related work in planetary exploration and “next best view” modeling. Section 3 discusses the approach to complementary flyover and surface modeling for point features where 3D detail matters. Specifics of the experiments conducted are presented in 4. Results are presented in 5. Sections 6 and 7 discuss conclusions and directions for future research.

2 Related Work

Modeling and localization are closely related: the robot location when a given frame of data was captured must be known to fit that data accurately into a model, and the most accurate localization estimate is often produced by building a model from multiple frames of data. Maps and 3D models of terrain have been created from a combination of orbiter, lander and rover imagery and used for rover localization, but not in a fully autonomous manner, and not for planetary features where 3D really matters. For the Mars Exploration Rovers (MERs), the DIMES system took three images of the landing site at about 1000m altitude during descent, aiming to determine the lander motion [11]. The MERs computed visual odometry onboard, although the computation was quite slow at 2 minutes per frame [11]. Visual odometry estimates of rover motion were more accurate than wheel odometry due to wheel slip, but position estimates still drifted over time, so bundle adjustment was performed on Earth to improve estimates of rover position. Tie points were selected automatically within a stereo image pair or panorama, and in some cases across different rover positions. DIMES imagery from the lander and HiRISE orbital imagery was used in localizing the rover and building maps, but the registration between rover and overhead imagery was done manually [12]. While the models built by MER provide a fascinating glimpse of Martian terrain, they do not take on point features with geometries that severely restrict visibility. Victoria Crater is perhaps the closest - it has been modeled from orbit and investigated extensively by the Opportunity rover [20, 10], but at 750m across and 75m deep, Victoria Crater is not a point feature, and does not have visibility-restricting geometry. In contrast, the Marius Hills Hole, a lunar skylight, is estimated to be 48 to 57m in diameter and approximately 45m deep [3]. See Fig. 2 for an example of how skylight geometry prevents viewing the floor from a surface perspective.

The MER waypoints were chosen by operators on Earth, but significant work done in autonomous mapping and modeling can be leveraged to automate this part


of the process. Work on laser scanning of unknown objects has used a “next best view” approach, choosing the next position from which to scan based on the amount of new information gained while maintaining overlap with existing data to facilitate model building [16]. This approach has also been applied to the robotic exploration of unknown environments [14].

Kruse, Gutsche and Wahl present a method for planning sensor views to explore a previously unknown 3D space [9]. This space is represented by a 3D grid, and each voxel in this grid is marked as either occupied, free or unknown. The value of a given view is evaluated by estimating the size of the unknown regions that become known after the measurement and determining the distance between that view and the current position in robot configuration space. The estimation of size for the unknown regions that can be seen in a given view is done using ray tracing, with a relatively small number of rays to limit computation time. This value function is re-evaluated after each view. The next view is chosen by following the gradient of the value function, starting from the current configuration. If the value function drops below a threshold, the gradient search is re-started from the best of a randomly chosen set of configurations.

Sawhney, Krishna and Srinathan use amount of unseen terrain visible and distance to determine the next best view for individuals in a multi-robot team. They find that the metric computed as (amount of unseen terrain)/distance is the most successful out of several evaluated [19].

Hollinger et al. use uncertainty to plan sensor views for a ship inspection robot [8]. They use a Gaussian process to model the surface of the ship hull. Because the shape of the ship is relatively well known before inspection, the approach assumes there will not be large changes to the model surface. This assumption would not hold in a skylight exploration case when it cannot be determined from the prior model whether a region inside the skylight is void space or collapsed ceiling.

Fig. 2 Skylight geometry restricts visibility from a rover perspective. Blue cone shows example of visible area from a rover positioned at the skylight edge.



3 Complementary Flyover and Surface Modeling Approach

3.1 Overview

This work combines lander flyover and rover exploration data to autonomously model point destinations where 3D detail matters. Lander and rover use both cameras and active sensors, such as LIDAR. Active sensing is needed to peer into shadowed regions, but active sensors are range-limited by available power and lack the high resolution of cameras.

Satellite imagery is used for terrain relative navigation, enabling landers to precisely position themselves as they fly over the features of interest. This technology enables landers to fly within 30m of their intended trajectory within the final 500m of descent and model regions on order of 50m across from very low altitude. Hazard detection and avoidance technology, combined with precise navigation, enables safe and autonomous landings near features even without guaranteed-safe zones of landing-ellipse size.

Rover modeling begins at the lander location, providing a common tie-point between surface and flyover models. On-board hazard detection and avoidance ensure safety as a rover moves. Rover paths and sensor views can be autonomously chosen, using a “next best view” approach, to fill holes in a lander model.

Lander flyover captures detailed overview data, as well as perspectives that cannot be observed from a rover viewpoint. Rovers can capture close-up images of the terrain, and they can linger to capture multiple views from stationary locations, though always from low, grazing perspectives. Alternately, landers can acquire bird’s-eye views but with less detail and resolution since their one-pass, always-moving trajectories are constrained by fuel limitations. Lander and rover data are combined, using lander data to localize and plan rover paths, to autonomously construct quality 3D models of point destinations.

3.2 Lander and Rover Trajectory and Sensing

For complementary flyover and surface modeling, the portion of the lander trajectory of interest is the final 500 m of descent. By this point, the lander has already canceled most of its forward velocity. It pitches over to a vertical orientation and cancels gravity to maintain a constant velocity. The lander points its sensors toward the feature of interest. After passing over the feature, the lander uses its LIDAR to detect hazards and follows a trajectory to avoid detected hazards in the landing zone. Above its target landing point, it cancels the rest of its forward velocity and descends straight down.

There is a trade-off between time to capture data and fuel used: flying slowly over a feature leaves more time to capture data but requires more fuel to maintain altitude for a low flyover; flying quickly over the feature saves fuel but may result in

sparse data coverage. With this architecture, sparse data collected by the lander can be filled in by the rover, enabling the lander to move quickly and save fuel.

For complementary flyover and surface modeling, rover trajectories are described by waypoints. Desired view angles at those waypoints are specified. A rover drives from one waypoint to the next and stops to capture desired views at each waypoint.

3.3 Complementary Flyover and Surface Modeling

Following lander flyover, a point cloud model built from lander data is binned into a voxel array. Voxels which contain at least one point are marked as occupied, and all others as unseen. Ray tracing is then done from each camera perspective, and all voxels that a ray passes through before it hits an occupied voxel are marked as seen and empty. The Amanatides and Woo algorithm is used to determine ray-voxel intersections [1]¹.

The area of interest is gridded into possible rover positions. For each possible position, a discrete set of possible views are evaluated. The voxel array with seen, unseen and empty voxels is used to predict unseen areas that can be observed from each rover view. The list of previously unseen voxels visible from each view is stored, as well as the total number that can be seen in all of the views from a given position.

A metric is then computed as (number of unseen voxels visible)/(distance from current position). The distance is computed along a straight line rover path. The position with the highest value for this metric is the next rover position. The voxels that were predicted to be visible in views from the new rover position are then marked as seen, and the metric is recomputed. This is repeated until there are no rover positions for which previously unseen voxels are visible. This approach does not always result in a particularly efficient path between waypoints. Given a set of rover positions with planned views which cover the space of visible but as-yet unseen voxels, the order in which those positions are visited can be changed without affecting the total number of as-yet unseen voxels observed, so the positions were re-ordered using a greedy algorithm to pick the next-closest position among the set of selected rover positions.

¹ The code for this operation is a vectorized version of an implementation by Jesus Mena-Chalco, available on MATLAB Central: <http://www.mathworks.com/matlabcentral/fileexchange/26852>.

4 Experiments

4.1 Overview

A mission to a Lunar skylight is examined as an archetypical example of the complementary flyover and surface modeling approach. Lander and rover data for skylight modeling are collected in simulation. Three mission scenarios for modeling a skylight are evaluated: 1) lander alone, 2) rover alone, 3) lander and rover, where lander model informs rover path. Case 1 mimics the architecture in which there is no rover, and the lander is the sole robotic explorer. Case 2 mimics the case in which the lander does not attempt to fly over or near the feature of interest to capture overhead views. A rover travels from the landing site to the feature of interest and is the only robotic explorer to view the feature. Case 3 follows the architecture presented in section 3, autonomously selecting the rover path to improve the model generated by the lander. Data for cases 1 and 2 were generated independently. For case 3, the lander-only model from case 1 was used to determine the rover path and views. The data for this alternate rover path was then generated and combined with the case 1 model.

The simulated skylight in this work is 30m in radius and approximately 40m deep. Modeling is restricted to a target area 100m square, centered on the skylight, although sensor views may go beyond this area.

Models built using each of the three scenarios are compared using the metric of percent coverage. In these experiments, percent coverage is defined as the percent of voxels that are filled in the ground truth model that are also filled in the generated model.

For robotic planetary landers and ground vehicles, the size of a feature considered to be a hazard is on order of 0.2m to 1m [7][21][5]. For a terrain model to be valuable for robotic operations, the resolution should be in this range or better. In the experiment presented here, the voxelization of the lander model used in the complementary flyover and surface modeling method was 0.5m, and the models for each scenario were also discretized to 0.5m for comparison with ground truth. Higher resolution models can easily be built by using a smaller voxel size in the complementary flyover and surface modeling method.

4.2 Simulation

A 3D model of a lunar-like site containing a skylight was fabricated for this work. A 2m/post digital elevation map (DEM) of the Apollo 11 region, created using a stereo pairs from the Lunar Reconnaissance Orbiter, was used as the base terrain. Small craters and rocks were overlaid following a statistical distribution derived from Surveyor data[15]. A procedural model was used, so that generated features did not need to be explicitly stored and only details visible in a given image were processed.


A static 3D mesh of the skylight was manually modeled using Blender software [4], based on imagery and measurements of the Marius Hills skylight [2]. Terrain components were combined, detail textures were added, and simulated color and depth data were generated using Blender's built-in raytracer. Images were synthesized using a pinhole camera model with the appropriate field of view and resolution for each sensor. For LIDAR, depth channel instead of color channel information was used. Range data were transformed from the orthographic projection of the depth buffer to a perspective projection with rays from the LIDAR center.

4.3 Lander Modeling

For purposes of this study, it is assumed that the lander has a camera and a flash LIDAR mounted on a gimbal. The camera has a field of view (FOV) of 15° by 15° and a resolution of 1944x1944 pixels. The LIDAR also has a 15° FOV, but its resolution is 128x128 pixels. The optical centers of the camera and LIDAR are assumed to be co-located, which could be accomplished using a cold mirror.

The lander targets a constant glide slope of approximately 7.5 m/s along-track and 3.5 m/s down while scanning the landing zone and the skylight. After passing over the skylight it zeros its forward velocity and begins a vertical descent to the ground. The ground track follows a straight line across the center of the hole to the landing site. Due to the limited visibility that would occur when the sensors were mounted on a physical lander, it is assumed that the sensor views are restricted to look forward along the lander's direction of travel and to the sides - no backward-facing sensor views are allowed. Lander camera and LIDAR views are specified with the camera position and sensor view direction. The data from each view are first combined assuming that the lander position and orientation are known for each sensor view. LIDAR point clouds are stitched into a single model using ground truth lander poses, and the combined model is compared to the ground truth 3D hole model to determine coverage. In the case where the lander trajectory is not known, SIFT features [13] can be detected in each image, and the relative camera poses can be determined by matching features between images. Figure 4 shows an example of images stitched using this method. These poses could also be used to stitch together the LIDAR data, although that is not explored in this work.

Fig. 3 Bird's eye view of simulated terrain containing a skylight. Section shown is 600m long. Skylight diameter is 60m, depth is 40m.






Fig. 4: Lander images stitched together using SIFT features




Fig. 5: Simulated camera image from rover's-eye view, looking at the skylight edge

4.4 Rover Modeling

The rover is assumed to have a 45° FOV flash LIDAR with a resolution of 128×128 pixels. It is also assumed to have a pair of stereo cameras with 45° FOV and 1000×1000 pixel resolution. LIDAR and cameras are located on a pan-tilt head.

A naive rover path is selected to achieve full coverage of the skylight walls and surrounding terrain. The tilt angle of the rover sensors is set to achieve a desired resolution of data on the ground given the sensor resolution and FOV. The desired percentage overlap in sensor frames along the rover's path and between neighboring parallel paths is set to ensure that no gaps will be left between neighboring paths, and a resulting zig-zag pattern is generated. So long as the desired resolution can be achieved with the selected FOV, the FOV does not significantly affect the sensor coverage on the ground.

In order to maintain rover safety, there is a keep-out zone around the hole which the rover cannot enter. For this work, the width of the keep-out zone is 7.5m. When the rover's coverage pattern intersects the edge of the keep-out zone for the first time, the rover makes a circle of the skylight, aiming its sensors inward to capture the ground in the keep-out zone. As it travels around the skylight rim, it stops periodically to scan the opposing skylight wall. The frequency of these stops is determined by the portion of the far wall that would be visible given the sensor FOV, assuming the hole is cylindrical. If the sensor resolution on the far wall is predicted to be lower than the desired resolution, multiple overlapping frames of data are captured to achieve the desired resolution. Once the rover completes the circle, it continues its zig-zag pattern.

Data from each view are combined assuming that the rover position and sensor pointing direction are known for each sensor view, and ground truth values are used for these parameters. See Fig. 5 for an example rover camera image.

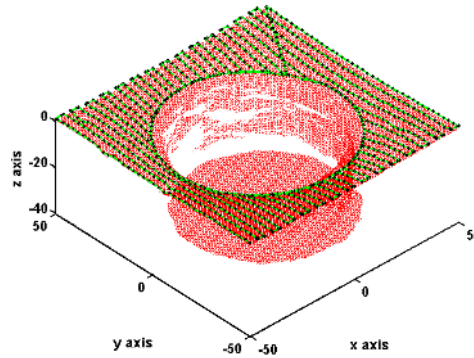


Fig. 6 Rover camera positions (black) and path (green) chosen naively to achieve full coverage, overlaid on voxel model of skylight. Axes are marked in meters.

4.5 Lander and Rover Modeling with Informed Path Selection

For this experiment, the method in section 3 is applied to the lander model from section 4.3. The rover is assumed to have camera and LIDAR with the FOVs and resolutions described in section 4.4. In this experiment, the grid of possible rover positions had squares of 5m and excluded the skylight and the keep-out zone. Sixteen views were evaluated at each rover position - 8 pan angles to cover a full circle with a 45° FOV, and 2 tilt angles for each pan angle. Rover paths between waypoints are straight lines, unless the straight-line path would intersect the skylight or the keep-out zone, in which case the path skirts the skylight until it can continue in a straight line toward the target waypoint.

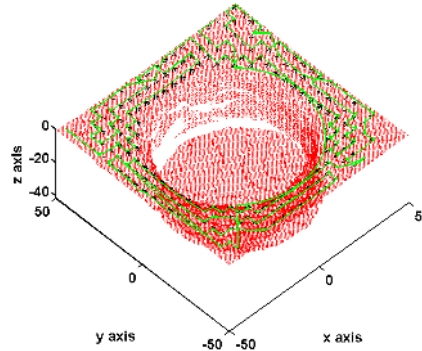


Fig. 7 Rover path planned using method from section 3, overlaid on voxel model of skylight




Fig. 8: Side and top views of point clouds constructed from (a) flyover LIDAR data, (b) rover LIDAR data, and (c) flyover and rover LIDAR data. The flyover-only model misses detail on the walls of the skylight. The rover-only model has dense wall coverage, but misses detail on the floor. Combined flyover and rover modeling provides dense coverage of the entire feature.

5 Results

5.1 Results for Lander Modeling

Figure 8a shows a model built from lander flyover LIDAR data. Given a voxelization of the model and the ground truth at 0.5m resolution, the coverage of this model is 46%.

5.2 Results for Rover Modeling

Figure 8b shows a model constructed from rover LIDAR data. These data were generated using the naive rover coverage pattern. Note how the skylight walls are densely covered, as is most of the terrain around the skylight, but the central portion of the floor cannot be seen by the rover. Given a voxelization of the model and the ground truth at 0.5m resolution, the coverage of this model is 85%. The naive rover path length is 2152m (See Fig. 6).

5.3 Results for Lander and Rover Modeling with Informed Path Selection

Figure 8c shows a model constructed from flyover and rover LIDAR data using the “next best view” approach for rover path planning. The lander model is used to generate the rover plan. Skylight walls and surrounding terrain are densely covered. The resulting model covers 92% of the ground truth with a rover path length of 1281m, a 40% reduction in path length compared to a the naive coverage pattern. Figure 7 shows the planned rover path and views.

6 Conclusions

This paper has presented approach and analysis of a new mission concept: complementary flyover and rover sensing for planetary features. In order to explore and model planetary terrain features, rovers will be delivered to the surface by landers. In complementary flyover and surface modeling, the lander is leveraged to serve a dual purpose as a flyover explorer, capturing 3D data at a much higher resolution than possible from orbit and obtaining views of the feature not possible from a rover perspective. By planning rover sensor views from this prior lander model, complementary flyover and surface modeling achieves greater coverage with a significantly shorter rover traverse.

Experiments performed for a simulated skylight reconnaissance mission demonstrate significant improvements in coverage and quality relative to a rover-only or lander-only mission. The model built by combining lander and rover data following the complementary flyover and surface modeling method achieves 92% terrain coverage, compared to 46% coverage for the lander alone and 85% coverage for a rover-only approach. The rover-only approach, using a path planned with very limited prior information about the skylight feature, has a path length of 2152m. The complementary flyover and surface modeling approach, by planning the rover path and views using lander data, is reduced to 1281m, a 40% reduction in distance traveled. Although not explored in this work, high resolution data from the lander can also be used to precisely localize the rover and identify hazards, further improving the efficiency of rover operations.

7 Future Work

Lander and rover positions were assumed known for this work, but in the future, accuracy of localization and effects of localization error will be investigated. If only relative navigation techniques are used and models are built purely from these navigation estimates, position will drift over time. This means, for example, that a longer

rover traverse will tend to result in a less accurate model. Model building can also contribute to localization, as was done for the MER rovers.

The effects of noise in the LIDAR data and in the camera and LIDAR commanded orientations will be investigated in future work. Noisy data would tend to reduce the accuracy of the model, and it could also reduce coverage.

Obstacle avoidance was not considered for these simulations. For a real mission, however, a rover would need to detour to avoid obstacles. These detours would change the path length and could provide obstructions to viewing certain areas of terrain that had been predicted to be visible. Additional methods of path planning for the rover views given the flyover data will also be explored.

Acknowledgements This work is partially supported by the NASA Innovative Advanced Concepts program under contract NNX11AR42G. The authors would also like to thank NVIDIA for donation of a machine that was used in the high-fidelity simulation of terrain done for this work.

References

- [1] Amanatides, J., Woo, A.: A fast voxel traversal algorithm for ray tracing. In: Proc. of Eurographics (1987)
- [2] Ashley, J.W., Boyd, A.K., Hiesinger, H., Robinson, M.S., Tran, T., van der Bogert, C.H., LROC Science Team: Lunar pits: Sublunarean voids and the nature of mare emplacement. In: LPSC. The Woodlands, TX (2011)
- [3] Ashley, J.W., Robinson, M.S., Hawke, B.R., Boyd, A., Wagner, R.V., Speyerer, E.J., Hiesinger, H., van der Bogert, C.H.: Lunar caves in mare deposits imaged by the LROC narrow angle camera. In: First Intl. Planetary Caves Workshop. Carlsbad, NM (2011)
- [4] Blender Foundation: Blender 2.59. www.blender.org (2011)
- [5] Braun, R.D., Manning, R.M.: Mars exploration entry, descent and landing challenges. In: IEEE Aerospace Conference (2006)
- [6] Cushing, G., Titus, T., MacLennan, E.: Orbital observations of Martian cave-entrance candidates. In: First Intl. Planetary Caves Workshop. Carlsbad, NM (2011)
- [7] Darms, M.S., Rybski, P.E., Baker, C., Urmson, C.: Obstacle detection and tracking for the Urban Challenge. IEEE Transactions on Intelligent Transportation Systems 10(3) (2009)
- [8] Hollinger, G.A., Englot, B., Hover, F., Mitra, U., Sukhatme, G.S.: Uncertainty-driven view planning for underwater inspection. In: ICRA (2012)
- [9] Kruse, E., Gutsche, R., Wahl, F.M.: Efficient, iterative, sensor based 3-d map building using rating functions in configuration space. In: ICRA. Minneapolis, Minnesota (1996)
- [10] Li, R., Di, K., Hwangbo, J., Chen, Y., Team, A.S.: Rigorous photogrammetric processing of HiRISE stereo images and topographic mapping at mars exploration rover landing sites. In: LPSC. League City, TX (2008)

- [11] Li, R., Squyres, S.W., Arvidson, R., Archinal, B., Bell, J., Cheng, Y., Crumpler, L., Des Marais, D.J., Di, K., Ely, T.A., Golombek, M., Graat, E., Grant, J., Guinn, J., Johnson, A., Greeley, R., Kirk, R., Maimone, M., Matthies, L.H., Malin, M., Parker, T., Sims, M., Soderblom, L.A., Thompson, S., Wang, J., Whelley, P., Xu, F.: Initial results of rover localization and topographic mapping for the 2003 Mars Exploration Rover mission. *Photogrammetric Engineering & Remote Sensing* 71(10), 1129–1142 (2005)
- [12] Li, R., He, S., Chen, Y., Tang, M., Tang, P., Di, K., Matthies, L., Arvidson, R.E., Squyres, S.W., Crumpler, L.S., Parker, T., Sims, M.: MER Spirit rover localization: Comparison of ground image–and orbital image–based methods and science applications. *JGR* 116 (2011)
- [13] Lowe, D.: Distinctive image features from scale-invariant keypoints. *IJCV* 60(2), 91–110 (2004)
- [14] Nagatani, K., Matsuzawa, T., Yoshida, K.: Scan-point planning and 3-d map building for a 3-d laser range scanner in an outdoor environment. In: *FSR* (2009)
- [15] NASA Marshall Space Flight Center: Lunar surface models. NASA space vehicle design criteria (Environment) NASA-SP-8023, NASA (1969)
- [16] Pito, R.: A sensor-based solution to the next best view problem. In: *Proc. Intl. Conf. on Pattern Recognition* (1996)
- [17] Riris, H., Cavanaugh, J., Sun, X., Liivia, P., Rodriguez, M., Neuman, G.: The Lunar Orbiter Laser Altimeter (LOLA) on nasa’s lunar reconnaissance orbiter (LRO) mission. In: *Intl. Conf. on Space Optics*. Rhodes, Greece (2010)
- [18] Robinson, M., Brylow, S., Tschimmel, M., Humm, D., Lawrence, S., Thomas, P., Denevi, B., Bowman-Cisneros, E., Zerr, J., Ravine, M., Caplinger, M., Ghaemi, F., Schaffner, J., Malin, M., Mahanti, P., Bartels, A., Anderson, J., Tran, T., Eliason, E., McEwen, A., Turtle, E., Jolliff, B., Hiesinger, H.: Lunar Reconnaissance Orbiter Camera (LROC) instrument overview. *Space Science Rev.* 150, 81–124 (2010)
- [19] Sawhney, R., Krishna, K., Srinathan, K.: On fast exploration in 2d and 3d terrains with multiple robots. In: *Proc. Intl. Conf. on Autonom. Agents and Multiagent Systems* (2009)
- [20] Squyres, S., Knoll, A., Arvidson, R., Ashley, J., Bell, J.I., Calvin, W., Christensen, P., Clark, B., Cohen, B., de Souza, P., Edgar, L., Farrand, W., Fleischer, I., Gellert, R., Golombek, M., Grant, J., Grotzinger, J., Hayes, A., Herkenhoff, H., Johnson, J., Jolliff, B., Klingelhofer, G., Knudson, A., Li, R., McCoy, T., McLennan, S., Ming, D., Mittlefehldt, D., Morris, R., Rice, J.J., Schroder, C., Sullivan, R., Yen, A., Yingst, R.: Exploration of Victoria Crater by the Mars rover Opportunity. *Science* 324(5930), 1058–1061 (2009)
- [21] Wagner, M., Heys, S., Wettergreen, D., Teza, J., Apostolopoulos, D., Kantor, G., Whittaker, W.: Design and control of a passively steered, dual axle vehicle. In: *Intl. Symposium on Artificial Intelligence, Robotics and Automation in Space* (2005)