

TECHNOLOGIES ENABLING THE EXPLORATION OF LUNAR PITS

Virtual Conference 19–23 October 2020

Jordan Ford¹, Heather Jones¹, Khaled Sharif², Red Whittaker³

¹Robotics Institute, Carnegie Mellon University, Pittsburgh, PA 15213, USA, [jsford,hjones]@andrew.cmu.edu

²Ames Research Center, Mountain View, CA 94035, USA, khaled.sharif@nasa.gov

³Robotics Institute, Carnegie Mellon University, Pittsburgh, PA 15213, USA, red@cmu.edu

ABSTRACT

Lunar pits are windows to great unknowns beneath the surface of the Moon. The steep pit walls expose a vertical cross-section of stratified lunar crust, a natural historical record of the Moon's formation. The shadowed pit floors conceal possible entrances to caves or lava tube networks where astronauts might shelter from the hazards of the lunar surface. While orbital imagery has led to the discovery of more than 200 lunar pits, surface exploration is needed to capture the close-range, long exposures required to address questions of lunar formation and habitation. We envision an autonomous microrover for rapid surface exploration and *in situ* modeling of a lunar pit. We survey the specific autonomy, mapping, and roving technologies required, and we demonstrate the accuracy and coverage achievable by pit modeling in an evaluation performed at a terrestrial pit.




Figure 1: Artist's depiction of a pit exploration rover imaging the opposing wall from a vantage on the rim.

1 INTRODUCTION

Lunar pits are compelling exploration destinations that could change the future of human presence on the Moon and reshape our understanding of the Moon's past.

Of the hundreds of known lunar pits, a rare few might conceal entrances to caves or lava tube networks. Many melt pits are likelier to access overhung caverns of smaller scale. If accessible, lunar caves could shelter humans from the thermal fluctuations, cosmic radiation, and micrometeorite impingement they experience on the lunar surface [1]. The permanently dark interior of a lunar cave might also harbor frozen volatiles that could provide fuel and water for future utilization [2].

Lunar pits are scientific destinations in their own right. The steep, exposed pit walls reveal stratified timelines of lunar geological events, and the pit interior may preserve the only remaining biosignatures of ancient lunar life [3].

While orbital imagery has been used to discover more than 200 lunar pits, orbiters cannot capture the close-range, long exposure images required to address questions of cave access and lunar science. For this reason, a variety of surface exploration approaches have been advanced. Some propose deploying a tethered rover over the rim and into the pit [4]. Direct landers, drones, walking robots, and bouncing balloon-bots have also been suggested [5]. Untethered approaches that enter pits are challenged by non-line-of-sight communication, lack of illumination for perception, and fractional coverage by sunlight for powering on pit floors.

We envision affordable, near-term exploration by microrovers that view and model pits from their rims to discern cavernous openings and scientifically investigate morphology and volcanology. The scenario is to navigate the apron, discover and occupy overlooks, capture telescopic images of the far wall and visible floor, repeat this while circumnavigating the apron, then processing the imagery to compute a high-fidelity 3D model from which cave discoveries and the science derive. We derive requirements for the specific roving,

autonomy, and modeling technologies required and demonstrate the coverage and accuracy achievable by *in situ* modeling of planetary pits.

2 MISSION CONCEPT

Pit exploration mission ConOps proceed in three phases (Fig. 2).

If a landing is close enough to a pit then earliest actions and downlinks can be commanded from and relayed to Earth. Otherwise, the rover autonomously proceeds beyond comm. range.

In phase one, the rover traverses to the edge of the pit and captures panoramic pit images. These images are triaged, compressed, and a few are streamed through the lander to Earth to ensure an earliest result before proceeding with the comprehensive imaging, mapping, and modeling. The rover then proceeds to capture imagery from as many vantage points as possible while remaining in communication with the lander.

In phase two, the rover ventures autonomously beyond communication range. It identifies and occupies vantage points overlooking the pit. At each vantage, it acquires imagery and distills many images into few by applying HDR fusion and panoramic mosaicing. The data reduction is greater than hundredfold. The rover returns within communication range of the lander to upload images for modeling and continuous data downlink to Earth.

In the third phase, the rover circumnavigates the pit, capturing panoramic imagery at more and more vantages. This proceeds in sectors during several treks to support incremental refinement of the pit model and continuous downlink of imagery and improved model to Earth. Caverns, if any, are revealed in the models. Science investigations relating to morphology, vulcanology, and geology derive from the models and imagery.




Figure 2: Pit exploration proceeds in three phases. After each phase the rover returns to lander communication to upload imagery for modeling.

3 HIGH-LEVEL REQUIREMENTS

Frequent, affordable, small lunar missions of our time call for minimalist rovers that lack direct-to-Earth communication, must achieve goals in a single lunar daylight period, and must exhibit a leap of autonomy to accomplish significant exploration.

Small Weight and Power—

Rovers relevant for these missions must be <20kg, stow in <0.2m³ volume, cost <\$10M, and be energy-rich, computationally powerful, and solar powered [6].

In Situ Pit Modeling—

The pit model must be computed on the Moon since small landers provide only limited bandwidth that could not possibly convey the vast, high-resolution imagery to Earth where the modeling on high-powered terrestrial computing would be straightforward. The innovative incremental *in situ* photogrammetry developed in this program enables the lander to combine all image and telemetry data into a thin data stream with maximal explorational and scientific utility.

Explorational Autonomy—

Without an isotopic heat source or a massive battery, the rover and lander do not survive lunar night. Therefore, the multi-kilometer mission must

complete in a single illumination period. This is vastly distinct from nuclear powered Mars rovers that aspire to decadal duration. This requires two orders-of-magnitude average speed over prior planetary rovers. This is only achievable by autonomy capable of guiding the rover in near-continuous motion.




Figure 3: A terrestrial prototype pit rover carries a high-resolution camera on a pan axis atop a tilting solar panel for acquiring imagery within a pit.

Telescopic Pit Camera—

A telescopic camera is essential for capturing the scientifically relevant high-resolution detail on distant walls and floors. The individual images have a small field of view. The pan-tilt is essential for capturing the myriad overlapping small-footprint images required to cover the walls and floors.

Powerful Computing—

Powerful computing beyond that of any prior rover is absolutely essential for achieving the required autonomy, high-speed navigation, and image processing to complete the exploration in a single illumination period.

Steep-Slope Mobility—

Many favorable vantages are accessible for some pit geometries, but negotiating steep slope matters for occupying general pit overlooks. Ascent of steep, soft slope determines the quality and quantity of overlooking vantages that are accessible. Weathered pit aprons exhibit funnels that challenge this mobility. Significant coverage is achievable without undue risk for many pits. The lowest visible ray governs the depth that can be viewed on the far wall and the area that can be viewed on the far floor. It is determined by the achievable pitch and the height to which the camera can be elevated above the rim.

4 PIT MODELING

Modeling is the means by which thousands of images captured by a rover are aggregated into a highly-detailed, high-fidelity textured triangular mesh model of a pit. The following example models a terrestrial pit, the West Desert Sinkhole (39.2047° N, 113.2705° W), to illustrate the method in action and exhibit its coverage and accuracy.

4.1 Image Capture

At each overlook, the 20 Megapixel rover camera panned $\pm 90^\circ$ and tilted 80° below the horizon in increments of 10° . The maximum height of the camera above the rim was 20cm. At each camera pose, five images were captured with exposures ranging from 625 μ s to 500 ms.

4.2 High Dynamic Range

Immediately post-capture, Debevec's weighting scheme was used to combine each sequence of five images into a single high dynamic range image [7]. This preserves detail in both the dark pit interior and the brightly lit pit rim (Figure 3), and it reduces the per-vantage image count from 720 to 144.



Figure 3: Fusing images from multiple exposures produces a single HDR image with detail in brightly and dimly lit regions of the pit.

4.3 Panorama Construction

The 144 HDR images from each vantage overlap significantly, so much of the image data is redundant. Stitching the 144 images into a panorama (Figure 4) removes much of the redundant image data and significantly speeds up future transmission and processing [8].



Figure 4: Stitching the HDR images into panoramas prior to pit modeling reduces redundant image data and simplifies photogrammetry.

4.4 Photogrammetry




Figure 5: Photogrammetry estimates camera poses and constructs a pit model from panoramic imagery.

Panoramas from 26 vantages around the rim of the analog pit were processed using COLMAP [9,10], an open-source photogrammetry package, to produce the textured triangle mesh shown in Figure 6.




Figure 6: This textured triangle mesh is the result of photogrammetric pit modeling of the West Desert Sinkhole.

4.5 Ground Truth

In order to evaluate the geometric accuracy of the photogrammetric pit model, a ground truth LIDAR point cloud was created. 44 overlapping LIDAR scans of the pit and surrounding terrain were captured and co-registered. This produced a single point cloud with 100 million points and 5cm precision.

4.5 Geometric Evaluation

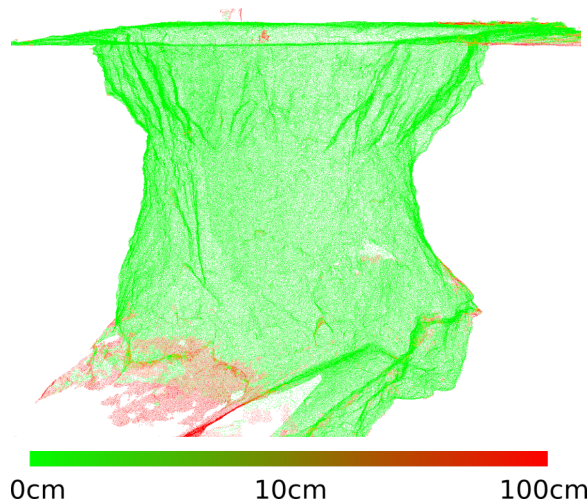


Figure 7: Point-to-plane deviation between the photogrammetric model and the LIDAR ground truth is less than 10cm across 90% of the pit surface.

The photogrammetric model and the LIDAR model were registered in the same coordinate frame, and the point-to-plane distance was calculated for each LIDAR point. Greater than 90% of the photogrammetric model surface was found to be closer than 10cm to the LIDAR ground truth (Figure 7). As expected, the largest deviations occur in the deepest region of the pit where visibility is poor for both cameras and LIDAR.

5 PIT EXPLORATION AUTONOMY

Autonomy is critical since conventional teleoperation is not possible without direct-to-Earth radios. Stopping and waiting for human guidance contributes to why prior rovers have made slow progress. Continuous exploration planning and autonomy enable a hundredfold increase of average speed over historical mission norms.

5.1 Viewpoint Selection

The highest level of the autonomy generates behaviors, plans, and decisions for acquiring the required images. The autonomy directs the rover to occupy strategic overlooks, acquire thousands of relevant images, and periodically return to the lander to transmit. Overlooks are identified by balancing the value of imagery against the risks encountered while

descending to the rim. At each overlook images are acquired at many pan/tilt combinations to view the far wall and floor.




Figure 8: The rover models local terrain as a triangle mesh for planning and safeguarding purposes.

Periodically, the rover shuttles back to the lander to transmit the latest model data. The rover then returns to the pit to continue exploration. In this way, serial models are downlinked early and often to Earth, and the modeling progresses continuously. The longer that the rover and lander remain operational, the greater the fidelity and coverage of the model.

5.2 Driving Autonomy

High-cadence stereo perception and fast replanning achieve the rigor and cadence to guide the rover with high quality navigation in continuous motion at two centimeters per second. Given a landing reasonably close to a pit, this is double the speed required to achieve exploration of any largest pits, a length of pit chain, or a cluster of melt pits. We extend the RASM autonomous roving framework to perform safe exploration near the edge of a lunar pit [11]. The navigation and mapping must exhibit a leap of quality, speed and traverse accuracy beyond prior rovers.

6 SUMMARY

This work describes a pit exploration mission scenario and derives high-level requirements for autonomous roving, *in situ* photogrammetry, and micro-roving. We provide a pit modeling example demonstrating the coverage and accuracy achievable by *in situ* HDR photogrammetry of a terrestrial analog pit, and we survey the autonomy and roving technologies required to accomplish microrover pit exploration.

Acknowledgement

The developments presented here are supported by NASA's Institute for Advanced Concepts under grant #80HQTR19C0034.

References

- [1] Horz, F., 1985, Lunar Bases and Space Activities of the 21st Century, Lunar Planet. Sci. Inst. p. 405.
- [2] Buczkowski, D. L. and Wyrick, D. Y. (2020). Tectonic Caves on Small Bodies: Potential In Situ Resource Reservoirs. *3rd International Planetary Caves*.
- [3] Boston PJ, Spilde MN, Northup DE, et al. Cave biosignature suites: microbes, minerals, and Mars. *Astrobiology*. 2001;1(1):25-55. doi:10.1089/153110701750137413
- [4] Kerber, L. et al. (2020). Moon Diver: Journey into the Ancient Lavas of the Moon. *3rd International Planetary Caves*.
- [5] Lee, P. et al. (2020) Ice-Rich Caves on the Moon and Mars: Prospects and Pragmatic Recommendations for Exploration. *3rd International Planetary Caves*.
- [6] Astrobotic Technology Inc. (2020) Payload User's Guide. p 5.
- [7] Paul E Debevec and Jitendra Malik. Recovering high dynamic range radiance maps from photographs. In *ACM SIGGRAPH 2008 classes*, page 31. ACM, 2008.
- [8] Dersch, H. (2007, June). Panorama tools: open source software for immersive imaging. *International VR Photography Conference*, Berkeley.
- [9] Schönberger, J.L., and Frahm, J. "Structure-from-Motion Revisited." . In *Conference on Computer Vision and Pattern Recognition (CVPR)*.2016.
- [10] Schönberger, J.L., Zheng E., Pollefeys, M., and Frahm, J. "Pixelwise View Selection for Unstructured Multi-View Stereo." . In *European Conference on Computer Vision (ECCV)*.2016.
- [11] Wettergreen, D. and Wagner, M.D. (2012) Developing a Framework for Reliable Autonomous Surface Mobility. *Proceedings of International Symposium on Artificial Intelligence, Robotics and Automation in Space*.