Experimental Results of Autonomous Microrover Exploration and Mapping of an Analog Planetary Pit

Jordan Ford
Carnegie Mellon University
Pittsburgh, PA 15213
jsford@andrew.cmu.edu

Khaled Sharif
NASA Ames Research Center
Mountain View, CA 94035
khaled.sharif@nasa.gov

William L. "Red" Whittaker
Carnegie Mellon University
Pittsburgh, PA 15213
red@cmu.edu

Warren "Chuck" Whittaker
Carnegie Mellon University
Pittsburgh, PA 15213
warrenw@andrew.cmu.edu

Heather Jones
Carnegie Mellon University
Pittsburgh, PA 15213
hjones@andrew.cmu.edu

Uland Wong
NASA Ames Research Center
Mountain View, CA 94035
uland.wong@nasa.gov

Abstract - **Hundreds of pits on the Moon and Mars expose unique unweathered geology on their walls, and some may access habitable caves. Small, autonomous rovers could explore, image, and map planetary pits from the rim, capturing close-range, low-angle, long-exposure imagery that cannot be captured from orbit. The small but high fidelity models must be processed in-situ for the short-duration small missions of our time with limited bandwidth since the raw imagery required for modeling is at least an order of magnitude greater than what can be downloaded in a lunar day. The vast raw imagery is the basis for modeling pits with the accuracy, resolution and lighting corrections necessary for morphology study and search for cavernous openings. This work develops and evaluates several component technologies for these explorations. A safe-approach behavior is developed and employed to navigate pit rims to vantages for acquiring images. Images from multiple vantages are fed to a specialized incremental pit modeling pipeline that computes the high-resolution 3D pit models. The technologies are manifested aboard a prototype microrover and demonstrated at a terrestrial pit having size and shape comparable to known planetary pits. In multiple end-to-end mission simulations, these techniques are shown to reliably produce unprecedented, accurate, high-coverage pit models.**

Table of Contents

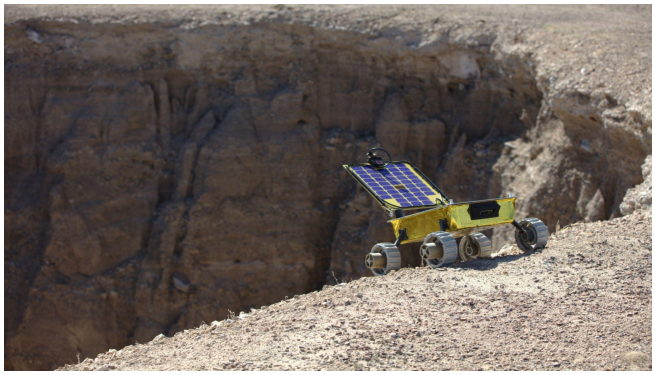

Figure 1. An autonomous microrover approaches a pit rim, images the interior, and constructs a high resolution 3D model of a planetary pit.

1. Introduction

Planetary pits are compelling exploration destinations that could change the future of human presence on the Moon and Mars and reshape our understanding of their pasts Visible pristine bedrock exposed on pit walls can reveal the only planetary geologic timelines that are unobscured by space weathering processes [1]. Rarer pits could provide access to sheltered alcoves or caves where long-term habitats could be established and protected from cosmic radiation, extreme temperatures, and micrometeorite impacts [2].

To-date, orbiters have discovered hundreds of pits on the Moon, Mars, Venus, Phobos, Eros, Gaspra, Ida, Enceladus, and Europa [3]–[5]. The Lunar Reconnaissance Orbiter alone has revealed more than 300 pits on the surface of the Moon [3], but sufficiently resolved grazing views of pit walls are challenges, and many views are obscured by shadow. The long-range, short-exposure imagery available from orbit cannot resolve details in the shadowed interior of pits. Reconstructions of pit geometries from these views are profoundly limited in coverage, resolution, and accuracy. For this reason, surface exploration is needed as a means for cave access discovery and for study of pit morphology.

2022 IEEE Aerospace Conference (AERO) | 978-1-6654-3760-8/22/$31.00 ©2022 IEEE | DOI: 10.1109/AERO53065.2022.9843392

Small, fast, autonomous rovers hold the transformational promise of modeling planetary pits. Where prior planetary rovers have been slow, large, and expensive, pit exploration rovers could be fast, small, and economical. A fast microrover could circumnavigate, image and model an immense pit or a string of smaller pits in a single short mission. The microrover pit modeling concept leverages modern computing for the autonomy, imaging and 3D modeling to generate and return detailed triangle mesh models of pit interiors [6].

The exploration paradigm supported here is responsive to the small, short, economical, single-daylight-period lunar missions of our time such as PRISM (Payloads and Research Investigations on the Surface of the Moon) and CLPS (Commercial Lunar Payload Services). These currently accommodate rovers massing less than 20 kilograms and data rates on the order of 100 kbps. These preclude scenarios that downlink massive data that would be possible with greater duration and higher bandwidth.

This work develops specific capabilities for autonomous microrover pit exploration and evaluates the performance of these capabilities in end-to-end mission experiments performed at the West Desert Sinkhole. This is a terrestrial pit with dimensions and morphology comparable to known pits on the Moon and Mars, but the technologies demonstrated in this work are applicable to pits of any dimension.

First, we develop a technique for creating a sub-meter resolution elevation map from pit flyover imagery captured by a lander during its final descent to the surface. Then, we describe an exploration autonomy algorithm that processes the flyover elevation map to select and sequence safe vantages around a pit rim. Next, we introduce "brinkmanship", a safeguarding subsystem that enables a microrover to safely navigate to the edge of a pit to image its interior. Finally, we demonstrate in situ 3D modeling of a planetary pit using rover-acquired imagery.

These capabilities are evaluated in end-to-end mission simulations performed around the West Desert Sinkhole using the prototype microrover shown in Figure 1. In multiple evaluations, vantage point selection and brinkmanship are demonstrated. The resulting rover-generated terrain models are compared to a ground truth LIDAR survey performed in [7] and shown to exhibit excellent coverage and geometric accuracy.

2. Experimental Context

The experiment explores and models a terrestrial sinkhole as an analog to exploration of a planetary pit. It emulates key elements of a class of economic, near-term, from-rim microrover missions that can autonomously explore and model planetary pits for cave discovery and science on other worlds.

Landers for pit exploration differ from their predecessors by exhibiting the maneuverability and precision for delivery adjacent to a discrete feature, not just to a region. Accurate delivery precludes a long trek just to reach a target pit. This maneuverability and precision enables unprecedented functionality in a lander as an imaging instrument used during pit flyover and landing descent. This is the basis for **Pre-Modeling From Flyover Imagery** which is emulated in this experiment by 3D mapping of the surrounding terrain, pit apron, and upper pit walls from drone flyover imagery.

The robot utilizes this pre-model to (1) identify favorable vantages on the rim from which views of the pit are likely to provide the image coverage and perspectives for modeling, and (2) plan admissible routes for reaching those vantages. The resulting vantages are typically distributed around a pit. Typical route strategies are a combination of circumnavigation to reach azimuths of vantages, then scalloping or radial spurs to occupy the vantages. Sensing, safeguarding and navigational behaviors then autonomously guide the rover to reach those vantages while mapping terrain. Intended imagery is acquired at each vantage. These capabilities comprise the robot's **Pit Exploration Autonomy.**

Unique to from-rim rover exploration is the need to approach precipices on steep terrain while trading off ambition of approaching the brink for best viewing. The rover must balance this against risks such as inability to egress after imaging or a mission-ending fall into the pit. This 'brinkmanship' is distinct from general-purpose rover safeguarding. Traditional missions are risk-averse and prioritize least-challenging terrain with quasi-planar slope. The body of technology that governs the essential risk for pit exploration is here connoted as **Safeguarding at the Rim.**

The cave discovery and science observables from this class of mission derive primarily from illumination-corrected high-fidelity 3D pit models. These exhibit a leap of resolution, accuracy and coverage unachievable by any other means. The models are processed in-situ from 2,520 rover-acquired very high resolution images comprising 9.2 gigabytes of data. Pre-processing to eliminate overlapping imagery and moderate illumination reduces the size of the data by two orders of magnitude before submitting it to the modeling. Algorithm optimizations and accelerations compute the models in-situ in a few hours. The resulting models are a trivial amount of data relative to the acquired imagery, and they are readily conveyed to Earth. We connote the body of technology that achieves this as **In Situ Pit Modeling**.

The technical bases, implementations and evaluations of **Pre-Modeling from Flyover Imagery**, **Pit Exploration Autonomy**, **Safeguarding at Pit Rims**, and **In Situ Pit Modeling** are more fully developed and evaluated in the following sections.

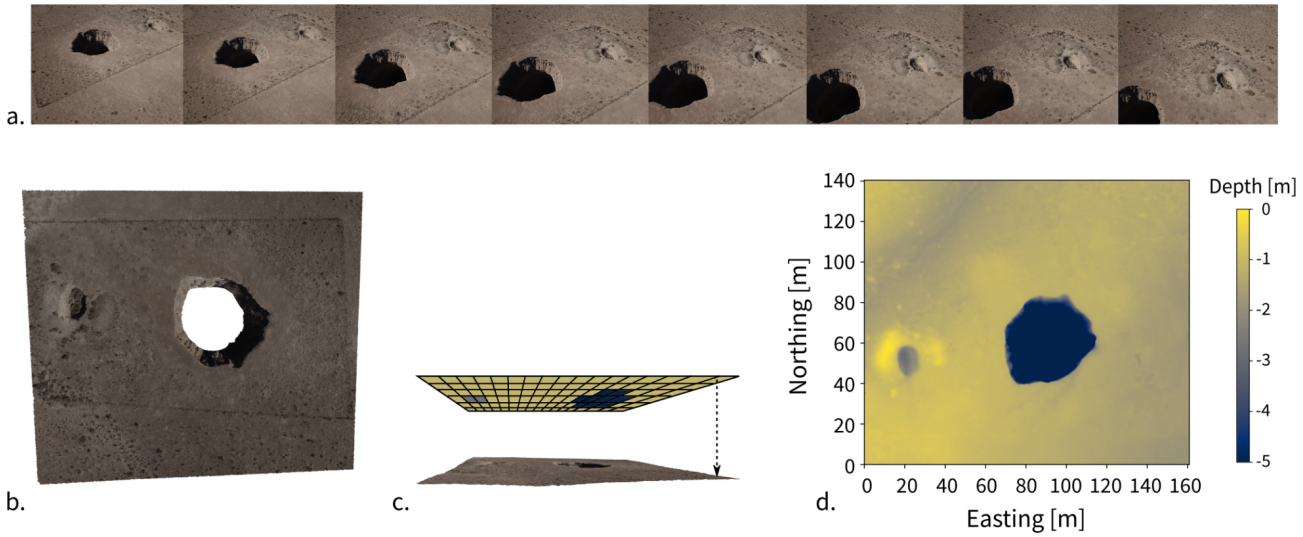

Figure 2. During descent the lander is directed to fly over the pit to acquire a series of down-looking birds-eye images (a). The imagery is processed to create a model of the pit (b). Note that the core of the pit is not modeled since the black imagery of the unlit depth provides no information to the vision algorithms. The resulting mesh is rasterized (c) into a 10 cm/pixel digital elevation map (d). The rover uses the resulting map and model to plan its exploration routes and select the vantages on the rim from which it will observe the pit.

3. Pre-Modeling From Flyover Imagery

Orbital elevation maps are too coarse to be relevant for microrover exploration planning. Lunar digital elevation maps (DEMs) derived from the the Lunar Reconnaissance Orbiter's (LRO) Narrow Angle Cameras (NAC) typically exhibit 3-5 m/pixel spatial resolution [8], but long-range exploration planning for meter-scale rovers requires resolution an order of magnitude greater than is available from orbit.

Fortunately, since a pit mission's lander is necessarily precise and maneuverable, it can capture high-resolution bird's-eye imagery of the pit, the landing site, and the surrounding terrain during descent. These images are passed to the rover and processed into a DEM of the site and pit rim that exhibits accuracy, resolution and utility for mission planning which far exceed the quality of anything possible from orbital data.

To simulate lander flyover imaging of a planetary pit, a drone acquired 48 down-looking 12 megapixel images while transiting across a pit at an altitude 1.5 times the pit's diameter descending to a landing site approximately one diameter from the center of the pit. A 3D model of visible portions of the pit and its surrounding terrain were generated by applying the modeling methodology of this research to the flyover imagery. When implemented on an actual mission, the pre-model's scale would be even more precisely scaled by associating pose of the source images to the lander's celestial/IMU referenced state.

The rover computed a 140m by 160m triangulated mesh model of the pit, landing site, and a wide swath of terrain in 53 minutes (Fig. 2b). The pit's interior is void since flyover imagery like satellite imagery cannot peer into darkness. The mesh was converted to a cartesian elevation map having a 10cm grid suitable for pit exploration planning and rover navigation. The conversion floated a 10 cm grid above the triangulated mesh, then assigned an elevation to each grid point as the elevation encountered in the triangulated mesh by ray casting downward from above. The resulting DEM is then normalized such that the height of the terrain at the landing site is zero, and all other heights are offsets from the elevation of the landing site. In practice during a mission the datum for the flyover-generated pre-model might be fit to the Moon's regional elevation DEM.

The resulting DEM that encodes pit terrain at 10 cm/pixel resolution is more than an order of magnitude greater resolution than is available from orbit.

4. Pit Exploration Autonomy

Autonomy is essential to microrovers that cannot carry direct-to-Earth radios or endure nights with isotope heating and hence must accomplish missions in a single illumination period. Autonomy is the key to continuous motion that is the

key to speed-made-good and to multi-kilometer exploration during a short mission window. Because microrovers cannot carry, aim, or power direct-to-earth radio transceivers, they are forced to relay their communications to Earth through their landers. Without autonomy, microrovers cannot explore beyond the limited range of lander communications.

Pit exploration autonomy is the on-rover capability to select safe vantage points on the rim of a planetary pit. It analyzes the high-resolution flyover DEM to select reachable poses with overlapping view coverage of the pit interior and sets high-level goals for the rover's underlying navigational autonomy to achieve.

Exploration autonomy utilizes the previously-described high-precision 10 centimeter gridded flyover DEM (Fig. 2d). The DEM is differentiated to derive a slope map. This is thresholded at 20 degrees to reflect the rover's tested mobility capability. A flood-fill of navigable terrain beginning from the landing site is used to convert the slope map into a binary map of reachable, safe terrain which appears as (Fig. 3).

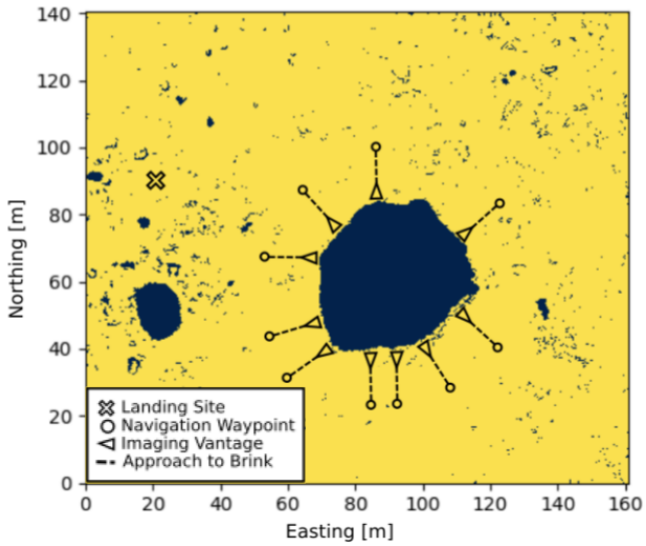

Figure 3. Exploration autonomy operates on flyover pre-modeling to identify imaging vantages around a pit that appear advantageous for imaging and reachable by the robot. Planning formulates radial vectors to approach and egress from the vantages along paths of steepest descent. The vectors embark from safe, navigable distance and approach radially to (a) preclude turning or side-sloping on steep ground, and (b) facilitate sensing and safeguarding behaviors when approaching the brink.

Geometric operators of dilation and erosion are applied to the binary map to identify a pseudo-center of the pit.

Vantages around the pit rim are selected by randomly sampling rover poses in a disk surrounding the pit. The disk is centered on the pit, and its radius is set to encompass the full extent of the pit as well as an additional 10m wide band of terrain. Up to 100 vantages are selected by uniform sampling from the reachable terrain inside the disk. Candidate vantages are rejected if the meter wide strip of terrain extending 5 meters radially from the rim and through the sampled point intersects any non-navigable terrain. This removes candidate vantages that cannot be reached by an approach that is radial to the rim.

Figure 4. In one field experiment, the rover traversed to nine reachable, safe vantages selected by pit exploration autonomy for their overlapping views of the pit interior and for their reachability by radial approach. .

The final set of vantages are selected from the pool of reachable candidates. From each candidate vantage, the angular extent of the visible pit rim is estimated. Ten thousand random subsets of vantages are generated, and any subset of vantages that does not achieve complete overlapping coverage of the pit rim is discarded. From the remaining subsets, the final set of vantages is selected as the set that is both radially reachable and achieves greatest total coverage of the pit rim.

Circumferential path segments connect the outer embarkation points of the radial approach vectors that are 5 meters away from the rim. These segments are planned over relatively easy terrain that is not challenged by the slopes or softness characteristic of the brink of a rim. These segments constitute a pseudo-polygon whose vertices are the outer embarkation points that lead to the vantage points. The global exploration plan occupies a polygon vertex, travels radially to its vantage point, images the pit, egresses by reversing rather than turning (to preclude risks of maneuvering on steepness), re-occupies the vertex of the polygon from where it embarked for that vantage, proceeds to another vertex, and repeats this until all vantages are

visited. This high level pit exploration plan is conveyed to the rover and executed by its navigational autonomy. The resulting exploration trajectory from exploration of the West Desert Sinkhole is shown in Figure 4 as it visits nine vantages.

5. Safeguarding at Pit Rims

Imaging deep inside a pit requires a rover to advance to the rim and pause at the brink. This balances the ambition for best imagery and modeling against conservatism for robot safety and egress. The algorithm that controls this behavior, dubbed "Brinkmanship", processes stereo image pairs and rover inertial state orientation to model the brink's terrain and evaluate the rover's mobility risk for progress versus egress.

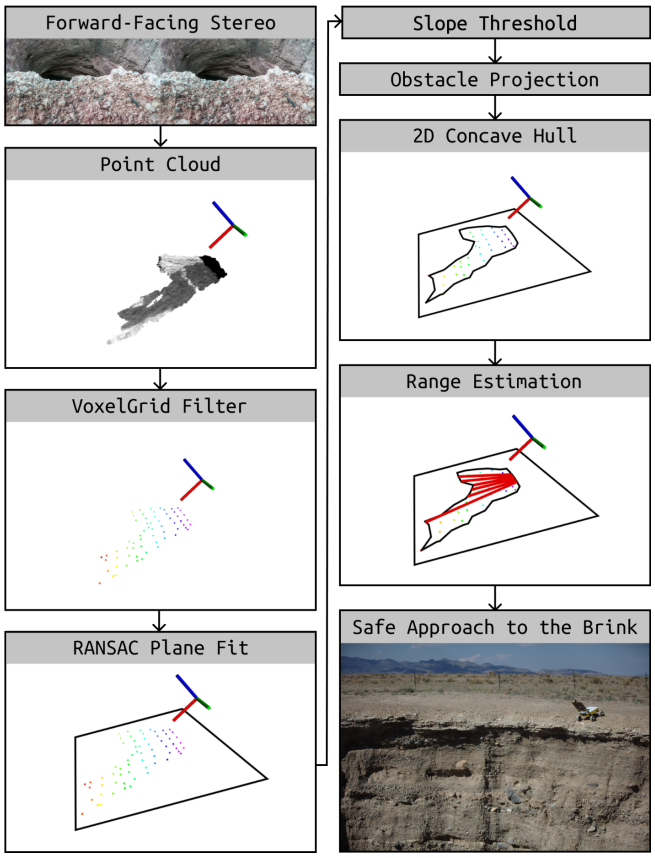

Figure 5. The brinkmanship algorithm analyzes terrain near a precipice relative to a rover's mobility capability. It continuously decides whether to proceed for the ambition of better pit viewing or stop due to concern for excessive risk.

Each iteration of the brinkmanship algorithm (Figure 5.) begins with a sparse point cloud derived from the rover's forward-facing stereo navigation camera. First, the point cloud is downsampled using a 5 centimeter voxel-grid filter to reduce the computational cost of future processing steps. Then, a range filter is used to remove points farther than 1m ahead of the rover. These long-range points appear when the rover approaches the brink and begins to sense the surface of the far pit wall. These points can be safely removed because they are not relevant to the task of estimating the range to the brink.

Next, a plane is fit to the filtered point cloud using Random Sample Consensus (RANSAC) [9]. If necessary, the plane normal is flipped to align with the +Z axis. The fitted plane's roll and pitch angles are calculated and compared to the IMU gravity vector. If the angle between the terrain and the gravity vector exceeds the rover's mobility limit, the brinkmanship algorithm reports a range of zero, indicating that it is not safe for the rover to proceed. In practice, these terrain slope limits are rarely triggered, as the rover's navigation autonomy effectively avoids steep terrain. At the brink, these limits are almost never invoked because more conservative safeguards cause the rover to stop long before it plunges over the edge.

Brinkmanship leaves detection and handling of positive obstacles to the rover's navigation autonomy, but it must account for their presence when searching for negative obstacles such as pit rims. Points above the RANSAC plane are projected along a ray from the camera origin onto the plane itself. If instead these points are removed or projected orthogonally onto the plane, it creates a large region containing no points. This is to be avoided because the edge of a large empty region is difficult to distinguish from a pit rim.

All remaining points are now on or below the fitted plane. They are projected orthogonally onto the plane, creating a 2D region whose point density represents the presence or absence of terrain. The concave boundary of this region is computed as the alpha-shape ($\alpha = 0.1$) of the 2D point set [10]. The result is a boundary polygon whose interior represents navigable terrain.

It is necessary to compute the concave hull of the point set because the intersection of the camera's viewing frustum with the convex rim of a pit is concave. If the convex hull is used instead, the safe range will frequently overestimate the true range to the pit rim.

To estimate the safe driving range using the 2D boundary polygon, a fan of five rays is projected forward from the front center of the rover to the forward boundary of the polygon. The length of the two shortest rays are averaged and reported as the brinkmanship estimate.

The brinkmanship estimate is a measure of the known-safe terrain in front of the rover. As the rover approaches a pit

rim, the brinkmanship range decreases gradually. As the range decreases below 1m, the rover slows its driving speed to two centimeters per second. Once the safe range decreases below 45 centimeters, the rover stops moving forward and begins imaging the pit interior.

The brinkmanship algorithm exposes two parameters that depend on rover geometry and mobility. Pitch and roll thresholds used to identify steep terrain were determined in mobility testing of the prototype pit rover. In soft terrain, the rover reliably ascends a 30 degree slope and traverses perpendicular to a 25 degree slope. Brinkmanship classifies terrain as unsafe if the terrain pitch or roll angles exceed 80% of these values.

The operating frequency of the algorithm is determined by the rover's maximum driving speed. The top speed of the prototype rover used in this work is 12cm/sec. In field tests, an iteration rate of 6Hz was sufficient to guard the rover from falls.

The alpha parameter used to construct the concave hull of the boundary polygon must be tuned, but the performance of the brinkmanship algorithm is insensitive to its exact value. In field testing, values ranging from 0.05 to 0.25 were found to adequately account for concave boundaries. For this work, the value of alpha was fixed at 0.1.

In six field trials, the rover drove approximately 1.2 kilometers and performed 90 approaches to the pit rim. In no instance did the rover fail to safeguard at the pit rim.

Twice, brinkmanship reported a pit rim while driving in flat terrain. These false positive detections were traced to transient self-shadowing by the rover's solar panel. As the front-facing stereo camera entered the rover's shadow, its autoexposure algorithm reacted slowly to the rapid decrease in illumination. The stereo camera temporarily produced over-exposed images and low-quality point clouds that triggered the brinkmanship safeguards. This problem was resolved by low-pass filtering the brinkmanship range estimate across time. A moving-average filter of the past three range estimates removed the transient detection without degrading the rover's ability to detect the pit rim.

6. In Situ Pit Modeling

In situ modeling generates accurate, detailed and realistic 3D pit geometry of a pit - including in shadowed and overlit regions - from two-dimensional images acquired from multiple viewpoints around the pit's rim.

The pit exploration rover developed in this initiative carries a 20 Megapixel camera with a narrow-angle lens atop its hinged solar panel (Fig. 1). At each imaging overlook, the rover tilts its camera from horizontal to 70° below the horizon. At each tilt angle, the camera pans from 45° left to 45° right, capturing overlapping panoramic imagery of the opposing pit wall.

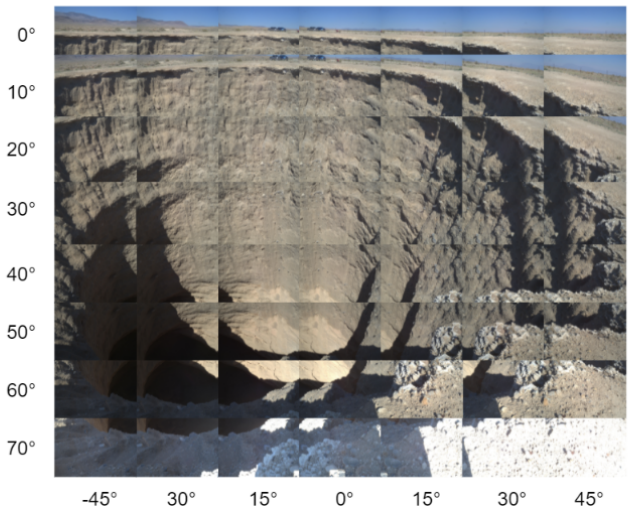

Figure 6. At each of 56 pan and tilt locations, a short, medium, and long exposure image was acquired. These were post-processed to moderate burnout and brighten dark regions resulting in 56 illumination-moderated images. The combined mosaic spans a 118° (horizontal) by 90° (vertical) field of view and contains 672 million floating-point pixels.

The pit imaging camera's horizontal field of view is 28°, and its vertical field of view is 20°. To capture the 50+% overlap required for photogrammetry [11], the camera panned in increments of 15° and tilted in increments of 10°. In total, the rover camera captured images from 56 distinct camera poses at each of the rim overlooks (Fig. 6).

At each camera pose, the rover captured an exposure bracket containing three images (Fig. 7). The exposure of the central image in the bracket was selected using the mean grey auto-exposure algorithm. The remaining two image exposures were set to +/- 1 EV above and below the center image exposure.

Bracketed exposures are critical for modeling deep pits with extremely wide dynamic range. Long exposures capture detail in the shadowed pit interior while short exposures retain detail along the brightly illuminated rim.

Immediately post-capture, the pit modeling process running onboard the rover uses Debevec's weighting scheme to merge each exposure bracket into a single High Dynamic Range (HDR) image (Fig 7.). This reduces the number of images that must be processed to synthesize a pit model by a factor of 3x, due to the 3-image exposure bracket. For

planetary operations, the bracket size may be increased to 5 or more, leading to even larger reductions between the size of the image set at initial capture and post-HDR processing.

Figure 7. Fusing images from multiple exposures (left) produces a single HDR image (right) that suppresses burnout, brightens darkness, and retains detail in both brightly and dimly lit regions of the pit.

Overlapping HDR image mosaics captured from 9 vantages around the pit rim were used to synthesize the pit model shown in Figure 8. The photogrammetric reconstruction ran for ~190 minutes on a mission-relevant Nvidia TX2i embedded computer to produce a 640 MB textured mesh model of the terrestrial pit covering 2850 square meters of the pit interior surface.

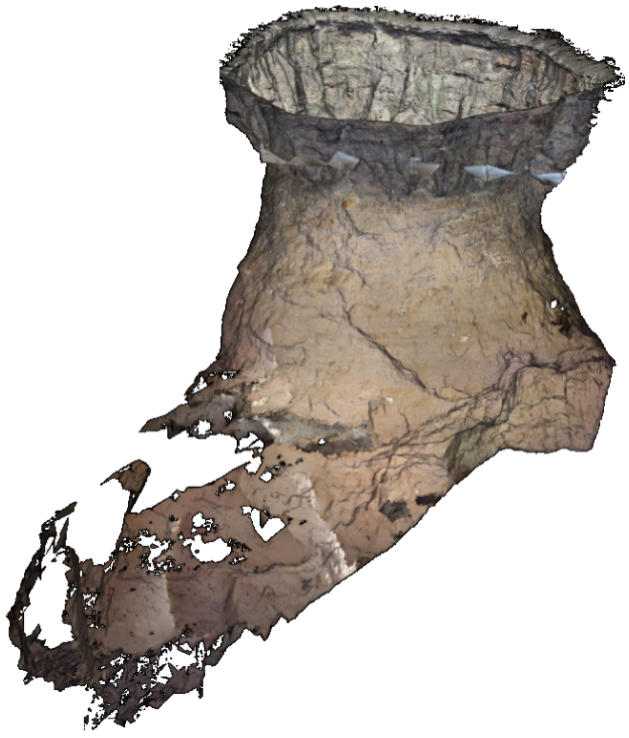

Figure 8. This textured triangle mesh is the result of pit modeling applied to the 2,520 images acquired by the pit rover around the rim of the West Desert Sinkhole.

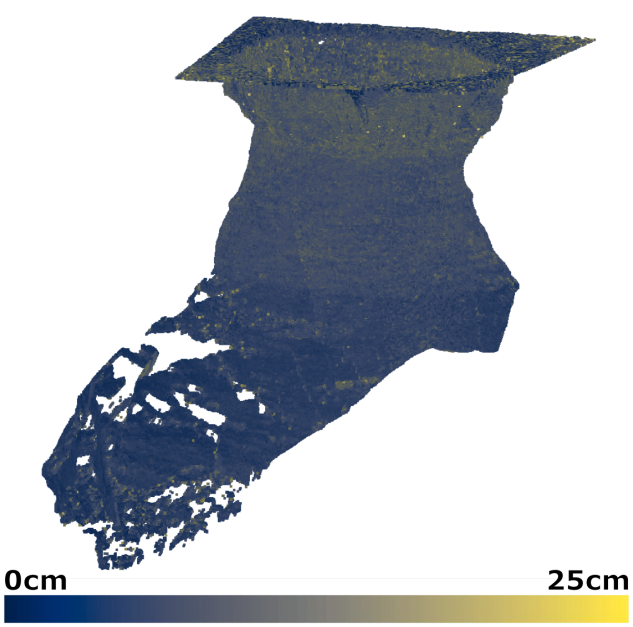

Figure 9. Point-to-plane deviation between the rover-acquired pit model and LIDAR ground truth is less than 10cm across more than 90% of the pit surface.

The geometric accuracy of the rover-generated pit model was evaluated by comparison to a comprehensive LIDAR survey of the West Desert Sinkhole [7]. The rover-generated pit model was registered against the LIDAR survey model using the Iterative Closest Point (ICP) algorithm with adjustments to scale disabled. The accuracy of the pit modeling method was quantified by computing the minimum distance from all vertices in the rover-generated model to their nearest neighbor in the registered LIDAR survey. Across more than 90% of the rover-generated pit surface, the minimum distance to the LIDAR point cloud was less than 10 centimeter (Fig. 9). For depths less than 19 meters, deviations were less than 6 centimeters (Fig. 10).

After compression, the pit model geometry and texture data are 640 MB. Using 100 kbps as a reasonable bandwidth estimate for a commercial lander payload, this requires only 12 hours to transmit to Earth. This is a drastic reduction over the 8 days that would otherwise be required to transmit the original 9.2 GB of pit imagery, and is much more reasonable for short duration, low-bandwidth pit exploration missions.

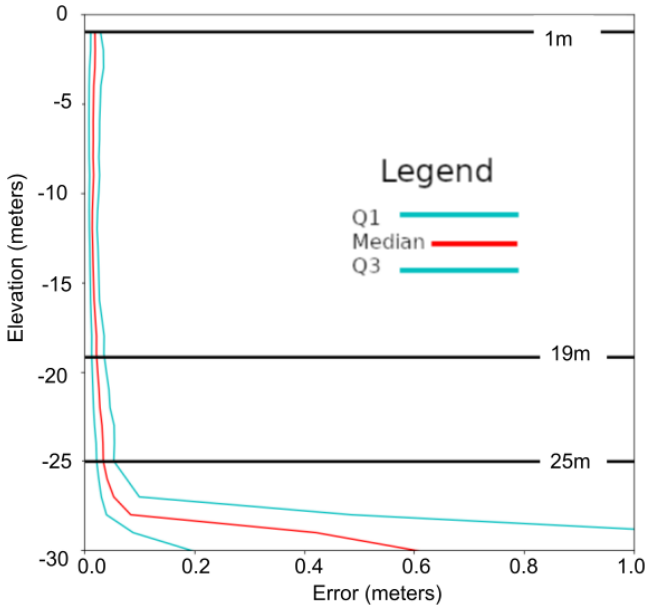

Figure 10. As depth increases, median reconstruction errors increase. Below 25m depth, view coverage becomes sparse, and point-to-mesh deviation from ground truth increases to a maximum of 60 centimeters.

7. Conclusion

This paper chronicled and evaluated the first autonomous microrover exploration of a pit and detailed its enabling technologies. The experiment autonomously explored and modeled Utah's West Desert Sinkhole as a terrestrial analog of a planetary pit.

The experiment demonstrated exploration autonomy, navigation, brinkmanship and image acquisition for pit modeling. The rover traversed 200 meters to and around the apron, with nine approaches to the brink of the pit, which took a total of 115 minutes. The rover acquired 2520 raw images at the nine stops overlooking the pit. These images were then processed onboard the rover for a span of 4 hours to generate a 3D model of the terrestrial pit. Pre-processing corrected illumination to both minimize burnout and view into under-illuminated depths of the pit. The resulting 3D model of the pit exhibited accuracy within 10 cm over 90% of the visible area. Notably, the model revealed a dark, overhung cavernous opening. The end-to-end compression from total raw imagery to triangulated pit model was 14:1.

The outcomes of this experiment conclude that:

1. Pre-modeling by pit flyover is viable and useful.

2. Processing of robotically-acquired imagery from a pit's rim generates exceptional, accurate, high-coverage, high-fidelity 3D photorealistic models.

3. High Dynamic Range pre-processing is effective at precluding burnout and viewing regions that would otherwise be obscured in darkness.

4. The requisite autonomy, imaging and pit modeling for this class of exploration are viable by microrover.

5. Microrover exploration is viable as a means of planetary pit modeling for purposes of cave discovery and science.

8. Future

Deeper post-processing of the data acquired in this experiment is a significant opportunity for future agenda—so much more is possible beyond what was achieved and reported in the short period between the field experiment and submission of this publication.

Results reported here independently generate (1) the flyover pit model, (2) reconstruction of the rover's subtended route, and (3) pit model as generated from rover-acquired imagery. These independent and sequential computations produce a superb model, but much more might be possible from cumulative and incrementally progressive co-processing of these three. At minimum, this offers the opportunity of maximizing the incorporation of all data as it is acquired to generate a continuum of successively more complete, detailed, accurate, and illumination-corrected models than is possible by any of the three in isolation. This could go so far as investigating a continuum of incrementally updating the pit model with imagery from each vantage point in a drive-image-model loop versus the all-drive, all-image, then-model sequence presented here. The ultimate interleaving might additionally include interactive exploration planning - including planning for different or additional vantages that might be compelled to remedy model weaknesses.

The model evaluated in this paper incorporates the imagery at every exposure that the robot acquired from all nine vantages at the rim. A study is merited to investigate the sensitivity of the model's coverage, accuracy and lighting information as a function of number and location of viewpoints from the rim, image resolution, and lighting correction methodology. This would rationalize sufficiency of achieving mission goals. For example, if a planetary mission were exploring a pit string, the necessary quality of

modeling could be assured relative to ambitions for visiting many pits.

There is a significant opportunity to further deepen and evaluate brinkmanship technology. The brinkmanship deployed for this experiment is predicated only on the brink's geometry. Planetary pit rims, however, are generally far steeper and much weaker since planetary pits are altered after formation by eons of impact. The current brinkmanship does not incorporate slip estimation or terra-mechanic considerations, and hence it is not comprehensive for the very steep, weak conditions that are likely to pertain near brinks of planetary rims.

The objective of this experiment is to ensure mission success of an early quality model before progressing successively to a penultimate model as a function of ongoing time, data and risk. A future experiment might progress from conservative standoff to aggressive approach to the rim by successively circling a pit. A first safe transit would generate a quality model, but would likely miss some near-field coverage and perspectives that are only determined from a more proximate approach. Successive transits might improve the modeling by tolerating greater risk for an increasingly bold approach to the precipice.

A future experiment might add an even longer image exposure or two to achieve even better imaging in darkness. The most intriguing aspect of the model from this experiment is arguably its modeling of the cavernous 'toe' of the boot that is analogous to discovering a cave portal in darkness. This capability might be vastly enhanced by better imaging in darkness.

The ultimate future is to seek caves and undertake science at planetary pits through autonomous microrover exploration.

Acknowledgements

The authors thank NIAC: NASA's Innovative Advanced Concepts program for support of this research (NIAC 80HQTR19C0034). The authors also thank Tim Angert and Daniel Scher for their assistance with development and evaluation of the rover that deployed the technologies exhibited in this work.

References

[1] J. W. Ashley *et al.*, "Lunar Pits: Sublunarean Voids and the Nature of Mare Emplacement," presented at the 42nd Lunar and Planetary Science Conference, 2011.

[2] F. Horz, "Lava tubes - Potential shelters for habitats," *Lunar Bases Space Act. 21st Century*, pp. 405–412, 1985.

[3] R. V. Wagner and M. S. Robinson, "Occurrence and Origin of Lunar Pits: Observations from a New Catalog," presented at the 52nd Lunar and Planetary Science Conference 2021.

[4] "Exploration of Planetary Skylights and Tunnels." Accessed: Oct. 14, 2021. [Online]. Available: https://www.nasa.gov/sites/default/files/atoms/files/whittaker_niac_exploration_planetaryskylights_tunnels_phase2.pdf

[5] G. E. Cushing, C. H. Okubo, and T. N. Titus, "Atypical pit craters on Mars: New insights from THEMIS, CTX, and HiRISE observations," *J. Geophys. Res. Planets*, vol. 120, no. 6, pp. 1023–1043, 2015, doi: 10.1002/2014JE004735.

[6] W. L. Whittaker, H. L. Jones, J. S. Ford, K. Sharif, and U. Y. Wong, "SKYLIGHT: A MISSION CONCEPT FOR IN-SITU INVESTIGATION OF THE MORPHOLOGY, GEOLOGY AND MINERALOGY OF LUNAR PITS," presented at the 52nd Lunar and Planetary Science Conference, 2021.

[7] J. . Ford, K. Sharif, H. L. Jones, and W. L. Whittaker, "PLANETARY PIT PHOTOGRAMMETRY: ACCURACY, COVERAGE, AND SCIENCE VALUE," presented at the 52nd Lunar and Planetary Science Conference, 2021.

[8] LROC, "Lunar Orbital Data Explorer - NAC_DTM_LACUSMORT01 Data Product." https://ode.rsl.wustl.edu/moon/productPageAtlas.aspx?product_id=NAC_DTM_LACUSMORT01&product_idGeo=24284964 (accessed Oct. 15, 2021).

[9] M. A. Fischler and R. C. Bolles, "Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography," *Commun. ACM*, vol. 24, no. 6, pp. 381–395, Jun. 1981, doi: 10.1145/358669.358692.

[10] H. Edelsbrunner, D. Kirkpatrick, and R. Seidel, "On the shape of a set of points in the plane," *IEEE Trans. Inf. Theory*, vol. 29, no. 4, pp. 551–559, Jul. 1983, doi: 10.1109/TIT.1983.1056714.

[11] P. Arias, J. Herráez, H. Lorenzo, and C. Ordóñez, "Control of structural problems in cultural heritage monuments using close-range photogrammetry and computer methods," *Comput. Struct.*, vol. 83, no. 21, pp. 1754–1766, Aug. 2005.

Biography

Jordan Ford *is a PhD candidate in the Carnegie Mellon Field Robotics Center studying exploration autonomy and photogrammetric mapping for robotic exploration of lunar pits. Previously, he has developed motion planning algorithms for autonomous vehicles and contamination mapping algorithms for nuclear cleanup. He earned a B.S. and M.S. in Computer Engineering from the Georgia Institute of Technology in 2016.*

Khaled Sharif *works at the NASA Ames Research Center with the Intelligent Robotics Group, and has a background in computer and electrical engineering. He is a robotics engineer working on Astrobee, a free-flying robot on-board the International Space Station. Khaled earned his M.S. in Computer Science from the Georgia Institute of Technology.*

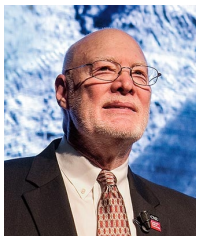

William L. "Red" Whittaker *is the SCS Founders Research Professor of Robotics at Carnegie Mellon and Director of its Field Robotics Center. His B.S. is from Princeton. His M.S. and PhD are from CMU. Red leads the pit exploration program described in this paper.*

Warren "Chuck" Whittaker *is a Senior Field Robotics Specialist at Carnegie Mellon University. He led the field experiment and lidar modeling reported here. He has decades of experience designing, fabricating, integrating, and testing robots for planetary exploration and other field applications. He also served as engineering manager for RedZone and chief operating officer for Workhorse Technologies, LLC., working with both companies to develop robots for exploration in hazardous environments. He earned his B.S. from University of Pittsburgh at Johnstown.*

Heather Jones *is a Senior Project Scientist at Carnegie Mellon University. She has developed and tested robots for planetary and nuclear applications. Prior to her graduate studies, She worked at the NASA Johnson Space Center doing analysis of operations for the Canadarm2 robot on the International Space Station. She earned her B.A. and B.S. from Swarthmore College. Her M.S. and PhD were earned from Carnegie Mellon with perception and planning research for robotic modeling of pits and other planetary terrain features.*

Uland Wong *is a Senior Computer Scientist at the NASA Ames Research Center with the Intelligent Robotics Group. He earned a B.S., two M.S. degrees, and a PhD from CMU. He has spent more than 15 years giving robots the ability to see at the frontiers of exploration - from dark caves and planetary poles to icy surfaces.*