



Carnegie Mellon University

Exploration of Planetary Skylights and Tunnels

NASA Innovative Advanced Concepts (NIAC) Phase II

FOR

OFFICE OF THE CHIEF TECHNOLOGIST
NATIONAL AERONAUTICS AND SPACE ADMINISTRATION
GRANT NUMBER: NNX12AQ55G

AWARD DATE: SEPTEMBER 10, 2012
END DATE: SEPTEMBER 9, 2014

PI: Red Whittaker

Email: red@cmu.edu

Home Page: astrobotic.com

Co-I: Uland Wong

Email: uland.wong@gmail.com

Program Manager: Steven Huber

Email: steven.huber@astrobotic.com

Summary Report

December 8, 2014



Report prepared by:

Uland Wong²

Heather Jones²

Steven Huber¹

Christopher Cunningham²

Warren C. Whittaker²

Steve McGuire¹

Xuesu Xiao²

Rick Shanor²

Ander Solorzano^{1,2}

Tom Carlone¹

Wennie Tabib^{1,2}

Christopher Greve¹

Lauren Schneider¹

Nathan Otten²

William L. "Red" Whittaker^{1,2}

1. Astrobotic Technology Inc.

2515 Liberty Ave.

Pittsburgh, PA 15222

Phone: 412-682-3282

2. Carnegie Mellon University

5000 Forbes Ave

Pittsburgh, PA 15237

1 EXECUTIVE SUMMARY

While planetary pits and caves have been fiction for a century, they have been seen from orbit only in the last few years. These discoveries exceed the fantasies in diversity, scale, and abundance. For pits and caves, this is the age of discovery, ranging from a few pits on the Moon and Mars




Figure 1: Pits have been identified throughout the solar system, including on the Moon (1), Mars (2), Phobos (3), Eros (4), and Gaspra (5).

in 2009 to hundreds within the time of this research, with many more to come. Pits with subsurface voids have been confirmed on the Moon and Mars and indicated on Venus, Phobos, Eros, Gaspra, Ida, Enceladus, and Europa. Compelling next steps are surface and subsurface exploration.

Pits and caves are opportunistic study targets for unique origins, geology, and climate that will broadly impact planetary science. Holes on Mars are of particular interest because their interior caves are relatively protected from the harsh surface, making them good candidates to contain Martian life. Pits are prime targets for possible future spacecraft, robots, and even human interplanetary explorers. Caves and caverns could be ready-made shelters for future Moon and Mars explorers and colonists.

Discoveries to date look down from on high with satellites but cannot reveal the wonders of caves. They cannot enter, touch, or view pits up close. Genuine exploration is only achievable through surface missions. Robotic missions can assess suitability for safe entry and habitation, plus inform techniques for developing subsurface infrastructure.



Figure 2: Pits are prime targets for possible future spacecraft, robots, and even human interplanetary explorers

Missions into planetary voids redefine the future of exploration, science, and habitation beyond Earth. We can reach this future only by targeting specific technological advancement now. Prior missions and current roadmap priorities target regions of benign terrain. While in-cave concepts have been postulated, the critical technologies have not been identified and demonstrated.

While robotic exploration of skylights and caves can seek out life, investigate geology and origins, and open the subsurface of other worlds to humankind, it is a daunting venture. Planetary voids present perilous terrain requiring innovative technologies for access, exploration, and modeling. These same technologies are broadly applicable to explorations of rough and/or subsurface planetary environments, including caves, craters, cliffs, and rock fields.

This research speculates on the possibilities and means of such exploration with fundamental contributions to exploring, modeling, and visualizing this new class of large-scale, highly three-dimensional concave planetary features.

High impact results of this work include:

- Innovative techniques for super-resolution modeling of pits and caves by fusing camera and LIDAR data and demonstrated those techniques in analog terrestrial environments.
- Innovative visualization techniques to communicate pit and cave models to humans, including building 3D-printable models from robot data and rendering mesh models with high-contrast illumination, and applied these and other visualization techniques, such as point rendering for cloud displays, mesh modeling, and stereo anaglyph, to models generated in field experiments.
- Conception and demonstration of access to and modeling of pits by perimeter traverse, Tyrolean traverse, and sloped descent.
- Field experiments at analog pit and cave sites, capturing thousands of images and hundreds of gigabytes of other sensor data
- Development of mission concepts and technology roadmap to achieve pit and cave exploration.

Pit and cave exploration enables a future of robots and humans collaboratively exploring and living in caves on other planets, and potentially enabling discovery of life on Mars. This report lays out the future made possible by pit cave exploration and presents the critical technologies and proofs of concept of those technologies.




Figure 3: Research conceived and demonstrated access and modeling of pits by Tyrolean traverse (left & top right) and sloped descent (bottom right).



Figure 4: Research conducted extensive field testing and explored existing and new methods for visualization of robot data, including mesh models (left), point rendering (center left), 3D printing (center right), and anaglyph stereo (right).

2 TABLE OF CONTENTS

1 EXECUTIVE SUMMARY	III
2 TABLE OF CONTENTS	V
3 INTRODUCTION	1
3.1 ELEMENTS OF PIT AND CAVE EXPLORATION	1
3.1.1 MODELING AND VISUALIZATION METHODS FOR PIT AND CAVE EXPLORATION	1
3.1.2 ORBITAL RECONNAISSANCE	2
3.1.3 FLYOVER AND SURFACE RECONNAISSANCE	2
3.1.4 PIT ACCESS AND DESCENT / SUBSURFACE EXPLORATION	2
3.2 MOTIVATIONS FOR SKYLIGHT/CAVE EXPLORATION	4
3.3 WHAT IS KNOWN ABOUT PLANETARY CAVES?	5
3.4 RELATED WORK	9
4 MISSION OBJECTIVES FOR PIT AND CAVE EXPLORATION	11
4.1 SCIENCE OBJECTIVES	11
4.1.1 PIT CHARACTERIZATION AND MORPHOLOGY	11
4.1.2 VOID CHARACTERIZATION AND MORPHOLOGY	11
4.1.3 FORMATION MECHANISMS	11
4.1.4 MINERAL DEPOSITS IN PITS AND VOIDS	11
4.1.5 VOLATILE DISTRIBUTIONS IN PITS AND VOIDS	12
4.1.6 ASTROBIOLOGY	12
4.1.7 SCIENCE OBJECTIVES ENGAGED BY THIS RESEARCH	12
4.2 HUMAN EXPLORATION MISSION OBJECTIVES	13
4.2.1 IDENTIFICATION AND CHARACTERIZATION OF FAVORABLE ENVIRONMENTS FOR HABITATION	13
4.2.2 METHODS AND INFRASTRUCTURE FOR CAVE AND TUBE EXPLORATION AND DEVELOPMENT	13
4.2.3 HUMAN EXPLORATION OBJECTIVES ENGAGED BY THIS RESEARCH	13
5 CHARACTERIZATION OF PIT AND CAVE EXPLORATION	14
5.1 VANTAGE POINTS FOR PIT AND CAVE EXPLORATION	14
5.1.1 ORBITAL RECONNAISSANCE AND MODELING	14
5.1.2 FLYOVER RECONNAISSANCE AND MODELING	14
5.1.3 SURFACE RECONNAISSANCE AND MODELING	15
5.1.4 PIT ACCESS/DESCENT	15
5.1.5 SUBSURFACE EXPLORATION	15
5.2 EXPLORATION TAXONOMY OF PITS	15
5.2.1 RAMP	15
5.2.2 CAVE ACCESS	16
5.2.3 LATITUDE	16
5.2.4 DEPTH TO DIAMETER RATIO	17

6	ROBOT ROLES & CONFIGURATIONS FOR PIT AND CAVE EXPLORATION	18
6.1	DEFINING ROBOT ROLES IN PIT AND CAVE EXPLORATION	18
6.1.1	ORBITAL EXPLORER	18
6.1.2	LANDER	18
6.1.3	FLYOVER EXPLORER	18
6.1.4	SURFACE EXPLORER	18
6.1.5	DESCENT INFRASTRUCTURE EMPLACER	19
6.1.6	PIT DESCENDER	19
6.1.7	PIT EXPLORER	19
6.1.8	PIT INFRASTRUCTURE LINK	19
6.1.9	SUBSURFACE EXPLORER	19
6.2	PIT DESCENDER AND PIT EXPLORER: ROBOTIC CONFIGURATION STUDY	20
6.2.1	PRECISION LANDING ON A PIT FLOOR	20
6.2.2	LEAPING INTO A PIT WITH A CAVEHOPPER ROBOT	20
6.2.3	TETHERED RAPPEL DESCENT	22
6.2.4	ROBOTIC DESCENT OF A SCREE SLOPE	23
6.2.5	DESCENT FROM A TYROLEAN LINE	23
6.2.6	ROBOTIC CONFIGURATION FEASIBILITY STUDY CONCLUSIONS	26
6.3	SUBSURFACE EXPLORER: MOBILITY CONFIGURATION TRADE STUDY	27
7	MISSION CONCEPTS FOR EXPLORATION OF PITS AND CAVES	30
7.1	SCOUT: FLYOVER AND SURFACE RECONNAISSANCE	30
7.1.1	MISSION OBJECTIVES	30
7.1.2	CONCEPT OF OPERATIONS	30
7.1.3	SCIENCE AND HUMAN EXPLORATION CONTRIBUTIONS	31
7.1.4	POTENTIAL LUNAR MISSION TARGET	32
7.2	WAYFARER: SURFACE RECONNAISSANCE	33
7.2.1	MISSION OBJECTIVES	33
7.2.2	CONCEPT OF OPERATIONS	33
7.2.3	SCIENCE AND HUMAN EXPLORATION CONTRIBUTIONS	34
7.2.4	POTENTIAL LUNAR MISSION TARGET	34
7.3	SPELUNKER: PIT ACCESS AND SUBSURFACE EXPLORATION	36
7.3.1	MISSION OBJECTIVES	36
7.3.2	CONCEPT OF OPERATIONS	36
7.3.3	SCIENCE AND HUMAN EXPLORATION CONTRIBUTIONS	37
7.3.4	POTENTIAL LUNAR MISSION TARGET	38
7.4	ONE-SHOT SPELUNKER: FLYOVER AND SURFACE RECONNAISSANCE, PIT DESCENT, AND SUBSURFACE EXPLORATION	38
7.4.1	MISSION OBJECTIVES	38
7.4.2	CONCEPT OF OPERATIONS:	38
7.4.3	SCIENCE AND HUMAN EXPLORATION CONTRIBUTIONS	39
7.4.4	POTENTIAL LUNAR MISSION TARGET	39
8	MODELING METHODS FOR PITS AND CAVES	40



8.1	SENSORS	40
8.1.1	CAMERA	40
8.1.2	LIDAR	43
8.1.3	SPECTROMETER	44
8.1.4	GROUND-PENETRATING RADAR	44
8.1.5	GRAVIMETER	44
8.1.6	INERTIAL MEASUREMENT UNIT (IMU)	44
8.1.7	OTHER SENSORS	45
8.2	COMMON MODELING METHODS	45
8.3	SENSOR FUSION	47
8.3.1	SENSOR FUSION BACKGROUND	47
8.3.2	LUMENHANCEMENT	48
8.3.3	LUMENHANCEMENT ALGORITHMS	49
9	CHARACTERIZATION OF EXPLORATION METHODS	56
9.1	ORBITAL RECONNAISSANCE & MODELING	56
9.1.1	STEREO FOR MODELING PIT SURROUNDINGS	56
9.1.2	MULTIVIEW STEREO MODELING OF A PIT FROM ORBIT	58
9.2	FLYOVER RECONNAISSANCE & MODELING	62
9.2.1	PRECISION LANDING AND HAZARD AVOIDANCE	62
9.2.2	FLYOVER MODELING	63
9.3	SURFACE RECONNAISSANCE AND MODELING	66
9.3.1	BUILDING 3D PIT MODELS FROM SURFACE ROVER IMAGES	67
9.3.2	HIGH DYNAMIC RANGE IMAGERY OF A PIT FROM A SURFACE ROVER	76
9.4	PIT ACCESS AND MODELING	78
9.4.1	PIT ACCESS MISSION SIMULATIONS	78
9.4.2	TYROBOT MODELING FIELD TEST – SCAN PATTERNS AND SPEEDS	80
9.4.3	TYROLEAN MODELING PERFORMANCE	84
9.5	SUBSURFACE EXPLORATION	92
9.5.1	CAVE MODELING BY STRUCTURE FROM MOTION WITH A HAND-HELD SENSOR	93
9.5.2	CAVE MODELING BY STRUCTURE FROM MOTION WITH ROBOT-MOUNTED SENSORS	95
9.5.3	CAVE MODELING WITH SOLID STATE FLASH LIDAR	97
9.5.4	HIGH RESOLUTION MODELING OF CAVE WALLS	100
10	VISUALIZATION OF MISSION DATA	104
10.1	POINT RENDERING FOR CLOUD DISPLAYS	104
10.2	MESH MODELS	107
10.3	ENHANCING USER UNDERSTANDING WITH HIGH CONTRAST ILLUMINATION IN RENDERING	109
10.4	3D PRINTING	113
10.5	ANAGLYPH STEREO	114
11	ROBOT DEVELOPMENT	115
11.1	TYROBOT	115
11.1.1	CONFIGURATION & ROBOT DEVELOPMENT	116



11.1.2	DRIVE SYSTEM DEVELOPMENT & CHARACTERIZATION	120
11.1.3	TYROLEAN LINE TENSION TESTING – OPTIMAL LINE TENSION	122
11.1.4	CABLE LAUNCHING AND ANCHORING	123
11.1.5	ROBOTIC RECHARGE FOR FLOOR EXPLORATION	124
11.1.6	SENSING PAYLOAD CONFIGURATION	126
11.2	KRAWLER	127
11.2.1	KRAWLER MOBILITY CONFIGURATION: EVALUATING SLOPE TRAVERSE	128
11.2.2	KRAWLER MOBILITY CONTROL: LEARNING TERRAIN TYPES FOR ACTIVE CONTROL COMPENSATION	128
11.2.3	KRAWLER POWER CONSIDERATIONS: ENERGETICS FOR LONG FORAYS	131
12	TECHNOLOGY ROADMAP	134
12.1	ENABLING TECHNOLOGIES FOR PIT & CAVE EXPLORATION	134
12.2	MISSIONS ROADMAP	135
12.3	TECHNOLOGY AREA DESCRIPTIONS	136
12.3.1	T1. AUTONOMY	136
12.3.2	T2. SENSING & MODELING	138
12.3.3	T3. POWER & COMMUNICATIONS	140
12.3.4	T4. MECHANISMS	141
12.3.5	T5. ANCILLARY BENEFICIAL TECHNOLOGIES	142
13	IMPACT	144
13.1	PUBLICATION	145
13.2	RELATED THESIS WORK	146
13.3	MISSION DEVELOPMENT	146
13.4	RELATED RESEARCH AND FUNDING	148
13.5	SCIENCE EFFORTS	150
14	CONCLUSIONS	151
15	ACKNOWLEDGEMENTS	151
16	REFERENCES	152
APPENDIX A: ANALOG FIELD TESTING SITES		158
MINE PITS		158
PLUTO’S CAVE		159
KING’S BOWL PIT		161
INDIAN TUNNEL		162
APPENDIX B: CRATERS OF THE MOON FIELD DEMONSTRATION AND SURVEYING		165
KING’S BOWL		165
ADDITIONAL SURVEY DATA COLLECTED AT KINGS BOWL		167
INDIAN TUNNEL		168

3 INTRODUCTION

Caves on other planetary bodies are enticing but challenging targets for exploration. Skylights, formed by cave ceiling collapse, are access points proven to exist on the Moon and Mars and presumed to exist on other planetary bodies. These subsurface features can be windows into a planet's geology, history, and perhaps biology, and may one day be safe havens for human missions. We must learn how to safely model, enter, and explore pits and caves to unlock their treasures.

This research developed technologies and mission architectures for robotic exploration and modeling of pits (i.e., dead ends and potential skylights) and caves. Missions are motivated by both science and human exploration objectives (Section 4). Science objectives include visual and non-visual characterization of skylights and caves, plus techniques to study their origins, morphology, and volatiles and mineral deposits. Human exploration requires assessing a skylight's suitability for safe entry and habitation, plus techniques for developing subsurface infrastructure.

3.1 ELEMENTS OF PIT AND CAVE EXPLORATION

There are multiple elements of pit and cave exploration, which progressively advance our understanding of these features (Section 5.1). Each element builds on the ones that precede it, amassing scientific and exploratory knowledge while informing development of the craft and sensors for the next. The exploration elements are orbital reconnaissance, flyover reconnaissance, surface reconnaissance, pit access and descent, and subsurface exploration.

A taxonomy of pits and caves based on exploration-relevant features (e.g., ramp type, subsurface cave access, pit latitude, and pit depth to diameter ratio) guides decisions on craft, sensors, and mission objectives. (Section 5.2)

Based on the exploration elements and the pit taxonomy, this research defined robot roles for exploration of pits and caves (Section 6). Applying these robot roles to the exploration elements and taxonomy yielded a suite of mission concepts of varying complexity (Section 7). *Scout*, a flyover and surface reconnaissance mission, models an individual pit and the surrounding terrain and searches for evidence that the pit provides cave access. *Wayfarer*, a surface reconnaissance mission, seeks to model multiple pits and surrounding surface features or a single pit when flyover is not possible. *Spelunker*, a pit access and subsurface exploration mission, seeks to model a pit and any associated subsurface voids. The *One-Shot Spelunker* mission combines the Scout and Spelunker missions to address scenarios in which there is strong orbital evidence that a pit leads to a cave, or where the distance from Earth makes multiple missions impractical.

3.1.1 Modeling and Visualization Methods for Pit and Cave Exploration

Each of these missions depends on accurate modeling of surface and subsurface structure. This research evaluated a wide range of modeling methods, including some specifically developed for planetary exploration (Section 8). Stereovision triangulates position by identifying common features in images taken by multiple cameras. Structure from motion extends stereovision to handle the case where the relative orientations of the cameras are not known beforehand, enabling the use of a single camera and multiple perspectives to model a target. Lumenancement (Section 8.3.2), a fusion of LIDAR and visual data using a controlled lighting source, improves model quality by an order of

magnitude – previous results have indicated a 40x increase in measurement density and a 40% increase in range accuracy. Section 8.3.3 explains the algorithms behind Lumenhancement.

Humans must be able to understand the models generated by robotic explorers. This research evaluated a range of visualization methods. Point Rendering for Cloud Displays (Section 10.1) provides an immersive model that can be generated in real-time and streamed from a robot. Mesh rendering (Section 10.2) is better supported by general-purpose renderers and is therefore a better solution for Education and Public Outreach, for feeding 3D models from a reconnaissance mission into pit exploration robot design, and for 3D printing (Section 10.4). Anaglyph stereo (Section 10.5) enables 3D perception from a single image that can work in printed materials and with very inexpensive viewing hardware.

3.1.2 Orbital Reconnaissance

The goals of orbital reconnaissance are to identify pits on the surface, determine optimal sites for information gain, and map hazardous and safe terrain for landing and rover planning. This research evaluated modeling methods with real lunar data by developing a digital elevation map (DEM) with a 5m/pixel resolution and errors of less than 50m using the LRO Narrow Angle Camera images of a Lacus Mortis pit (Section 9.1) then comparing it to a 100m/pixel DEM from the LRO Wide Angle Camera.

3.1.3 Flyover and Surface Reconnaissance

Flyover and surface reconnaissance can model significant portions of a pit's structure (Sections 9.2 and 9.3). Modeling methods for flyover and surface reconnaissance were evaluated using simulation or terrestrial analog environments. Flyover modeling of a pit was evaluated on simulated LIDAR data (Section 9.2.2). Flyover modeling is enabled by precision landing technology (Section 9.2.1) that reliably guides landers within 30m of their intended trajectory during the final 500m of descent.

Field tests for surface modeling in a terrestrial analog environment determined that the vast majority of the points in these models have distance errors in the single-digit centimeter range. This research also developed a method to autonomously plan views for surface modeling that account for transient illumination (Section 9.3).

3.1.4 Pit Access and Descent / Subsurface Exploration

Pit access and descent robots provide detailed information about the pit structure and provide unique perspectives relative to surface robots. These robots can also deliver payloads to the pit floor to investigate any caves that are discovered. This research evaluated a series of robot configurations (Section 6.2) and selected a Tyrolean robot for further study. The Tyrolean robot was field tested and optimal scan patterns and speeds were determined (Section 9.4). Post calibration and alignment results demonstrated an average measurement error of .47m and a median error of .31m.

This work further developed Tyrobot (Section 11.1), which employs an adjustable idler pulley to enable the robot to deploy to an existing Tyrolean line. Field and laboratory tests were conducted to determine optimal pulley design and line tension (Sections 11.1.2 and 11.1.3) as well as evaluating a method for Tyrolean line deployment (Section 11.1.4).

A crawling robot configuration, *Krawler* (Section 11.2), was further developed for pit access and subsurface exploration. *Krawler* employs a light and flexible chassis, elastic suspension, high ground clearance, and low center of mass to navigate extreme terrain. Field tests demonstrated that, for small



diameter pebbles, Krawler was able to ascend and descend the piles even at 30% power using either a linear or swerving s-curve trajectory strategy (Section 11.2.1). Larger rocks introduced slip and bouncing that altered performance in heading and trajectory. Using machine learning techniques in these field tests, Krawler was able to classify the surface materials it traversed by size with 81% accuracy and by size and contour with 70.25% accuracy (Section 11.2.2). This information can be used to improve mobility and control over different terrain. With power a driving constraint for subsurface exploration, an energetic model for mobile robots for long forays was constructed, relating robotic power consumption, mobility power consumption, mission time, and range (Section 11.2.3).

A robotic recharging technology and sensor configuration were developed to support subsurface exploration and pit modeling (Sections 11.1.5 and 11.1.6). Field tests in expected and extreme conditions demonstrated an overall recharging success rate of 62.5%. With conditions controlled by conops and reasonable contamination levels used, the success rate increased to 100%.

Finally, this research generated a technology roadmap toward subsurface exploration missions, including key enabling technologies and precursor missions (Section 12).

3.2 MOTIVATIONS FOR SKYLIGHT/CAVE EXPLORATION

Cave entrances have been conclusively shown to exist on Mars (Cushing, Titus and Maclennan 2011) (G. E. Cushing 2012) and the Moon (Ashley, et al. 2011). There is also evidence supporting their existence on other planetary bodies throughout the solar system (Ashley, et al. 2011). Despite astonishing discoveries of skylights and cave entrances, and their inevitable exploration, missions are not yet planned. Skylights and the voids below are so unknown that it is too risky to send humans to explore them without prior robotic reconnaissance and modeling.




Figure 5: Possible skylights on Mars (Images from a presentation by Glen (Cushing, Titus and Maclennan 2011))




Figure 6: Three views of the Mare Tranquillitatis skylight on the Moon. In the first image the camera is close to the nadir direction; three boulders can be seen marking the position of the skylight wall. As the viewing angle increases, void space under an overhanging ceiling can be observed. (Images from a presentation by James Ashley (Ashley, et al. 2011))

While robotic exploration of skylights and caves can seek out life, investigate geology and origins, and open the subsurface of other worlds to humankind, it is a daunting venture. Planetary voids present perilous terrain that requires innovative technologies for access, exploration, and modeling. The robots that venture into caves must leap, fly, or rappel into voids, traverse rubble, navigate safely in the dark, self-power, and explore autonomously with little or no communication to Earth. Exploiting these features necessitates a leap of technology from current planetary missions, which land with large error ellipses in statistically safe terrain, rove slowly and cautiously across the surface, depend on the Sun for power and light, and rely on constant human oversight and control.

3.3 WHAT IS KNOWN ABOUT PLANETARY CAVES?

Even now, when caves have been proven to exist on the Moon and Mars, Earth analogs are the best sources of information about planetary caves as satellites provide only limited and low-resolution views into subsurface features. Known mechanisms for cave formation on Earth – lava flows, volcano-tectonic fractures, and chemical dissolution – are likely to form caves on other planets as well.

Lava tube caves are formed by volcanic activity; the top layer of a channel of lava cools and forms a crust, leaving a void space when the hotter lava in the center of the channel flows out. Lava tubes tend to have smooth floors, and they may have “soda straw” stalactites formed by lava dripping from the ceiling. Sinuous rilles visible on the lunar surface were likely formed by lava tube collapse (Oberbeck, Quaide and Greeley 1969), and lava tube



Figure 7: Lava tube cave (Photo courtesy USGS)

structures have also been identified on Mars (Bleacher, Greeley, et al. 2007) (Bleacher, Greeley, et al. 2007). Due to the lesser gravity, it is predicted that lava tubes on Mars or the Moon may be much larger in diameter than those found on Earth (Coombs and Hawke 1992).

Caves can form when tectonic plates shift relative to each other and leave void spaces. In contrast to lava tubes, volcano-tectonic fracture caves are less sinuous; they are likely to be straight or slightly curved (G. E. Cushing 2012). The fractures can extend kilometers beneath the surface and may be partially filled from the bottom by magma (G. E. Cushing 2012).




Figure 8: Sinuous rilles on the Moon. Location of the Marius Hills pit is marked (Ashley, et al. 2011).




Figure 9: Karst-like features on Titan (top) compared to Karst on Earth (bottom) (Mitchell & Malaska, 2011)



Figure 10: Stalactites, stalagmites and columns in limestone cavern

Caves can also form when rock is dissolved by chemical means. Limestone caverns commonly found on Earth result when limestone is dissolved by water that has become slightly acidic through absorption of carbon dioxide. Karst is a name for the rock formation caused by dissolution of bedrock – the same dissolution that causes caves also results in karst formations. Karst-like features have been observed on Titan (Mitchell and Malaska 2011). Limestone caves on Earth tend to include sequences of chambers at many different levels, as opposed to the long, continuous, gently sloping caverns in lava tubes. They often have many stalactites and stalagmites, formed when minerals are deposited by the flow of the dissolving liquid.




Figure 11: Volcano-tectonic fractures on the Earth (top) and Mars (bottom) with potential cave entrances [Cushing 2011]

Skylights¹, formed by cave ceiling collapse, can provide entrance into caves. Several skylights on the Moon and Mars have been characterized from orbital image data. Skylight diameters can be determined by counting pixels in an image of known resolution. Shadow measurements provide rough estimates of skylight depth. More detailed information can be gained from stereography – matching features between images taken from different perspectives. A digital elevation model (DEM) of a funnel-shaped pit on Mars was generated through this method (McEwen and Mattson 2013). In high-resolution images, the dimensions of large blocks on a skylight floor can be measured, and terrain roughness on a scale below image resolution can be estimated from the standard deviation of surface reflectance, with a higher standard deviation indicating rougher terrain (Robinson, et al. 2012).

Of the three lunar skylights that have been studied in detail (see Figure 12), diameters range from 49m (short diameter of Marius Hills skylight) to 104m (long diameter of Ingenii skylight), and depths range from 38m (shallow end of Ingenii skylight) to 107m (Tranquillitatis skylight) (Robinson, et al. 2012). A fracture cave skylight examined on Mars (see Figure 13) has diameters from 68m to 48m; its depth was measured at 37m, but may be as shallow as 19m in the skylight center (G. E. Cushing 2012). A more circular Martian skylight (see Figure 14) (a) has a diameter of approximately 65m and a depth 45m or greater (G. E. Cushing 2012). One particularly interesting Martian skylight, shown in Figure 15) (b), sits at the bottom of a pit crater. This skylight is approximately 40m across, 50m below the surface and 25m deep (G. E. Cushing 2012).

High-energy impact can also cause flows of molten rock. A number of pits have also been identified in lunar impact melts. While these pits are smaller and less well understood than the three skylights discussed above, they may also lead into caves (Robinson, et al. 2012).




Figure 12: Lunar Skylights: Mare Ingenii Skylight (left); Mare Tranquillitatis Skylight (center); Marius Hills Skylight (right).

¹ For clarity in this work, a skylight is defined as an entrance to a cave from above, without regard to the formation

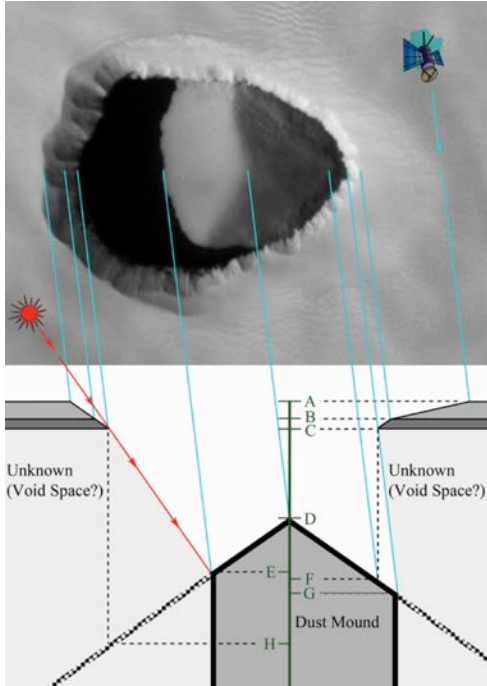


Figure 13: Fracture cave skylight on Mars




Figure 14: Martian skylights

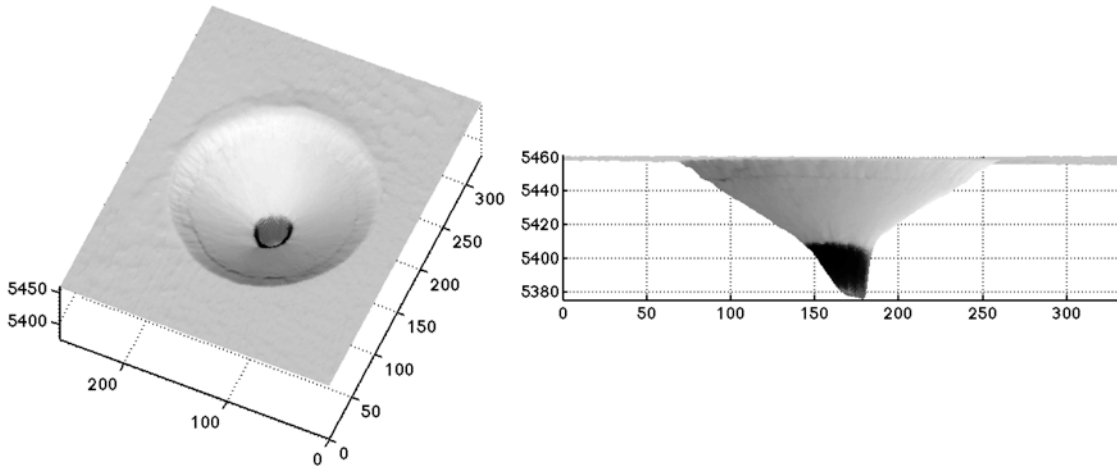


Figure 15: This NIAC team created a 3D mesh from a stereo DEM of a Mars pit (McEwen and Mattson 2013) to facilitate visualization of pit shape. The dark portion of the mesh in the right figure may be a cave entrance. Axes denote meters.

Caves are unique environments where both secondary mineral precipitation and microbial growth are enhanced by stable physico-chemical conditions. They represent excellent locations where traces of microbial life, or biosignatures, are formed and preserved in minerals. Caves on Mars may contain a record of secondary mineralization that would inform us on past aqueous activity. They may also represent the best locations to search for biosignatures (Léveillé and Dattab 2010).

3.4 RELATED WORK

A prior NIAC Phase I study by Werker, et al. (Werker, et al. 2001) studied the scientific value of exploring caves on other planets. This research speculated on planetary cave value by comparing to scientific knowledge gained through investigation of terrestrial caves. This study specified devices and infrastructure required to execute subsurface planetary exploration, including communication networks, biological sensing, and drilling capabilities.

Dubowsky, Iagnemma, and Boston (Dubowsky, Iagnemma and Boston 2005) (Dubowsky, Plante and Boston 2006) proposed exploration of subsurface voids with a large team of expendable robots. These robots were self-contained spherical hopping robots weighing approximately 100g with a 100mm diameter. The rationale behind this development was that traditional wheeled rovers such as *Sojourner* or *Curiosity* are not well suited to navigate through extremely rough terrain or access the highly sloped surfaces anticipated in subsurface environments. Additionally, Dubowsky, Iagnemma, and Boston opted for a large team of small-scale, low-cost robots because large rovers were deemed too valuable to risk entrapment (Dubowsky, Iagnemma and Boston 2005).

Prior academic research has addressed robotic model generation of terrestrial voids. Carnegie Mellon University has performed extensive research in this domain, publishing algorithms to solve localization, feature extraction and scan matching problems in a cavern-like environment. Wong et al. (Wong, Garney, et al. 2012) demonstrated significantly improved modeling in caves using range scanners and sampling the scene with a Nyquist criterion. When venturing into unknown cave environments with no access to absolute localization methods such as GPS, a robot must solve the Simultaneous Localization and Mapping (SLAM) problem. Fairfield, Kantor and Wettergreen presented approaches for SLAM applied to a robot exploring underwater caves (Fairfield, Kantor and Wettergreen 2005) (Fairfield, Kantor and Wettergreen 2006) (Fairfield, Kantor and Wettergreen 2007) (Fairfield, Kantor and Wettergreen 2010). Robot motion on natural surfaces has to cope with changing yaw, pitch, and roll angles, making pose estimation a problem in six mathematical dimensions. Nutecher et al. (Nutecher, et al. 2004) developed a fast variant of the Iterative Closest Points algorithm that registers 3D scans in a common coordinate system and re-localizes the robot. Consistent 3D maps can then be generated using a global relaxation. Zlot and Bosse coupled measurements from a spinning, scanning LIDAR with data from an inertial measurement unit to achieve SLAM from a moving platform that built a 3D model for 17km of mine tunnel (Zlot and Bosse 2012). Prior work also encompasses planning for subterranean exploration and mapping (Morris, et al. 2006) (Thrun, Thayer, et al. 2004) and science autonomy (Wagner, et al. 2001) (Wettergreen, et al. 2005).




Figure 16: Axel, a tethered two-wheeled rover to descend steep slopes [35].

(Bares and Wettergreen 1999) developed and fielded *Dante II*, a tethered, four-legged walking robot that successfully descended into the Mt. Spurr volcanic crater. Dante II carried seven video cameras, a scanning laser rangefinder, and gas detection sensors in order to conduct its science mission along the walls and on the bottom of the crater. Dante II used its self-contained winch to rappel down a slope of up to 90° with a cross-slope angle of up to 30°. It could step onto a 1.3m boulder on flat terrain and a 1.0m boulder on a 30° cross-slope. The mission into the crater was a success; however, on its way out of the crater, Dante II fell onto its side when it ran into a steep slope

coupled with cross-slope angle and a soft, unstable slope material. The robot was fully functional but unable to right itself, needing to be retrieved by helicopter.

(Nesnas, et al. 2008) at NASA's Jet Propulsion Laboratory developed the Axel family of two-wheel robots to provide versatile mobility in difficult planetary terrain with minimal complexity. The Axel rovers have only three actuators, one for each wheel and one for the trailing link, which can connect to other robots. One configuration of this system connects a single Axel robot to a larger rover, overcoming steep terrain by lowering the Axel robot by a tether. This enables exploration down steep slopes and cliff faces.

(Spenko, et al. 2008) developed RiSE, a legged machine capable of locomotion on both the ground and a variety of vertical building surfaces (e.g., brick, stucco, and crushed stone) without the use of suction, magnets, or adhesives. It achieves these capabilities through a combination of bio-inspired and traditional design methods. RiSE combines innovative body morphology, hierarchical compliance in the legs and feet, and sensing and control systems that enable robust and reliable climbing on difficult surfaces.



Figure 17: Cliff-bot utilizes two robots as anchors with tethers to assist in the motion of a rappelling robot. [37,38]

(Pirjanian, et al. 2002, Huntsberger, et al. 2003) developed Cliffbot, a robotic system to explore the surface of cliffs with up to a 70° slope. Cliffbot utilizes a trio of wheeled rovers; two robots act as the anchors at the top and use two tethers to assist in the motion of the third robot that rappels down and along a cliff. This system enables the robot to rappel down a steep slope and avoid obstacles by moving vertically and laterally on the slope with the assistance of the two tethers from the anchor robots.

4 MISSION OBJECTIVES FOR PIT AND CAVE EXPLORATION

Pit and cave exploration missions are valuable both for science and for human exploration. This section proposes specific objectives in each area.

4.1 SCIENCE OBJECTIVES

The general science objectives for robotic missions to planetary subsurface structures are to advance our understanding of these structures throughout the solar system, to use them as windows into each planet's geology and formation, and to advance the search for life. This research classifies specific science objectives into six categories:

1. Pit Characterization and Morphology
2. Void Characterization and Morphology
3. Formation Mechanisms
4. Mineral Deposits in Pits and Voids
5. Volatile Distributions in Pits and Voids
6. Astrobiology

4.1.1 Pit Characterization and Morphology

Pit characterization will address fundamental questions: How deep are they? How smooth or blocky are pit floors, walls, and surroundings? How frequently do they open into voids? How large are the voids? Characterization of pits will also provide clues to possible pit origins and allow testing of different origin mechanisms. For example, some pits may form when the ceiling of a lava tube collapses, such as the 50m straight-walled pit in Marius Hills (Haruyama, et al. 2009). Several known pits, including the Marius Hills pit, harbor void spaces (Robinson, et al. 2012).

4.1.2 Void Characterization and Morphology

Voids accessible by open pits are of greatest interest. Missions can identify and characterize subsurface voids through thermal and gravimetry data through flyovers, surface exploration, and then subsurface examination. Flyovers and surface exploration can determine pit size, geometry, and possible zones of collapse in relative safety prior to direct subsurface exploration.

4.1.3 Formation Mechanisms

The recent discoveries of pits in orbital images are astounding and important, but not sufficient to reveal the underlying physics and events of their volcanic or geomechanical origins. Pit exploration missions will develop detailed physical models that inform study of the collapse mechanisms that open voids to form surface pits, and the mechanisms that soften the aprons and create scree features over time. Alternatively, these models may show that some planetary pits are formed by escaping gases, as does happen on Earth..

4.1.4 Mineral Deposits in Pits and Voids

Lava tubes offer surfaces on which unusual minerals may be deposited during or shortly after the lava level has begun to drop. The wall surfaces left behind are a vast, accessible observation domain for robot explorers.

4.1.5 Volatile Distributions in Pits and Voids

Pit exploration missions will also search for volatiles. Do pits serve as cold traps? How might the appearance of and access to volatiles be different inside caves? For example, if a lava tube on the Moon, with no atmosphere to drive oxidation, develops a glassy interior surface when cooled, might volatiles be observed on these surfaces as frost? Many tubes on Earth contain permanent or transient ices, but we do not yet know how this extrapolates to other planetary environments.

4.1.6 Astrobiology

Missions to pits on Mars or the icy moons of the outer solar system will seek secondary mineral precipitation and evidence of microbial growth or other biosignatures. Does secondary mineralization indicate past aqueous activity? Did caves at one time (or even now) contain microbial life? Pits and caves may have been more hospitable to life than the surface, or may simply have preserved the evidence better.

4.1.7 Science Objectives Engaged by this Research

From the categories outlined above, the following science mission objectives were engaged by this research:

- Characterize pit and cave morphology.
- Investigate pit/cave origins and dating.
- Determine distribution of blocks on the pit floor.
- Model tunnels and voids with ground penetrating radar, acoustic sensing, and/or gravimetry.
- Survey terrain precisely above potential tunnels to infer voids.
- Characterize existence/extent of void space accessible from pit.
- Map distribution of material types in walls and floor.
- Build scientific models of pit formation.
- Sample surfaces inside pit/cave to determine if volatiles are present. If they are, map distribution of volatiles inside pits and caves.
- Characterize layering and thermal properties.
- Identify layers in surrounding terrain.
- Investigate surface crystallization on pristine, non-regolith covered cave walls and floors.
- Examine lava tube cave walls and floor for gasses trapped during lava cooling.
- Search for signs of life (on Mars or the icy moons of the outer solar system).
- Investigate impact sites and any detectable remainder of meteorites.
- Conduct other scientific experiments enabled by robotic access.

4.2 HUMAN EXPLORATION MISSION OBJECTIVES

The study of pits and voids lays a foundation for human exploration of planetary bodies. Voids such as caves and tubes that are accessible through pits may provide shelter from harsh planetary environments, reducing the complexity of human survival beyond Earth. Robotic exploration can accelerate human planetary exploration missions by identifying and characterizing favorable environments for habitation, and by creating methods and infrastructure for cave and tube exploration and development.

4.2.1 Identification and Characterization of Favorable Environments for Habitation

Caves and tubes provide protection from micrometeoroid impacts, solar radiation, and thermal extremes. Robotic pit exploration methods will characterize these subsurface environments to inform the destinations and technologies developed for future human missions. This characterization will include the presence and size of subsurface voids, their entry geometry, and their structural integrity.

4.2.2 Methods and Infrastructure for Cave and Tube Exploration and Development

Human exploration and development of caves and tubes will require much of the same infrastructure as robotic exploration, including power, communications, and transportation. The methods for delivering an exploration robot to the pit floor and operating it throughout the subsurface can scale to the larger payloads required for human exploration.

4.2.3 Human Exploration Objectives Engaged by this Research

From the two categories outlined above, the following human exploration mission objectives were engaged by this research:

- Characterize structural integrity of pits and voids.
- Create 3D models of accessible void space to facilitate design of human habitats.
- Characterize the thermal environment in subsurface voids.
- Characterize the radiation environment in subsurface voids.
- Characterize existence/extent of void space accessible from pit; Demonstrate the transfer of payload to the floor of a pit.
- Demonstrate the transfer of power to the floor of a pit.
- Demonstrate communications from Earth to the floor of a pit / accessible void entrance.

5 CHARACTERIZATION OF PIT AND CAVE EXPLORATION

Exploring planetary pits and caves represents a new class of robotic mission that demands unique mission characterizations and exploration taxonomies. Once exploration objectives are identified, the first determination is the vantage point(s) from which the spacecraft or rover will explore the target. The main alternative vantage points – orbital, flyover, surface, pit descent, and subsurface -- drive selection of craft and choice of sensors, technologies and techniques for the mission. These issues are examined in Section 5.1. The second determination is the specific targets to be explored. A classification of these targets (discussed in Section 5.2) enables refinement of craft configuration(s), sensor suite selection, and other required technologies. Section 7 defines a series of mission concepts optimized for different vantage points and target classifications.

5.1 VANTAGE POINTS FOR PIT AND CAVE EXPLORATION

Pit/cave exploration and utilization is a campaign of diverse missions for diverse purposes. The choice of vantage point is a helpful organizing principle: orbital, flyover, surface, pit descent, and subsurface. As a general progression, each vantage point provides data to inform the design of mission, craft, and sensors for the next. Orbital data informs flyover and surface missions, which inform pit descents, which inform subsurface missions. However, missions will inevitably mix and match these vantage points.

5.1.1 Orbital Reconnaissance and Modeling

Orbital reconnaissance can identify pits on a planetary surface. Once a pit is identified, a reconnaissance satellite can be directed to collect more data about it and the pit can be coarsely modeled. Orbital data informs selection of a pit as mission destination. Models built from orbit will be used to select a mission concept, to map hazardous and safe terrain for landing, and to aid rover planning. An illustration of how orbital reconnaissance & modeling are performed is given in Section 9.1.

5.1.2 Flyover Reconnaissance and Modeling

The resolution of models from orbital data is necessarily limited by sensing from orbital distances. Even focused laser beams from lunar orbit spread into multi-meter spots on the surface. Orbital sensing on Mars and other planets is further hampered by atmosphere. Only a surface mission can build a higher resolution model.

A lander can carry other robots to the surface, but it can also serve as a flyover explorer, passing over a pit during descent. Flyover reconnaissance captures the same bird's-eye views seen from orbit but at much higher resolution. Data from flyover can be used for planning surface or descent exploration by other robots. Flyover modeling is discussed in greater detail in Section 9.2.




Figure 18: Orbital reconnaissance identifies pits and potential voids and gathers data to produce coarse models of the pits. Image Source:

<http://lunar.gsfc.nasa.gov/spacecraftillustrations.html>

5.1.3 Surface Reconnaissance and Modeling

Although flyover captures birds-eye views at high resolution, a lander’s descent to the surface happens quickly. Surface robots can linger and maneuver to capture data from more perspectives and lighting conditions. If scientists in mission control develop new goals after seeing flyover data, surface reconnaissance robots can respond. Due to viewing geometry, surface modeling is better than flyover for capturing details in pit walls. Surface reconnaissance may be able to confirm the presence of a cave. If so, surface reconnaissance models can help plan a mission to descend into the pit. A detailed description of surface reconnaissance and modeling is provided in Section 9.3.




Figure 19: Flyover and Surface Reconnaissance captures data of the same region, but from different perspectives.

5.1.4 Pit Access/Descent

Some pit views and science are simply not possible without descent into the pit. This is a major mobility challenge for a robot. The best approach depends on the shape of the pit. A few pits have ramps. Most have steep walls that require alternative solutions. Pit descent serves two purposes: to obtain data from inside the pit and to reach any accessible caves (or tubes). An evaluation of robots for accessing pits is presented in Section 6.2. Information on modeling evaluations performed in the field as part of this work is presented in Sections 9.4.

5.1.5 Subsurface Exploration

Once a subsurface robot reaches a cave (or tube) by descent into a pit, it can explore the cave’s extent. The distance that a robot can travel into a cave is limited by its ability to overcome rough terrain (such as rubble piles) and by its energetics, since solar power will not be available. Communication from inside a cave will be difficult or impossible once the robot passes beyond line of sight to the entrance, so subsurface explorers must foray and return to relay data. Subsurface exploration robots can model the cave morphology, evaluate habitability of the caves for human missions, and investigate subsurface minerals and volatiles. An evaluation of subsurface exploration robots is presented in Section 6.3. Field test results for subsurface exploration methods are presented in Section 9.4.1.

5.2 EXPLORATION TAXONOMY OF PITS

Robotic configuration to access and explore pits and caves is intimately intertwined with key aspects of the target voids. As an organizing principle to correlate target sites with mission capabilities and objectives, this research developed a taxonomy with the following axes: Ramp, Cave Access, Latitude, and Depth-to-Diameter Ratio.

5.2.1 Ramp

The Ramp axis leads to different choices for pit access/descent. Planetary pits have been observed along a spectrum from purely vertical descent to gentle slopes (see Figure 20). A ramp is defined here as

a continuous, non-vertical region of terrain that connects the surface with the pit floor. In other words, a ramp is terrain that can be followed from the surface to the pit floor without crossing significant gaps or overhangs. Pits can be divided into classes of (1) Gentle Entry Ramp, (2) Steep Entry Ramp, or (3) No Entry Ramp. For Gentle Entry Ramp pits like the Lacus Mortis pit (see Figure 20, right), a 4- or 6-wheeled rover with good terrainability could do a free descent of the slope. For Steep Entry Ramp pits, a 4- or 6-wheeled rover may still be a viable choice, but it would likely have to be tethered to one or more anchor points during descent. For No Entry Ramp pits (see Figure 20, left), robots would have to rappel, fly, leap, or be lowered into the pit.




Figure 20: Pit with no entry ramp (left) and Pit with gentle entry ramp (right)

5.2.2 Cave Access

The Cave Access axis affects a pit's applicability for subsurface science or safe haven. It also affects whether robots specialized for subsurface exploration are needed or useful for exploration of that pit. Pits can be divided into classes of (1) Large Cave Connected, (2) Some Cave Connected, or (3) No Cave. It is probably impossible to determine if a pit belongs to the Large Cave Connected class using orbital data. Distinguishing between the Some Cave Connected and No Cave classes may be possible from orbit, if the right orbital data are collected. For the Mare Tranquilitatis pit, for example, a cluster of 3 boulders can be observed near the edge of the pit floor in a nadir image, and a greater and greater distance between the boulders and the pit wall is observed in more oblique images (see Figure 6). This indicates that the wall overhangs the pit: there is at least some cave available. Not all pits have such obvious features to facilitate the detection of caves. Gravimetry or ground penetrating radar measurements may determine if there is un-collapsed void in the vicinity of a pit, but it will not necessarily determine if the void is accessible from the pit.

5.2.3 Latitude

The latitude of a pit affects the methods that can be used to detect and model it from orbit and to explore it with lander and rover reconnaissance. It also affects whether the pit could be a cold trap for volatiles. Pits can be divided into classes of (1) High Latitude, (2) Mid Latitude, or (3) Low Latitude.

For High Latitude pits, permanent shadows will cover most or all of the pit floor. It may not be possible to estimate the pit depth from shadow length, making it difficult to distinguish pits from craters, and thus difficult to detect pits in orbital images. For flyover and surface reconnaissance, active sensing

(LIDAR, or pairing cameras with illumination sources) would greatly increase probability of finding pits or quality of mapping.

For Mid Latitude pits, some portion of the pit floor will be in permanent shadow whether or not a cave is present. Little or no information about this shadowed region can be determined from orbit – at best, reflected sunlight from a pit's illuminated walls may provide some light in shadowed regions. Surface reconnaissance robots may be able to peer into shadowed regions with passive cameras if high dynamic range imaging is used. Shadows in orbital images of a pit can be used to estimate the depth of the pit and (with sophisticated computer vision techniques) to model pit shape.

For Low Latitude pits, the entire pit floor is illuminated at some point during the local day or year. Areas of permanent shadow only exist if there are significant caves or overhangs. Passive sensing is a viable technique for lander and rover reconnaissance, if done at an appropriate time of day. As with Mid Latitude pits, shadows in orbital images can be used to estimate pit depth and model pit shape.

5.2.4 Depth to Diameter Ratio

The depth to diameter ratio affects how observable a pit floor and walls are from the surface. This also, as a function of latitude, determines whether active sensing is needed to observe the pit floor, and whether the pit could be a cold trap for volatiles. Pits can be divided into classes of (1) High Depth/Diameter Ratio, (2) Medium Depth/Diameter Ratio, (3) Low Depth/Diameter Ratio.

For High Depth/Diameter Ratio pits, it is not possible to view the pit floor from a point on the surface without getting very close to the pit edge (which may be unsafe) or even extending sensors over the edge. A significant portion of the pit wall may not be visible from a surface rover. Flyover reconnaissance and modeling will be useful in capturing the pit floor, but will not be very helpful for the walls.

For Medium Depth/Diameter Ratio pits, some portion of the pit floor is viewable from a surface rover, but not all. Most or all of the pit wall is viewable from a surface rover. Flyover sensing is useful for the floor of these pits, and to a lesser extent for the walls.

For Low Depth/Diameter Ratio pits, the entire pit floor is viewable from a surface rover. Greater pit diameter means that views of the directly opposing wall will be relatively low resolution, which requires a change in optics (smaller field of view sensors) or in viewing strategy (collect and stitch multiple close, oblique views instead of taking images with small angles between the wall surface normal and the viewing direction). Pit access/descent strategies that require spanning the pit with one or more cables may require too much mass, unless a partial span (i.e. not stretching across the diameter) is sufficient. This class of pit is easier to model from orbit and flyover than other depth/diameter ratio classes, since the floor is a large component of the pit surface area, and walls can be viewed obliquely from a greater distance.

6 ROBOT ROLES & CONFIGURATIONS FOR PIT AND CAVE EXPLORATION

This research defined robot roles to perform each element of pit and cave exploration, based on pit taxonomy and mission objectives. These roles are unique to pit and cave exploration; each was analyzed in relation to existing robotic capabilities. In some cases, the roles required development of new robot configurations (Section 11) and new capabilities for the robots performing these roles (Section 9).

6.1 DEFINING ROBOT ROLES IN PIT AND CAVE EXPLORATION

Any given mission will likely require some subset of the available robot roles, depending on the exploration elements pursued. One robot may fulfill multiple roles. Some roles can achieve the required mobility with well-established robot configurations, while others will require mobility specific to pit and cave exploration.

6.1.1 Orbital Explorer

Collect measurements from orbit that can be used to identify and model pits and, in some cases, confirm the existence of caves. Role required for the Orbital Reconnaissance and Modeling.

Existing robotic satellites already fulfill this role. Although additional sensor modalities, higher resolution sensors, and better localization could improve performance; further investigation with satellite data that has already been captured for the Moon (or Mars) must come first. Existing mobility configurations are sufficient.

6.1.2 Lander

Deliver other robots to a target destination from orbit. Role required for all elements of pit and cave exploration, with the exception of Orbital Reconnaissance and Modeling.

A Lander is necessary for any surface mission. While improvements in precision landing and hazard detection enable landing closer to pits, the overall mobility concept is not specific to pit and cave exploration.

6.1.3 Flyover Explorer

Collect measurements of a pit, and possibly caves, while flying over a pit at close range. Role required for Flyover Reconnaissance and Modeling.

The same robot can act as a Lander and Flyover Explorer. While it requires precision landing capability and may benefit from specialized sensors and sensor pointing, this mobility concept is not specific to pit and cave exploration.

6.1.4 Surface Explorer

Approach a pit from the surface, collect measurements of the pit and any visible caves to model and conduct scientific investigations. Role required for Surface Reconnaissance and Modeling.

Successful planetary surface rovers have been demonstrated on the Moon and Mars. Mobility improvements are possible but not unique to pit and cave exploration. Specialized sensors and planning algorithms for pit exploration may be needed.

6.1.5 Descent Infrastructure Emplacer

Conduct non-exploration tasks to facilitate further pit and cave exploration (e.g., anchoring tethers or other lines). Role may be required for Pit Access and Descent.

The requirements for this role depend on the method chosen for pit descent. A Lander could serve as an anchor for a tether or one side of a line stretching across a pit. For a second anchor point, a rover might drive the line around or drive to locate the end of a line shot across the pit (see Section 11.1.4). This line could be anchored simply by attaching to the rover if it has sufficient weight, which is not necessarily the case for lightweight planetary rovers, or it could be buried (Skonieczny 2013). The mobility configuration for these approaches can be the same as for the Surface Explorer.

6.1.6 Pit Descender

Move from the surface to the floor of a pit in a way that enables robotic activity on the pit floor or in accessible caves. Role required for Pit Access and Descent.

This role faces mobility challenges unique to pit and cave exploration, and thus is a focus of this work. Steep wall and sloped descent mobility configurations are explored.

6.1.7 Pit Explorer

Collect measurements to model and conduct scientific investigations inside a pit. Role required for Pit Access and Descent.

A Pit Descender could perform this role, though the Pit Explorer may have to operate in more constrained conditions. This role faces mobility challenges unique to pit and cave exploration, and thus is a focus of this work.

6.1.8 Pit Infrastructure Link

Provide a link between resources available at the surface (i.e., power and communications) and robots on a pit floor or at cave entrances. Role may be required for Subsurface Exploration. The pit taxonomy classification (i.e., pit depth to diameter ratio, pit latitude, etc.) helps determine whether this role is needed.

The Pit Descender could perform this role.

6.1.9 Subsurface Explorer

Collect measurements to model and conduct scientific investigations inside a cave. Role required for Subsurface Exploration.

This role faces mobility challenges unique to pit and cave exploration, and thus is a focus of this work.

6.2 PIT DESCENDER AND PIT EXPLORER: ROBOTIC CONFIGURATION STUDY

Five concepts for descending into and exploring a pit were considered: precision landing on a pit floor, leaping into a pit with a Cavehopper robot, tethered rappel descent, robotic descent of a scree slope, and descent from a Tyrolean line.

6.2.1 Precision Landing on a Pit Floor

A propulsive landing vehicle could fly over the surface to a pit, then descend directly to the floor and deploy an exploration rover. This enables study of the surface, pit floor, and pit walls, as well as delivering payloads to the pit floor.

While propulsive landers can access any pit with a sufficiently large aperture, independent of ramp type, their ability to study subsurface caves is dependent on their relative size and their supply of fuel.

Landing on the pit floor presents a simplified and streamlined mission that maximizes mass delivered to the subsurface, but it encounters a number of significant issues. The propulsion byproducts and any unused propellant could contaminate the void environment and interfere with science



Figure 21: Precision landing on a skylight floor

objectives. This problem has been addressed for bodies with atmospheres with a system that compresses atmospheric gases for later use in propulsion. This approach is not applicable to airless bodies such as the Moon or asteroids. The floors of skylights are also laden with obstacles, making landings high risk. The aperture of the skylight places additional constraints on power and communications beyond those experienced on the surface.

Because propulsive lander descent is controlled, sensing the pit during descent is possible with this configuration. Sensing issues are similar to those encountered for the Flyover Explorer. See Section 9.2.2 for investigation of this type of sensing.

6.2.2 Leaping into a Pit with a Cavehopper Robot

The Cavehopper configuration provides a simple descent with no special-purpose descent robot required, regardless of skylight geometry. Cavehopper robots utilize wheeled mobility over relatively flat terrain and a mechanical hopping system for surmounting large obstacles and entering pits. The leaping capability enables the surmounting of any number of obstacles, including large boulders, on the skylight floor. Difficulty arises in planning and controlling the hopping motion, which limits the accuracy of the hop trajectory and increases the probability of the robot becoming trapped. This configuration also requires a separate system for communication and power due to limits imposed by the skylight aperture. Cavehoppers are capable of accessing any pit regardless of taxonomy, but pits with greater depth present greater risk to entry.

Unlike a propulsive lander, a Cavehopper's descent is uncontrolled. Thus, the configuration's suitability as a Pit Explorer role required further investigation. An initiative to consider modeling in free-fall developed a sensor pod to be dropped from a strip mine high-wall with a parachute to reach lunar-relevant descent velocities and induce a slow spin to capture 360-degree data in descent. The dropped sensor pod featured a side-facing, wide-angle, high-frame rate camera commonly utilized in sports video. This configuration images the walls as it falls, exploring the question of whether a quality model could be generated by something that hopped into a skylight.




Figure 22: Cavehopper robot ready to execute a hop (left). Cavehopper leaping into a pit




Figure 23: Free-fall sensor pod (left) is slowed with a parachute for lunar-like descent velocity (right).

Results indicate that the free-fall approach has limited capability for modeling. The ballistic trajectory provides only seconds to acquire data, induces motion blur in data, and allows only limited control of viewpoints and sensing angles. Success of this approach for skylights is unlikely.




Figure 24: Still capture from the descent video stream demonstrates high motion blur and uncontrollable perspectives.

6.2.3 Tethered Rappel Descent

In a tethered rappel descent, a robot would rappel down a pit wall to reach the floor. This offers a controlled descent and facilitates scientific investigations that require contact with the pit wall, although only along a narrow corridor that descends from the tether anchor point. A specialized robot configuration would be needed to perform this operation. JPL's Axel rover and Cliffbot, both designed to descend steep slopes and cliff faces, are examples of robot configurations for this descent concept (see Figure 25, discussion in Section 3.4).

Tether management is a difficult problem. For a controlled tether rappel, the robot may pass "points of no return" during descent, which preclude subsequent ascent. The tether line will also likely drag in the surface regolith and provide complications in the reel mechanism should the tether be reeled in. These issues are particularly troubling if the robot needs to ascend in order to reposition for an alternate descent path or to alter position for more optimal camera and sensor views. The uncertainty in surface conditions and the exact shape and structural integrity of the skylight walls creates the possibility of the robot getting stuck during descent or the tether pulling surface or wall material on to the descending robot.




Figure 25: Cliffbot descending a rocky slope (top). Axel rover descending a slope, viewed from the tether anchor point (bottom). [Images credit: NASA/JPL]

Tethered rappel descent is controlled, and the robot can stop to take sensor measurements to act as a Pit Explorer. The robot's view of the pit wall or slope along which it descends is limited. Depending on the wall shape, overhangs may leave the robot dangling and swinging freely for part of the descent. This free-swinging robot may then impact the wall in uncontrolled ways, so placement of sensors, such that

they could both view without occlusions from other robot components and be protected from impact, would require careful analysis. The tether in this concept could also act as a Pit Infrastructure Link.

6.2.4 Robotic Descent of a Scree Slope

For pits that have ramps, the robotic descent of a scree slope uses a wheeled robot to descend to the pit floor and explore the void below. Such a rover could optionally include a nominally non-load-bearing tether for a reliable power and communications connection to the surface. Robots with low centers of gravity and conformable suspensions are capable of descending dramatic, steep slopes one way, as either an untethered robot or one with a power and communications tether only. One significant advantage of this concept is that the same rover that descends into the pit could also act as a Subsurface Explorer. See Section 11.2 for development of a robot for scree slope descent and cave exploration.



Figure 26: Lunar pit with a scree slope (top left), terrestrial scree slope (top right), and scree slope descender prototype (right).

A robot descending a scree slope can stop to take sensor measurements and act as a Pit Explorer. The robot’s view of the slope it descends will be limited. If a power and communications tether is included for this concept, it could act as a Pit Infrastructure Link.

6.2.5 Descent from a Tyrolean Line

Tyrolean traverse is a mountaineering term for crossing through free space between two high points on a rope strung between the two points. A Tyrolean robot traverses a pit via a Tyrolean line and has the capability to lower itself or a payload to the pit floor. The Tyrolean descent configuration provides a unique vantage point for science modeling with accurate deployment, lift, and redeployment of robots anywhere along the Tyrolean line during operations. This configuration is capable of accessing any pit independent of ramp type.

This configuration also provides an enhanced communications link to the surface and a movable power link to robots on the skylight floor that is unaffected by the skylight aperture. It does require a considerable development effort in line deployment and anchoring techniques as well as cable management during deployment and operations. It also adds an




Figure 27: Tyrolean robot exploring a subsurface void

additional unique robot to the mission that cannot explore beyond line-of-sight down subsurface tunnels. Tyrobot, the Tyrolean robot developed in this work, is discussed in Section 11.1. Prior to robot development, preliminary investigations were performed with a single point anchor to evaluate sensing during controlled descent.

6.2.5.1 Vertical Descent Sensing Evaluation – Single Point Anchor

Descent from a Tyrolean line can be as slow as necessary to facilitate sensing as a Pit Explorer. Terrestrial borehole robots are state-of-art descent mappers. They lower hundreds of feet on rigid tubing made of light material such as fiberglass to collect 3D models. This rigid deployment is highly impractical for planetary scenarios – the Tyrolean Robot concept would have to use flexible cables. While it is accepted that a quality model can be generated from a rigid, static base, it was unknown if a quality model could be generated from a sensor that sways, rocks, and spins.



Figure 28: Experimental Setup for Cable-Descent Modeling. A skid-loader and boom rigging were utilized to suspend the device over the skylight wall (left). Ferret borehole sensor scans the wall face while free-hanging (right). A green tether is utilized for data communications.

A borehole sensor, *Ferret*, was available for use in testing (see 11.1.6). Ferret has laser ranging, a fisheye camera, and inertial sensing. The inertial sensor enables determination of the undesirable ego-motion of the robot, which facilitates removal of these effects from the resulting model. Ferret was incrementally lowered on a cable from a strip mine high-wall while data was acquired. 3D LIDAR data is illustrated in Figure 29, while imagery from the fisheye camera is shown in Figure 30.




Figure 29: 3D LIDAR scan from Ferret borehole sensor. Two views are shown, a panoramic view of the pit (top) and close-up detail of the wall near the sensors (bottom).




Figure 30: Fisheye Image from Ferret borehole sensor. The circular image represents a 180-degree field of view.

6.2.5.2 Comparison of Tyrolean and Single Point Anchor Descent

In contrast to the free-swinging, single-cable descent mechanism demonstrated in the single-point anchor, the Tyrolean provides rigidity against rotation, which is invaluable for sensing. Moreover, the spanning line enables greater reach to access different views across the skylight. A single-cable system would require a prohibitively large boom to reach the same range of views.

6.2.6 Robotic Configuration Feasibility Study Conclusions

Based on the feasibility study, both the Tyrolean robot and scree slope descender were selected for further development. The Tyrolean configuration enables access to pit floors regardless of taxonomy and minimizes the impact of the pit aperture on communications and power to sensors and robots studying the pit floor and walls. Its dual-point anchor system provides a significant reduction in sensor oscillation during descent and an improvement in motion control for camera perspectives over the single-point anchor Ferret borehole sensor. While the configuration does introduce an additional robot to the cave exploration mission, that robot remains useful and expands capability to act as the Pit Infrastructure Link. The scree slope descender configuration utilizes a single robot for pit descent and subsurface exploration, which reduces mission complexity. A low center of gravity and conformable suspension enables the descent of dramatic, steep slopes.

While preliminary evaluation indicated that the Cavehopper design was promising due to its versatility and efficiency, the evaluation and preliminary experimentation presented here identified significant issues that disqualified it from further development. The difficulties in ensuring accuracy of the hopping targeting, execution, and control could result in the robot getting stuck in the targeted terrain. The Cavehopper also provides limited camera views, since it is unlikely to hop out of a pit and data gathered during descent suffers greatly from high motion blur and uncontrollable perspectives.

The tether rappel descent concept was eliminated due to difficulty of tether management while moving over planetary terrain and the danger that the robot could be damaged by loose rocks or regolith dislodged from the pit walls by its tether. The precision landing on the pit floor concept was eliminated due to the high risk of landing on an obstacle-ridden pit floor combined with its additional limitations on communications and power.

Section 11 details further development of robot configurations for descent from a Tyrolean line and robotic descent of a scree slope. Section 9.4.1 describes demonstration of pit access via these two methods. Section 9.4 describes evaluation of sensing from a Tyrolean robot.

6.3 SUBSURFACE EXPLORER: MOBILITY CONFIGURATION TRADE STUDY

A range of robot configurations, from flight-proven rovers to conceptual robot designs, was explored to address the mobility needs of the subsurface explorer role. Table 1 below describes candidate configurations considered. Each candidate was evaluated against 12 criteria: mass, boulder field traversability, flat ground driving efficiency, simplicity, power duration, reliability, controllability, communication, technology readiness, model registration (localization accuracy), environmental impact (i.e. contamination), and quality of data collection. Table 2 shows the weighted value of each criterion, the scores for each configuration, and the total weighted score used for configuration selection.

Table 1: Subsurface Mobility Configurations Considered

<p>Four-Wheel, No Suspension: Wheeled locomotion with flexible wheels and no suspension. This simple robot design has limited ability to traverse complex terrain on void floors, but high efficiency on smooth areas of tunnels and low complexity.</p>	
<p>Six-Wheel Rocker-Bogie: The Rocker-Bogie system is the suspension arrangement used in the Mars rovers. This passive suspension maintains six wheel contact on complex and rocky terrain. This mobility concept is well proven and handles significant obstacles, but is limited in the size and complexity of obstacles it can overcome in a challenging void environment.</p>	
<p>Four-Wheel Rock Crawler Suspension: This wheeled robot suspension is based on rock crawlers, vehicles that surmount significant bouldered terrain on Earth. Rock crawler suspensions on Earth are designed for terrain similar to what is anticipated on void floors, making it natural to adapt them to planetary pit exploration. They use spring-damper systems that have been avoided on planetary surface missions, but that may be suited to the more consistent thermal environments in planetary caves. These vehicles could surmount significant obstacles, but large boulders may prevent passage.</p>	
<p>Legged Robot: Legged locomotion improves navigability in rough terrain by stepping over obstacles that would stop a wheeled robot of comparable size. They are still limited by obstacles that exceed a certain fraction of their height, and by potential leg obstruction or entrapment. Downsides include lower power efficiency on flat ground and complex control for foot placement.</p>	
<p>Wheeled & Legged Robot: Mobility with six wheels at the ends of six multi-degree-of-freedom limbs, based on NASA's ATHLETE vehicle concept. Wheels enable efficient driving over stable, gently rolling terrain. Because each limb has enough degrees of freedom for use as a general-purpose leg, the wheels can be locked and used as feet to walk out of excessively soft, obstacle rich, or other extreme terrain. This mobility is exceptionally complex, but provides excellent mobility on obstacles and efficient travel on flat terrain.</p>	
<p>Segmented Tracked Robot: Segmented tracked vehicles have demonstrated excellent mobility in boulder fields. These robots easily roll over obstacles. Downsides include low driving efficiency of tracks, lack of a state-of-art track mechanism that operates in space environments (complicated by challenges like cold-welding), and no clear ongoing development that would advance tracks to mission readiness.</p>	
<p>Climbot: Robot that climbs to traverse void floors, boulders, and walls in the cave. Could avoid debris and obstacles by climbing the obstacles or the wall. Requires high mechanism and autonomy complexity, and has limited range due to higher power and slower traverse. Requires advances in robotic climbing and anchoring technology not yet proven for mission feasibility.</p>	

Spherical Hopping Microbots: Spherical, baseball-sized hopping robots for cave exploration, based on (Dubowsky, Plante, & Boston, 2006). Their mission concept launched many of these small robots into a cave, accepting that many would not survive, and used the surviving population of microbots for exploration, comm, and data return.



Telescoping Ball Robot: Has two mobility modes: enclosed in sphere for launching or rolling to access the cave, and then deploying legs or wheels for subsurface mobility. For example, deployment could extend circular halves of the sphere as wheels in a dumbbell shape and also extend a tail. The robot could be launched or dropped from a rover or a lander to reach the skylight floor. Robot would be battery powered and could be part of a multi-robot team with each robot acting as a wireless communication node. Alternately, it be launched from another robot to access an otherwise inaccessible cave region.



Cavehopper: combines hopping to overcome large obstacles with wheeled mobility for efficient flat-terrain mobility. Cavehoppers may also leap into a pit for descent. Hopping robots can be small and light, making them effective for tight spaces and economical to launch. By hopping, especially in the lower gravity encountered on many planetary bodies beyond Earth, they can overcome obstacles many times their own size. However hopping is unpredictable and dangerous on uneven terrain. It also complicates localization and modeling.



Snake Robot: Snake-like locomotion improves navigability in rough terrain by slithering over boulders and through small gaps not traversible by wheeled and legged robots. Control and sensor placement are challenging. Mechanisms are complex and are difficult to automate or to control remotely.



Propulsive Flying Robot: Accesses and traverses by small thrusters. Easily surpasses boulder fields like those near many pit entrances to reach the pit floor or cave entrance. Could carry enough fuel to make multiple trips and return to surface. Could combine with a tethered node to reduce trips to surface for power and communication. However, if volatiles exist trapped at the bottom of a skylight, they could be contaminated by a vehicle's thruster plume. Similarly, living organisms inside a cave could be killed if a vehicle's thruster plume contained toxic chemicals; this limits the applicability of propulsion in scientific exploration of caves.



Aerobot: In environments with atmospheres, like Mars, it is possible to fly using atmosphere rather than limited propellant. This has the benefits mentioned in a propulsive flyer to traverse boulder fields, with added benefits of longer flights for model data collection and lower expended power to enhance duration of operations.




Table 2: Weighted Selection Matrix

Selection Criteria	Weight	4-Wheel	Rocker-Bogie	Rock Crawler	Legged Robot	Wheel-Legs	Segmented Tracks	Climbot	Spherical Microbots	Telescoping Ball Robot	Cavehopper	Snake Robot	Propulsive Flyer	Aerobot
Mass	3	3	3	3	2	1	2	4	5	3	4	3	3	3
Boulder Field Traversability	5	1	3	4	4	5	4	4	4	3	4	4	5	5
Flat ground Driving Efficiency	4	4	5	5	2	4	2	1	2	4	5	1	1	2
Simplicity	3	5	4	4	3	1	3	1	3	2	2	1	4	4
Power Duration	2	4	4	4	3	2	2	1	2	3	3	2	1	3
Reliability	3	2	4	4	2	2	3	1	2	3	2	2	3	3
Controllability	2	4	3	4	2	2	4	2	1	3	1	2	4	5
Communication	3	3	3	3	3	3	3	3	2	3	4	2	4	4
Technology Readiness	5	5	5	5	3	2	3	1	2	4	2	1	3	2
Model registration (localization accuracy)	4	3	4	4	4	4	4	3	1	3	1	2	3	4
Environmental Impact (i.e. contamination)	3	5	5	5	5	5	5	5	5	5	5	5	1	5
Quality of Data Collection	5	2	4	4	4	4	4	3	2	2	1	2	2	4
Total Weighted Score		138	167	174	134	131	139	104	109	133	118	94	121	152

From the above analysis of candidates for the subsurface explorer role, this research recommends highly mobile, wheeled robots such as Six-Wheel Rocker-Bogie or Four-Wheel Rock Crawler Suspension. These would be the primary robots for subsurface mission tasks of exploration, modeling, and collecting science data. They are capable of traversing rocky terrain and steep slopes while exploring surface, pit access, and subsurface environments. The Krawler, a Rock Crawler robot developed in this work, is discussed in Section 11.2.



Figure 31: Crawler robot exploring the pit floor.

7 MISSION CONCEPTS FOR EXPLORATION OF PITS AND CAVES


The following mission concepts are intended as templates that can be specialized to pits on any planetary body. Potential lunar targets are identified for each mission concept.

7.1 SCOUT: FLYOVER AND SURFACE RECONNAISSANCE

Mission Summary


Lander models a skylight by flying over during descent. Rover deploys, drives to the pit, and circumnavigates to capture low-angle views. A 3D model is reconstructed from lander and rover data. Scout can determine accessibility and, to some degree, extent of void spaces; map material-type distribution; and identify pit wall layers.

Uses Lander, Flyover Explorer, Surface Explorer



7.1.1 Mission Objectives

1. Build a 3D map of the surface surrounding the pit to enable
 - a. characterization of terrain morphology and identification of surface features that may indicate collapsed or un-collapsed tunnels; and
 - b. planning for future missions.
2. Map geophysical signature of the terrain around a pit.
3. Build a high-resolution colored 3D model of the pit to enable
 - a. characterization of pit morphology;
 - b. determination of size distribution of blocks on the pit floor;
 - c. situational awareness for scientists evaluating spatially-distributed measurements; and
 - d. planning for future missions.
4. Determine if void space is accessible from the pit and, if so, characterize its minimum extent. (E.g. determine that the void extends as far as the lander's and rover's limited lines of sight.)
5. Map distribution of material types in pit walls and on pit floor.



7.1.2 Concept of Operations

Scout uses a lander and a surface reconnaissance rover. The lander uses precision navigation and landing technology to navigate precisely over the target pit. It flies low over the pit during descent, capturing detailed images and LIDAR data of the pit floor and surrounding surface terrain. Less detailed

data is captured for pit walls. The lander then detects hazards in its intended landing zone, picks a safe site, and touches down.

After landing, camera and LIDAR data are processed to construct a 3D model of the pit. This model may be constructed onboard the lander or data may be transmitted back to Earth for processing, depending on communication bandwidth and latency and available onboard processing. This model is used to plan operations for the surface reconnaissance rover.




Figure 32: Flyover and Surface Reconnaissance mission concept. Lander captures LIDAR and camera imagery of terrain during flyover. Rover then captures data of the same region, but from a different perspective. Rover is localized within lander imagery to improve the combined model.

The rover deploys from the lander, departs the landing site, and traverses to the target pit, collecting data as it travels. The rover carries, at minimum, a color camera that can be pointed to change tilt and an inertial measurement unit (IMU). If budget and mass allocation allow, other sensors could include LIDAR, NIRCams (Krishnan, et al. 2013), NIR/visible light spectrometer, gravimeter, ground-penetrating radar (GPR), and additional cameras with different lensing or wavelength ranges. The rover follows a planned view trajectory that specifies the locations, view angles, and viewing times to build a high quality model of the pit walls and visible regions of pit floor (Jones, Tabib and Whittaker n.d.) Scientists may specify certain targets of interest for material determination in both the lander and the rover models as they are incrementally constructed. The rover may carry a spectrometer or NIRCams to perform material analysis. If not, the rover’s camera can capture images of the target at a range of viewing and illumination angles to provide some information on material type. As the rover travels, it stops periodically to measure the gravity vector with either a gravimeter or its IMU. A GPR can also be used to measure geophysical data as the rover moves. Navigation is done by registration of rover data to the lander model. Like the lander model, construction of the combined rover and lander model could happen onboard either the rover or lander, or be done on Earth. The rover communicates by relay through the lander. Depending on the mission destination, the lander either communicates directly to Earth or through a relay satellite.

For the Moon, primary mission objectives could be accomplished in one lunar day or less. On Mars, it would take multiple sols. In an extended mission, the rover drives a lawnmower pattern to survey the terrain around the pit, building up a 3D model of the surface and capturing geophysical data.

7.1.3 Science and Human Exploration Contributions

This mission constructs a high-resolution colorized 3D model, including high-precision 3D measurements with defined scale, which can be used to characterize pit morphology and map the size distribution of rocks on the pit floor. Birds-eye views from the lander are ideal for LIDAR modeling of the pit floor. Gravimetry and GPR measurements can identify subsurface void extents. GPR can also identify subsurface layers. High-resolution 3D modeling of surrounding terrain may indicate subsurface

geometry through observation of features like surface depressions over lava tubes. Data captured on color and material type can be used to create maps of material-type distribution and identify pit wall layers. Determination of the existence of a cave is useful for both science and human exploration.

7.1.4 Potential Lunar Mission Target

The Lacus Mortis pit (44.962 N, 25.61 E), the destination for Astrobotic’s planned inaugural mission, is a good target for a Flyover and Surface Reconnaissance mission. The pit is relatively large, which is advantageous for flyover because navigation does not have to be quite as precise as for a smaller pit and because a larger pit diameter gives the lander better potential view angles for the pit walls. The Lacus Mortis pit may have an accessible cave entrance but existence of a cave has not been confirmed. Visible layers are present in the pit walls, making it an interesting target for scientific study. The presence of a gentle ramp also means that if a cave is confirmed, a crawler robot could be used for both pit access and subsurface exploration.



Figure 33: The Moon’s Lacus Mortis pit is a good target for a Flyover and Surface Reconnaissance mission.

7.2 WAYFARER: SURFACE RECONNAISSANCE

Mission Summary

A surface reconnaissance rover traverses from a landing site to model multiple pits. Similar to Scout, the rover traverses to and circumnavigates pits to enable reconstruction of 3D models. Long-range roving enables characterization of multiple pits in one mission.



Uses Lander, Surface Explorer, possibly Flyover Explorer

7.2.1 Mission Objectives

1. Build a 3D map of the surface surrounding the pits to enable
 - a. characterization of terrain morphology and identification of surface features that may indicate collapsed or un-collapsed tunnels; and
 - b. planning for future missions.
2. Map geophysical signature of the terrain around and between pits.
3. Build a high-resolution colorized 3D models of multiple pits to enable
 - a. characterization of pit morphology;
 - b. determination of size distribution of blocks on the pit floor;
 - c. situational awareness for scientists evaluating spatially-distributed measurements; and
 - d. planning for future missions.
4. Determine if void space is accessible from the pit and, if so, characterize its minimum extent. (E.g. determine that the void extends as far as the rover's limited line of sight.)
5. Map distribution of material types in pit walls and on pit floor.

7.2.2 Concept of Operations

Wayfarer uses a lander and a surface reconnaissance rover. The landing may follow the pattern of a Flyover and Surface Reconnaissance mission (i.e., lander flyover modeling of the first target pit) or it may simply land at an identified safe location a reasonable distance from the first target pit.

The rover departs the landing site and proceeds to the first target pit. It carries, at minimum, a color camera that can be pointed to change tilt and an inertial measurement unit (IMU). If budget and mass allocation allow, other sensors could include LIDAR, NIRCam, NIR/visible light spectrometer, gravimeter, ground-penetrating radar (GPR), and additional cameras with different lensing or wavelength ranges.

The rover follows a planned view trajectory that specifies the locations, view angles, and viewing times to build a high quality model of the pit walls and visible regions of pit floor (Jones, Tabib and Whittaker n.d.). The rover view trajectory is computed from a lander model for the first pit, if available. Otherwise, and for all subsequent pits, view trajectories are planned from the best available orbital data. Scientists may specify certain targets of interest for material determination in the rover model as it is incrementally constructed. The rover may carry a spectrometer or NIRCam to perform material analysis. If not, the rover's camera can capture images of the target at various viewing and illumination angles to provide some information on material type. As the rover travels, both along its view trajectory for a

given pit or on stretches between pits, it stops periodically to measure the gravity vector with a gravimeter or with its IMU. A GPR could measure geophysical data as the rover moves.

Navigation around pits uses registration to orbital data and Simultaneous Localization and Mapping (SLAM) (Thrun 2003) for navigation and modeling. As the rover travels, it builds a 3D model of its route using structure from motion, which enables matching to orbital data even when the orbital data pixel size is much larger than the rover. Depending on the mission destination, the rover either communicates directly to Earth or through a relay satellite.

Certainly for Mars and likely for the Moon, it would take multiple days to visit and model multiple pits. For the Moon, the rover would need the ability to survive the night either by entering a powered-off hibernation state or using radioisotope heating.

7.2.3 Science and Human Exploration Contributions

This mission constructs a high-resolution colorized 3D model that can be used to characterize pit morphology and map the size distribution of rocks on the pit floors. Birds-eye views from the lander are ideal for LIDAR modeling of the pit floor. Gravimetry and GPR measurements can identify subsurface void extents. GPR can also identify subsurface layers. High-resolution 3D modeling of surrounding terrain may indicate subsurface geometry through observation of features like surface depressions over lava tubes. Data captured on color and material type can be used to create maps of material-type distribution and identify pit wall layers. Determination of the existence of a cave is useful for both science and human exploration.

As scientists seek to build analytical models of pit formation, evidence from one pit is helpful but evidence from multiple pits is much better. Surface reconnaissance missions can get this data by visiting more than one pit in a single mission.

7.2.4 Potential Lunar Mission Target

A surface reconnaissance mission would be ideal for investigating impact melt pits. Impact melts often result in many pits located in the same region. Figure 35 shows the locations of pits in Tycho Crater. Figure 34 shows one of these pits. A potential surface reconnaissance mission would land near one of the clumps of nearby pits in Tycho crater, model each pit in that clump, then move on to other pits, attempting to cover as many as possible before end of mission.



Figure 34: One of the impact-melt pits in the Moon's Tycho Crater. [Image credit: NASA/GSFC/ASU]




Figure 35: DEM of the Moon's Tycho Crater with pits marked by black dots. [Image credit: Robert Wagner, ASU]

7.3 SPELUNKER: PIT ACCESS AND SUBSURFACE EXPLORATION

Mission Summary

Spelunker targets a pit that has been identified by previous surface or flyover reconnaissance to have void access (i.e. a pit classified as Some Cave Connected or better). Spelunker descends into the pit, enters the void, models the void, and directly samples the pit floor. This mission is suitable for any ramp classification: Gentle Entry Ramp, Steep Entry Ramp, or No Entry Ramp.



Uses Lander, Pit Descender, Pit Explorer, Subsurface Explorer, possibly Surface Explorer, and possibly the supporting roles of Descent Infrastructure Emplacer and Pit Infrastructure Link

7.3.1 Mission Objectives

1. Capture high-precision 3D measurements of a pit and accessible cave regions.
2. Build a colorized 3D model of one or more attached cave regions to enable
 - a. characterization of cave morphology;
 - b. determination of size distribution of blocks on the pit and cave floors; and
 - c. situational awareness for scientists evaluating spatially-distributed measurements.
3. Characterize the accessible cave extent.
4. Map distribution of material types on pit and cave walls and floors.
5. Characterize cave thermal and radiation environment.
6. Search for signs of life (Mars only).
7. Plan for human exploration missions.

7.3.2 Concept of Operations

The Spelunker mission includes a lander and a crawler. For a pit with no ramp, the mission also includes a surface rover and a Tyrolean robot. The lander uses precision landing technology to navigate to the target pit, to fly over at least one pit closely enough for modeling, and to land at a location determined using data from a prior Reconnaissance mission.

Pit with No Entry Ramp

After the lander touches down near a pit with no entry ramp, it launches a tether to the opposite side. A rover egresses to circumnavigate the pit and anchor the tether. A Tyrolean robot egresses from the lander and attaches to the tether. The Tyrolean robot traverses the pit along the horizontal tether while raising and lowering sensors vertically to scan the pit walls, floor, and voids, capturing high-resolution visual, LIDAR, thermal, and spectral data. This data is used to generate a high-resolution model, which then drives choice of deployment site for the crawler robot. The crawler is lowered into position by the Tyrolean robot. The crawler ventures into the cave, returning to the deployment point to re-charge from the Tyrolean-supported tether and to transmit data.

Pit with Ramp

After the lander touches down near a pit with an entry ramp, the crawler egresses to descend into the pit. As the crawler descends, it generates high-resolution visual, LIDAR, thermal, and spectral models of the pit walls and floor. The crawler then ventures into the cave, returning to the deployment point to recharge from solar power and to transmit data.




Figure 36: A Tyrolean robot lowers a payload into a pit to access a subsurface cave (left). A Krawler robot explores a subsurface cave (right).

7.3.3 Science and Human Exploration Contributions

Pit descents offer more complete and detailed views of pit surfaces – vertical faces, barren monolithic rock, significant banding, bouldering and regolith dusting – by providing additional perspectives not possible from orbital, flyover, and surface modeling. Robotic imaging and sensing of thermal transients indicate geologic conditions and may indicate suitability to support frozen volatiles. The descent to the pit floor enables a more complete model of the pit and void, including assessment of its suitability for further subsurface exploration and for human habitation.

Science sampling requirements determine what parts of the cave to investigate, what sensors are necessary, and how far a robot must travel inside a cave to gather useful data. On Earth, floors are of particular interest in lava tubes, but walls and ceilings are more interesting in other types of caves. Of particular interest are measurements beyond the “twilight zone,” which is the transition from the region that is illuminated for some period during the day into the region of constant darkness. Observations through and beyond this transition are likely to be indicative of the cave’s variation of volatile contents and geological features, which may be impacted by sunlight and temperature variations. Moving beyond the blocky debris caused by skylight formation into un-collapsed lava tube sections will likely reveal a gradient of regolith thickness. Measurements and imagery can indicate constituent materials, volatile contents, thermal gradients, radiation environment, void stability, and surface roughness, and can inform void formation models. Sensors may include laser spectrometry, neutron spectrometry, a core drill, and a microscopic imager. A radiation assessment detector and thermal sensors can measure the environment for comparison to surface environments and habitability assessments.


7.3.4 Potential Lunar Mission Target

If the Lacus Mortis pit were proven by Flyover and Surface Reconnaissance to have accessible cave extents, it would be a good target for a Spelunker mission.

7.4 ONE-SHOT SPELUNKER: FLYOVER AND SURFACE RECONNAISSANCE, PIT DESCENT, AND SUBSURFACE EXPLORATION

Mission Summary

One-Shot Spelunker performs flyover and surface reconnaissance, pit descent, and void exploration in one mission with no prior landing. This would be ideal in destinations where a cave has been confirmed from orbit



or for likely caves at planetary destinations far from Earth. This mission is suitable for any ramp classification: Gentle Entry Ramp, Steep Entry Ramp, or No Entry Ramp.

Uses Lander, Surface Explorer, Pit Descender, Pit Explorer, Subsurface Explorer, and possibly the supporting roles of Descent Infrastructure Emplacer and Pit Infrastructure Link

7.4.1 Mission Objectives

1. Capture high-precision 3D measurements of a pit, the surrounding terrain, and any accessible cave regions.
2. Build a high quality 3D map of the surface surround the pit to enable
 - a. characterization of terrain morphology and identification of surface features that may indicate collapsed or un-collapsed tunnels; and
 - b. planning for future missions.
3. Map geophysical signature of the terrain around a pit.
4. Determine if void space is accessible from the pit, and if so, characterize the accessible extent.
5. Build a colorized 3D model of any accessed cave regions to enable
 - a. characterization of cave morphology;
 - b. determination of size distribution of blocks on the pit and cave floors;
 - c. situational awareness for scientists evaluating spatially-distributed measurements; and
 - d. planning for human exploration missions.
6. Map distribution of material types on pit and cave walls and floors.
7. Characterize cave thermal and radiation environment.
8. Search for signs of life (Mars only).
9. Plan for human exploration missions.

7.4.2 Concept of Operations:

One-Shot Spelunker includes a lander, a surface reconnaissance rover, a crawler and, for pits with no ramp, a Tyrolean robot. The lander uses precision navigation and landing technology to navigate precisely over the target pit. It flies low over the pit during descent, capturing detailed images and LIDAR data of the pit floor and surrounding surface terrain. Less detailed data is captured for pit walls. The lander then detects hazards in its intended landing zone, picks a safe site, and touches down.

After landing, camera and LIDAR data are processed to construct a 3D model of the pit. This model may be constructed onboard the lander or data may be transmitted back to Earth for processing depending on communication bandwidth and latency.

Pit with No Entry Ramp

After the lander touches down near a pit with no entry ramp, it launches a tether to the opposite side. A rover egresses to circumnavigate the pit and anchor the tether. A Tyrolean robot egresses from the lander and attaches to the tether. The Tyrolean robot traverses the pit along the horizontal tether while raising and lowering sensors vertically to scan the pit walls, floor, and voids, capturing high-resolution visual, LIDAR, thermal, and spectral data. This data is used to generate a high-resolution model, which then drives choice of deployment site for the crawler robot. The crawler is lowered into position by the Tyrolean robot. The crawler ventures into the cave, returning to the deployment point to re-charge from the Tyrolean-supported tether and to transmit data.

Pit with Ramp

After the lander touches down near a pit with an entry ramp, the crawler egresses to descend into the pit. As the crawler descends, it generates high-resolution visual, LIDAR, thermal, and spectral models of the pit walls and floor. The crawler then ventures into the cave, returning to the deployment point to re-charge from solar power and to transmit data.

7.4.3 Science and Human Exploration Contributions

The One-Shot Spelunker mission combines the science and human exploration contributions of the Scout, Wayfarer, and Spelunker missions.

7.4.4 Potential Lunar Mission Target

The Mare Tranquillitatis pit has been shown to have at least some accessible subsurface void extent. Since the void is already confirmed (Robinson, et al. 2012), completing flyover and surface modeling along with pit descent and subsurface exploration can save launch costs. Since this steep-walled pit has no ramp, a Tyrolean robot would be needed for this mission.




Figure 37: Mare Tranquillitatis pit [Image credit: NASA/GSFC/ASU]

8 MODELING METHODS FOR PITS AND CAVES

This section discusses the sensors that were considered in this work for modeling pits and caves (Section 8.1), gives a brief description of modeling methods in the “toolbox” for planetary pit and cave modeling that were used in this work (Section 8.2), and presents new *Lumenenhancement* methods for sensor fusion and exploiting illumination in modeling (Section 8.3). The algorithms behind the *Lumenenhancement* methods are presented in Section 8.3.3.

8.1 SENSORS

8.1.1 Camera

Cameras are simple and versatile sensors that will be carried by any robotic explorer of planetary pits and caves. Because humans are accustomed to perceiving the world through visual information, camera images are a great way to give mission controllers, scientists, and the public an understanding of the environment a robot encounters.

A digital camera converts light into electronic signals using a sensor chip, which is generally a Charge Coupled Device (CCD) or Complementary Metal–Oxide–Semiconductor (CMOS) sensor. The two sensor types vary in the way that they convert charge to voltage to read the signal for each pixel. Each has advantages and disadvantages, but both have been successfully used in space applications. Some CMOS sensors have a rolling shutter, in which rows of pixels are exposed to light sequentially, as opposed to a global shutter, in which all pixels are exposed to light at the same time. If cameras or objects being imaged move during image capture with a rolling shutter, that motion will cause distortions in the image (Litwiller 2001). Thus robots that move during image capture, (this will be the case with landers, is likely to be the case with Tyrolean robots, and may be the case with surface reconnaissance rovers and crawlers), should use cameras with global shutters. “Blooming” is another consideration – charge-buildup on one pixel can overflow onto nearby pixels. This is especially problematic if some areas of the image are much brighter than others, which would often be the case when imaging sunlit planetary pits with areas of dark shadow. Most CMOS sensors do not have this problem. Steps can be taken to mitigate it in CCD sensors, but not all cameras include that mitigation (Litwiller 2001). For both CCD and CMOS, the quality of the sensor is important and can determine the best attainable signal-to-noise ratio.

Camera parameters that can be adjusted on the fly include ISO, aperture, and exposure time. ISO determines how sensitive the imaging sensor is to light. High ISO values (e.g. 1600, 3200, 6400) make the sensor very sensitive to light, so it is easier to capture images in the dark. However high ISO also increases the amount of noise in the image. Low ISO values (100, 200) make the sensor less sensitive to light and reduce noise in the image. The aperture setting determines the size of the hole that lets light onto the camera sensor, which in turn determines the depth of field. Larger depth of field (smaller aperture) means that objects at a wide range of distances will be in focus, while shallower depth of field (larger aperture) means that only objects within a narrow distance range will be in focus. Small apertures are better when the field of view contains objects at a wide range of distances, but large apertures, because they let in more light, may be better for imaging in dark environments. Aperture settings are given in f-stops, with the f-stop number being inversely proportional to the aperture size. The exposure time determines how long the camera sensor will be exposed to light. Images in bright sunlight should be set to a short exposure time, while images in low light should have a long exposure

time. Long exposure time can result in blurred images if the robot moves while imaging. Exposure time can be traded off with ISO and aperture when imaging in low light, depending on the relative importance of low-noise imaging, depth of field, and imaging speed.

Dynamic range is also an important consideration. The dynamic range of an imaging sensor determines the range of brightness values it can capture. For the cameras on the LRO, for example, the imaging sensor has a much higher dynamic range, using 16 bits to store each pixel of data, than consumer-grade cameras, which have only 8 bits (per color) to store each pixel of data. The dynamic range of an image can also be made better than the camera's original dynamic range using a procedure called High Dynamic Range (HDR) imaging. In HDR, multiple images are taken at different exposure times (exposure bracketing): short exposures capture detail in bright areas of the frame, and long exposures capture detail in the dark areas of the frame. This set of images can then be combined into one in which more information is available at each pixel. Because HDR works best when the robot remains still while capturing all images in the HDR bracket, it is most applicable to surface reconnaissance robots and crawler robots, which are able to completely stop moving while imaging.

In choosing camera lenses for pit and cave exploration robots, quality is important, and so is field of view (FOV). For a given camera sensor resolution, a narrower FOV results in finer ground resolution on the imaged terrain, while a wider FOV covers a larger terrain area. For landers, if a camera used in flyover modeling is fixed, then the FOV should be calculated based on the image footprint on the pit terrain given the expected range of possible flyover trajectories. If such a camera is actively pointed, a narrower FOV could be used. For surface reconnaissance rovers, FOV should be chosen with careful consideration of the tradeoff between the number of images required to cover the pit (narrower FOV cameras require more images) and the desired ground resolution. Needing a large number of images to cover the pit may also create problems in stitching the images together, especially if the rover's estimate of its position is not good. At the extreme of wide-FOV lenses is the fisheye, which produces a circular image. Fisheye lenses are good for situational awareness, since they cover a wide footprint with a single image, but they have high distortion (see Figure 30 in Section 6.2.5.1) for an example image from a camera with a fisheye lens). Tyrolean robots and crawler robots face similar FOV tradeoffs. Situational awareness may be a much higher priority for crawler robots, since they travel over very rough terrain.

Cameras are generally used as passive sensors. They use existing light (which in planetary environments comes from the sun) in order to take images. For surface reconnaissance rovers, there may be enough light reflected off illuminated pit walls to effectively capture images in dark shadows, given a long enough exposure time. For subsurface exploration, however, a robot must bring its own light. The shape and placement of light sources relative to cameras is important. Camera centered point sources are ideal for detecting surface features. However, in practice, many materials and the atmosphere are backscattering. A small separation between a point source and the camera greatly increases backscattered light, which lowers contrast, rather than "good" light from the material surface. Wide separation in phase angle (observer to source) solves this but induces cast shadows that foil computer vision. Area sources can also mitigate this at the cost of contrast for human viewing.

The fact that the robot must bring its own light can also be an advantage. As discussed extensively by Wong in his thesis work (U. Wong, Lumenenhancement: Exploiting Appearance for Planetary Modeling 2012), if the robot brings the light, the characteristics of the light are known, and this knowledge can be

exploited. Structured light techniques – projecting a known pattern of light onto the terrain – can also help perceive 3D shape, as used by the Sojourner Mars rover (Schenker, et al. 1997).

Scientists often use imaging with a broader range of light wavelengths than that visible to the human eye, (called multispectral or hyperspectral imaging), in order to determine material types in planetary environments. Comparing camera responses in the ultraviolet, visible, and near infrared (NIR) ranges provides much more information than just examining a color image. NIRCcam, a prototype instrument developed at Arizona State University, pairs a digital camera with a set of filters in the visible and NIR ranges (Krishnan, et al. 2013). By taking an image with each filter, NIRCcam yields a coarse spectrometer reading for each pixel. This is cheaper than having separate cameras with different wavelength sensitivities, and it provides information on spatial distribution of materials more easily than a spectrometer. The NIRCcam requires a known, broad-spectrum light source. When imaging in sunlight, the sun can serve as this source, but for imaging in a cave, a portable broad-spectrum light source must be carried. A researcher on this NIAC project team experimented with the NIRCcam in a dark environment using different light sources. Light sources like LEDs (Light-Emitting Diodes) are considered “efficient” because the light they produce is primarily in the range of wavelengths visible to humans – there is little energy “wasted” producing other wavelengths. For the NIRCcam, however, these other wavelengths are necessary. The researcher observed that because each filter only passes a narrow range of wavelengths, a lot of light was needed to get a good NIRCcam measurement, so the light source consumed a lot of power. Utilizing a series of synchronized LEDs, which are tuned to the output of each filter and that emit the same spectral response as the camera, could dramatically reduce power for robots that explore subsurface caves.

Thermal cameras, which operate in the infrared range, could also be useful in pit and cave exploration. Information about material types and densities can be determined by watching how quickly an object’s temperature changes as incident solar illumination changes.




Figure 38: A sequence of NIRCcam images in a dark room using a halogen incandescent light source. Items in the image include a white Spectralon target (which has a calibrated spectral response), an aluminum plate, coins, leaves, and small rocks. The filter wavelengths in nanometers are 450, 523, 671, and 750 (top row) and 850, 920 and 1050 (bottom row)



Figure 39: The same scene illuminated by an LED array light source. The light source illuminates the scene effectively in the visible range (400-700 nm), but as the filters progress farther into the NIR range, the images go black.

Most consumer cameras and most machine vision cameras are frame cameras. They capture a 2D array of pixels in every image. Cameras used in orbital reconnaissance missions (such as the Lunar Reconnaissance Orbiter Camera, or LROC) are sometimes pushbroom or pushframe cameras. A pushbroom camera captures one line of pixels at a time. An image is built up by capturing successive rows of pixels as the satellite moves. LRO's Narrow Angle Cameras (NACs) are pushbroom cameras. A pushframe camera is similar, but instead of a single line of pixels, they capture a few lines at once. LRO's Wide Angle Camera (WAC) is a pushframe camera that captures 14 rows of 1024 pixels at a time (Robinson, et al. 2010). Because images in pushbroom and pushframe cameras are formed so differently than a for a frame camera, machine vision algorithms must often be specially adapted to work with them.

8.1.2 LIDAR

LIDAR sensing works by sending out light at a specific wavelength and measuring the reflection. Some LIDAR sensors work using time of flight: they measure the time from sending the pulse of light to receiving the reflection. From the time of flight, they compute the distance to the object causing the reflection. Other LIDAR sensors emit a continuous wave of frequency-modulated light, measuring the phase shift of the reflected wave, and using that to compute distance (Wong, Morris, et al. 2011). The Faro scanner used to collect ground truth data in this work is a phase-shift LIDAR.

A LIDAR using a single laser beam can effectively measure distance in one direction. This is useful as a laser altimeter, but not very good for modeling larger areas. With the addition of a spinning mirror, a LIDAR can capture a line of points instead of one individual point. The entire LIDAR with its spinning mirror can then be rotated on an axis perpendicular to the mirror rotation axis to collect a 3D point cloud, as described in Section 9.4. Alternately, some sensors tilt the mirror in two dimensions. Flash LIDAR uses a different approach, sending out a flash instead of a narrow beam of light and measuring return values on a grid of pixels to produce an image of depth values.

LIDAR sensors often use NIR light. An important consideration when choosing a LIDAR for use in a mission is which wavelength of light is being used and how that wavelength is reflected by the terrain to be modeled. For example, lunar mare basalt tends to have low reflectance in the NIR range (Heiken,

Vaniman and French n.d.). The effect of using a LIDAR with a NIR wavelength on lunar mare terrain would be a decrease in the maximum sensor range compared to operation on other materials. Thus, the manufacturer's quoted maximum range may have to be reduced for use in sensor selection and mission operations planning.

8.1.3 Spectrometer

A spectrometer profiles light over a range of wavelengths. If a light source with a well-characterized spectrum (the sun as observed from the Moon, for example) shines on a material, then a spectrometer can be used to measure the profile of the reflected light. The spectrometer determines the type of material from which wavelengths are most reflected and which are most absorbed. This is often a single-point measurement, though sensors like the NIRCcam can obtain some spectrometry information over an entire image. For a surface reconnaissance rover, a spectrometer for the UV, visible, and/or NIR ranges can be a passive instrument, using the sun as a light source. For subsurface exploration, however, a well-calibrated, broad-spectrum light source would be needed. Laser-Induced Breakdown Spectroscopy (LIBS) uses laser pulses to break down and excite the target material and then measures the spectral response. LIBS could be a good choice for subsurface exploration, since it is an active sensing method. Neutron Spectrometers count neutrons radiated off a surface, as opposed to UV, visible, or NIR light. They are especially useful when looking for signs of water.

8.1.4 Ground-Penetrating Radar

Radar uses reflected radio waves to measure distance. Ground-penetrating radar (GPR) uses a longer wavelength than radar used to detect obstacles, and it is able to detect boundaries between layers of different material underground. GPR could be carried by a lander or a surface reconnaissance rover and used to survey the ground above the pit and look for underground voids.

8.1.5 Gravimeter

A gravimeter measures gravity. If gravity is measured very precisely, the difference between solid rock (dense) and void space (not dense at all) might be detectable. Gravimeters have been used in orbit around the Moon and by aircraft and ground vehicles on Earth. An orbiter can build a 2D map of gravimeter readings over time, as can a surface rover. A lander is not a good platform for most gravimetry experiments, because a normal landing happens fairly quickly and an aerial survey would be fuel-prohibitive.

8.1.6 Inertial Measurement Unit (IMU)

An Inertial Measurement Unit (IMU) is a commonly used sensor for spacecraft and robots consisting of three accelerometers (to measure linear acceleration) and three gyroscopes (to measure rotational velocity). IMUs are often used in navigation, but because they measure 3-axis acceleration, they can also be used to collect gravity data. If a robot is not moving (i.e. a surface rover sitting on the ground), the acceleration vector measured by the IMU will be due to gravity. This acceleration vector is measured with respect to the robot's local frame, so any robot tilt due to slopes or rocks will affect the vector's direction. If the robot also builds a 3D model of the local terrain and monitors any active or passive degrees of freedom in its suspension, a map of gravity vector direction and magnitude relative to local terrain could be constructed. The degree to which this model would be useful for science is dependent on sensor sensitivity.

8.1.7 Other Sensors

Pit and cave exploration robots can also carry other sensors as desired for scientific investigations, so long as the mass and power requirements are not too high. A sun sensor, useful for navigation, could also be used to determine the angle of incident solar illumination on an imaged surface. Fluorescence imagers could be used in the search for life. Sensors to determine atmospheric composition inside caves could be useful on Mars.

8.2 COMMON MODELING METHODS

Stereovision is a method of determining 3D structure by triangulation using features identified in two or more images (see Figure 40). A common approach in robotics is to do stereo using two parallel cameras rigidly mounted on a robot with a fixed baseline (distance between the cameras). This is the configuration shown in Figure 40. The difference in position of a 3D point's projection between two images is the disparity. If the baseline is known, disparity values can be used to compute the depth of 3D points.

The closer a point is to the cameras, the larger the disparity will be. For more distant points, the disparity shrinks. As the disparity approaches the width of one pixel, the ability of stereo to resolve depth decreases. Increasing the baseline can increase the maximum range of stereo, but at some point it is no longer realistic for the two cameras to be rigidly mounted on one robot. For example, a 1-meter tall robot with a 5-meter baseline would not work. The concept of stereo still works in configurations other than parallel, fixed-baseline. One option is to take a picture, move the robot, and take another picture. Arbitrarily large baselines can be achieved in this way. The vision algorithm must apply a 3D transformation incorporating translation and rotation instead of just a 1D translation to get the baseline.

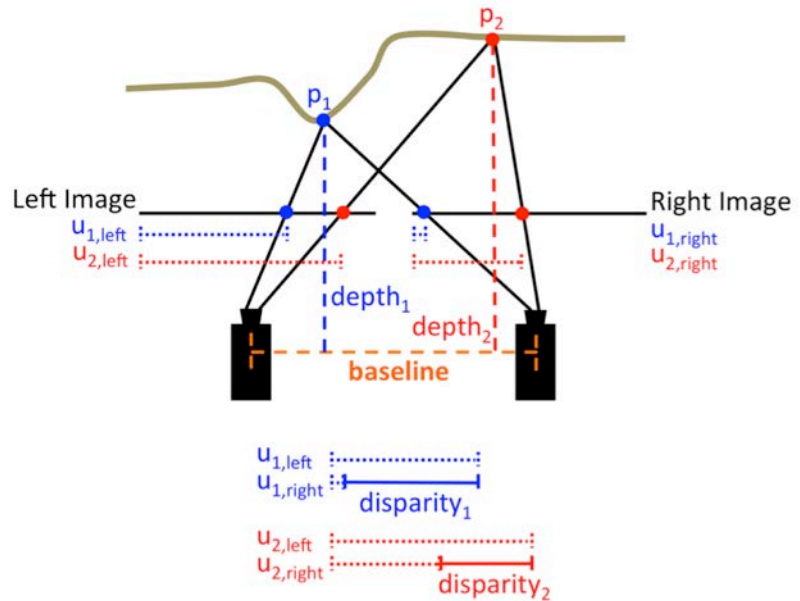


Figure 40: Illustration of the concept of stereo vision. A 3D scene is viewed by two cameras. The projection of scene points into the camera images is identified. The difference between the positions of a point's projections on the two images is the disparity. Knowing the baseline between the two cameras, the disparity can be used to compute a point's depth. Points closer to the cameras (p_1 in this figure) will have larger disparity values than points farther from the cameras.

In addition to the maximum range of stereo, the minimum range must also be considered. This is determined by the overlap between the frustums, or viewing pyramids, of the two cameras (see Figure 41). For large baselines, the cameras are often tilted toward each other to get better overlap. When that happens, consideration must also be given to the relative view angle between the two cameras. Too large an angle may prevent features from being effectively matched between images due to change in scene perspective and tilted planes of sharp focus between the two cameras. Stereo can also be implemented with more than two cameras. This is called multiview stereo.




Figure 41: Illustration of the overlap between a pair of stereo cameras. Objects in the yellow region can be perceived in 3D using stereo. Objects closer than depth d can never be perceived, but since the aim is generally to get a 3D reconstruction of more than one point, the effective minimum depth is larger.

Structure from motion is a stereo-based technique involving two or more cameras, where the 3D transformations between the camera views are not necessarily known beforehand. Structure from motion can compute 3D structure, but only up to scale if these 3D transformations are not known. It is not possible to distinguish between a camera moving a small distance around a small object or moving a large distance around a large object. When applying structure from motion techniques to robotic pit exploration, scale can be determined either from estimates of the robot's pose as it captured images, or from previously captured data with known scale. For example, if the diameter of a pit was measured in an orbital image, this can be used to scale a structure-from-motion model built by a surface reconnaissance rover.

The structure-from-motion pipeline for creating high-quality models involves first doing bundle adjustment to get a sparse reconstruction, and then doing multiview stereo to get a dense reconstruction. Bundle adjustment optimizes both the 3D structure of a set of points and the camera poses from which these points are visible (and perhaps the camera calibration) (Triggs, et al. 2000). Bundle adjustment is commonly used in orbital imaging to align images taken by a satellite on different orbits, or taken by different satellites. Traditionally, the points used for bundle adjustment were manually identified in each image, but the process also works with automatically detected and matched feature points. After bundle adjustment, the newly computed camera poses can be used for dense multiview stereo. This can produce a 3D position for each pixel in an image (in regions for which two or more cameras overlap), though the 3D positions are usually more accurate if they are computed over more than one pixel (i.e. a 2 pixel x 2 pixel, 3 pixel x 3 pixel, or larger patch). This is because the image may not contain sufficient information on the individual pixel scale to distinguish the correspondence between multiple images that finely.

Given two or more point clouds obtained from LIDAR, stereo, structure from motion, or some other means, it is useful to be able to align the two, either to combine multiple models or to compare models built by different methods. One good method is the Iterative Closest Point (ICP) algorithm (Zhang 1994). ICP uses one of the point clouds as the reference. The other point cloud, called the source or data point cloud, is rotated and translated to minimize the distance between the two. ICP works best when the two point clouds are of comparable density or when the reference point cloud is denser. It also requires a

large overlap between the two point clouds. If the source point cloud has points that are known to be outside of the range of the reference point cloud, those should be trimmed out before performing ICP. The 3D transformation computed by ICP can then be applied to all points in the source point cloud, including those clipped out for processing.

8.3 SENSOR FUSION

Cameras can produce very high-resolution intensity and color information: one or more megapixels per image. LIDAR measures 3D shape very precisely, but its resolution is often coarse. Fusing dense camera images with sparse LIDAR can exploit the advantages of both sensors. The key is that appearance in a 2D image provides clues to 3D shape. The simplest example is the observation that a contiguous patch of pixels of the same color and intensity in the image likely corresponds to a flat surface in the world, while a sharp change in color and intensity may correspond to a discontinuity in 3D space, especially if LIDAR points on one side of it are at a different range than those on the other. Lumenhancement, a family of camera and LIDAR methods developed by Wong for modeling in planetary environments, uses the additional insight that knowledge about the interaction of light with rocky surfaces can further improve the fusion of camera and LIDAR data (U. Wong, Lumenhancement: Exploiting Appearance for Planetary Modeling 2012). The following subsections provide an in-depth technical description of the Lumenhancement sensor fusion method. Section 8.3.1 describes the early development of sensor fusion, Section 8.3.2 describes Lumenhancement, and Section 8.3.3 details the Lumenhancement algorithms.




Figure 42: Sparse LIDAR range readings (left) and high resolution intensity images (center) can be fused to create super-resolution models (right) with a Markov Random Field.

8.3.1 Sensor Fusion Background

The fusion of range and imaging sensors to improve 3D model quality has been studied in (Diebel and Thrun 2005), (Torres-Mendez and Dudek 2008), (Gould, Baumstarck and Quigley 2008). A general model for fusing raw LIDAR and image data into super-resolution range images using a Markov Random Field (MRF) was explored in Diebel and Thrun’s seminal paper (Diebel and Thrun 2005). MRFs are undirected graphs that represent dependencies between random variables and have been used extensively in computer vision for noise removal, feature matching, segmentation and inpainting (Li

2009). The popularity of the MRF stems from the ability to model complex processes using only a specification of local interactions, relevance to the regular grid nature of CCD images, and the maximum *a posteriori* (MAP) solution requiring only direct convex optimization in many cases.

Diebel and Thrun surmised that higher resolution intensity (color) data could be used to texture range images and increase the range accuracy of interpolated points. The results in a uniformly and sufficiently illuminated regular office environment are quite compelling. Cameras are able to turn LIDAR scans into dense range images with very low computational overhead. However, the assumption that an image provides relative range information, even locally, is tenuous in unstructured environments. Generating 3D geometry from a general 2D projection is an ill-posed problem. The ability of Diebel's method to smooth point clouds using areas of flat image information was convincingly shown, but the converse of enhancing a point cloud using image texture was not.

8.3.2 Lumenenhancement

The work of Wong (U. Wong, Camera and LIDAR Fusion for Mapping of Actively Illuminated Subterranean Voids 2009) (U. Wong, Lumenenhancement: Exploiting Appearance for Planetary Modeling 2012) demonstrated *Lumenenhancement*, a family of related camera/LIDAR techniques for mapping planetary spaces. In particular, a Shape from Shading approach was coupled with Diebel's MRF to show that contributions of local shape estimated from imagery could generate an order of magnitude better quality of reconstruction. A key requisite of this approach was that Shape from Shading techniques were particularly effective in planetary spaces with their barren surfaces, rocky materials, and direct solar illumination (little or no atmospheric scattering). Results from Wong have shown a 40x increase in measurement density (optics dependent) and up to 40% increase in range accuracy for some lunar analogs.

8.3.2.1 Subsurface Modeling with Lumenenhancement

Robots must carry active illumination for imaging in dark, subsurface environments; Wong demonstrated that the form and distribution of such illumination can and should be designed for data enhancement beyond simple photography. Such methodology is directly applicable to barren environments like sub-planetary caves and lava tubes. By coupling point-source illumination with the assumption of diffuse surface reflectance in these environments, estimating geometry at every image pixel becomes a shape-from-shading (SFS) problem.

While the SFS framework can be solved with traditional statistical optimization methods, these are typically numerically unstable and inaccurate in the field. Variation of materials and albedo, complexity of intersecting geometry, and sensor noise makes estimation of shape from a single image severely ill-constrained. Sparse LIDAR data significantly simplifies the problem and bounds global errors, as LIDAR provides a way to directly validate geometry. This interaction of LIDAR and camera data is the fundamental idea of super-resolution fusion. Actively illuminated imagery is used to generate locally-consistent surface detail, which is then "textured" onto blocky, but globally-consistent 3D range models.



Figure 43: A mapping robot explores a terrestrial underground space using active illumination (left). An immersive virtual model of the environment is created using camera and LIDAR fusion techniques (right).

8.3.3 Lumenhancement Algorithms

This section elaborates on some of the algorithmic details, based on Wong’s work, which were used to create super-resolution models for this NIAC project.

8.3.3.1 Fusion in the Markov Random Field Framework

A range image is used as the common representation for fusion. The 3D range cloud data is registered to the pinhole of the camera, forming a range map (R) via projection of distances onto the $n \times m$ image plane at equivalent resolution. Many pixels in the range map will not contain range measurements; these holes are filled from nearby data through bilinear or nearest neighbor interpolation. The color image data can be then converted to intensity values or used as a raw RGB vector (I). A lattice MRF is formed where there is a single range and intensity measurement associated with each node (Figure 44). This is similar to the MRF fusion method documented in (Diebel and Thrun 2005); however, the image gradients are instead numerically integrated in this framework.




Figure 44: Markov Random Field Graphical Model. Green nodes (I) represent the image pixel data, brown nodes (x) represent the hidden true range value to be estimated, aqua nodes (R) represent the sparse range data and the blue node represents the interpolation uncertainty estimate. There is 1 pixel value for every hidden node (x), but there may be many nodes without a corresponding range value (R).

The range map potential (0.1) promotes agreement between the estimated variables and the interpolated range data.

$$\Psi = w_1 \sum_{i \in \mathcal{L}} (R_i - x_i) \quad (0.1)$$

The smoothness prior (0.2) regularizes large changes in the range estimate and like the image potential (0.3) connects potential transfer from a node to its neighbors.

$$\Omega = \beta \sum_{i \in \mathcal{L}} \sum_{j \in \mathcal{N}(i)} (x_j - x_i)^2 \quad (0.2)$$

$$\Phi = \alpha \sum_{i \in \mathcal{L}} \sum_{j \in \mathcal{N}(i)} (x_j + \nabla I_{ij} - x_i)^2 \quad (0.3)$$

where relative weights are given by:

$$\begin{aligned} \alpha &= w_2 \exp(-c \cdot \sigma) \\ \beta &= w_3 (1 - \exp(-c \cdot \sigma)) \end{aligned} \quad (0.4)$$

The image gradient is a reasonable predictor of depth change across neighboring pixels. However, integrating the gradient to produce depths over a large locality is prone to drastic shape distortions. The range estimate can be used to regularize numerical integration of the intensity gradient. Moreover, it can ensure that nodes with true range readings are never changed. The weights α and β are relatively scaled by an interpolation distance uncertainty (σ) for some weights w_1 and w_2 (0.4). σ can be generated from the range image during inpainting by using the Matlab command BWDIST, for example. The potential function corresponds to a Gibbs distribution of the form:

$$p(x | R, I, \sigma) = \frac{1}{Z} \exp\left(-\frac{1}{2}(\Psi + \Omega + \Phi)\right) \quad (0.5)$$

$$x_{mle} = \arg \min_x f(\Psi + \Phi + \Omega) \quad (0.6)$$

Solving for the MAP of the distribution requires running a gradient descent algorithm on the target variables x in (0.6), where Z is the partition function (Diebel and Thrun 2005).

8.3.3.2 Structure from Shading

The image gradient ∇I_{ij} in (0.3) can apply to either raw pixel data or better estimates of depth from the camera. As scene geometry cannot be ascertained from a single image without assumptions, often no better estimate exists. Definite reconstruction requires knowledge of image formation parameters like light field, surface reflectance (BRDF) and albedos. However, if assumptions like those commonly made in Shape-from-Shading are valid, as in the planetary domain, the amount of certainty is greatly increased. Robots in naturally dark environments like planetary caves can be fitted to carry small area light sources for photography which produce simple light fields. This was demonstrated in the program for underground Krawler exploration and mapping.

The MRF image observation (I) is estimated using Shape-from-Shading given the above assumptions. A lightness-based direct normal estimation method which uses range information is given below, but other techniques exist. This method factors range information to allow varying albedos and trades accuracy for feature preservation. The effect of the light source's irradiance fall-off is first removed from the raw image data (E_0). The following irradiance correction model for small area sources is assumed (0.7):

$$E_{unbiased} = \gamma(E_0) \cdot R^n \quad (0.7)$$

The radiometric function (γ) maps pixel values to irradiance, (R) is the interpolated depth estimate and (n) is the irradiance fall-off factor. For ideal point sources $n = 2.0$, while $n < 2.0$ for near-field area sources. The experimental setup described below exhibits an empirical decay of $n = 1.27$. The corrected image ($E_{unbiased}$) is devoid of a near-field illumination intensity bias from the use of an area source. Moreover, the compensation of intensity enables smooth color alignment when stitching several scans together.

Converting RGB color into a single intensity value provides compactness and symmetry, and also minimizes chromaticity effects. Color space transformations such as CieLAB or YCbCr are often used to heuristically isolate the lightness component of an image, discarding chromaticity and albedo. The SUV transformation (Mallick, et al. 2005) describes a class of physics-based specular-invariant color spaces produced by rotating the RGB space such that a single channel is aligned with the illuminant color vectors. This method has produced excellent results with single-source images and enables many Lambertian algorithms to handle a large set of environments with specularities. The specular invariant image, as defined in equation (0.8)-(0.9), is used in experimentation:

$$[s, u, v]^T = R_r(\theta) \cdot [E_{unbiased}^{(r)}, E_{unbiased}^{(g)}, E_{unbiased}^{(b)}]^T \quad (0.8)$$

$$E_{inv} = \sqrt{u^2 + v^2} \quad (0.9)$$

$R_r(\theta)$ is defined as a (3×3) rotation matrix that aligns the red channel of an $\{r, g, b\}$ triple with the source color. The magnitude of the $\{u, v\}$ components is taken to be the diffuse image.

An albedo map is subsequently generated from the diffuse image using Blake's method for lightness computation (Worthington 2005). Perceived intensity is a multiplicative relationship between surface slant angle and reflectance². The log image separates these components into additive terms. Scene albedos can be recovered from the gradient of the log diffuse image by thresholding to remove small changes and integrating. It is noted that the problem can be recast as finding the log albedo map (δ) that minimizes the following equation:

$$\arg \min_{\delta} \left| \frac{\partial}{\partial x} \delta - T_{\sigma} \left(\frac{\partial}{\partial x} \log E_{inv} \right) \right|^2 + \left| \frac{\partial}{\partial y} \delta - T_{\sigma} \left(\frac{\partial}{\partial y} \log E_{inv} \right) \right|^2 \quad (0.10)$$

² Recall the Lambertian BRDF: $E = \frac{\rho}{\pi} (\hat{n} \cdot \hat{l}) = \frac{\rho}{\pi} |n| |l| \cos \theta$

where (T_σ) is the threshold function. Exponentiating (δ) with the proper constant of integration produces the albedo values (0.11). The constant can be estimated from the range data to minimize depth discrepancy in the reconstruction.

$$\rho_{est} = \exp(\delta + c) \quad (0.11)$$

Lastly, surface normal approximations for every pixel are obtained by solving the Lambertian reflectance model:

$$E_{inv} = \rho |n| |l| \cos(\theta_{nl}) \quad (0.12)$$

$$\theta_{nl} = \arccos\left(\frac{E_{inv}}{\rho_{est}}\right) \quad (0.13)$$

The polar estimates (θ_{nl}) are combined with azimuth estimates (ϕ) from the range image. Range data is taken to be a reasonable indicator of the gradient direction while image intensities modulate the gradient magnitude. An integrable surface is constructed from these normals using the method of (Frankot and Chellappa 1988). The surface reconstruction is passed into the MRF as a second range image. An accurate surface is neither required nor preferred from this method. Instead, preservation of high frequency detail is preferred, while global consistency is enforced by the decoupled MRF pass.

8.3.3.3 Algorithm Summary

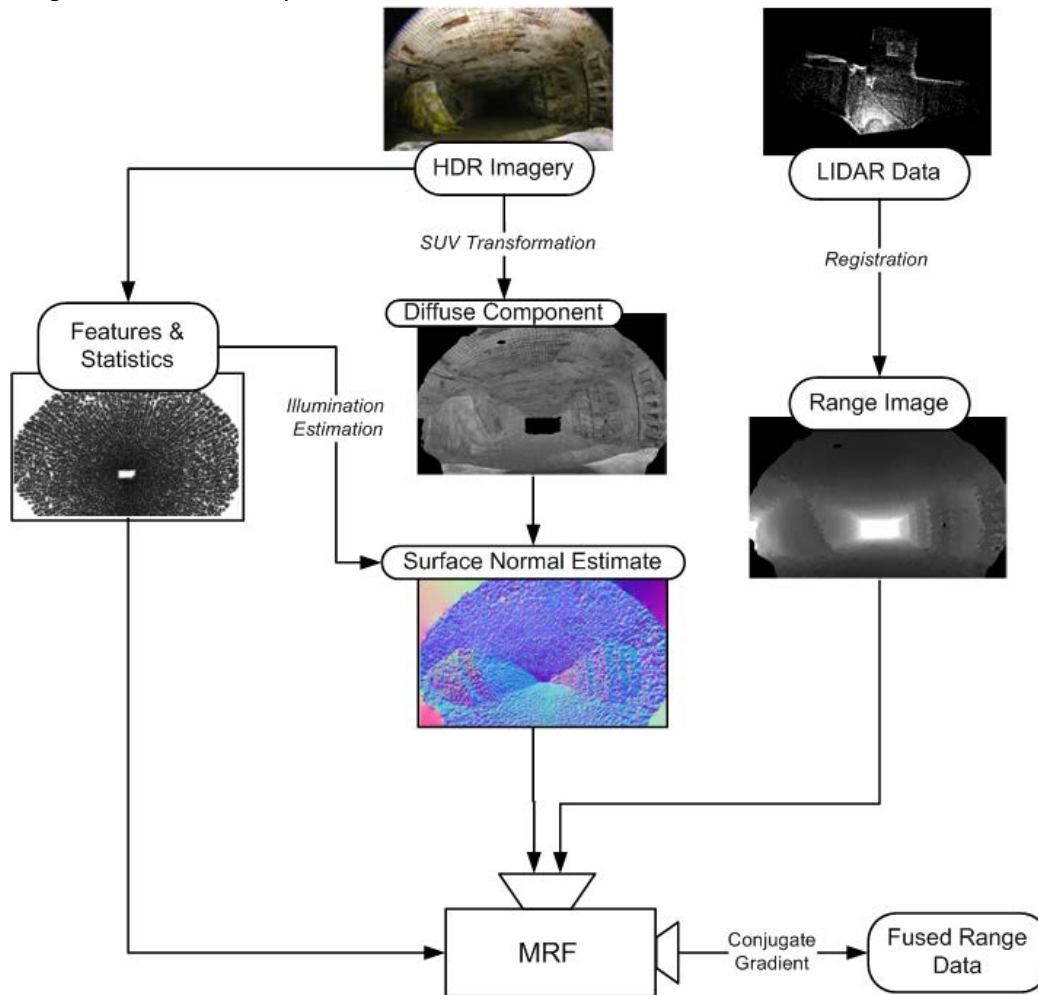


Figure 45: MRF Super Resolution Process. (1) Raw LIDAR point data is converted to a range image from the camera perspective. (2) Specularities are removed from the color HDR imagery to produce a diffuse image. (3) Surface normals are estimated from the diffuse image using shape from shading. (4) The surface normals and the range image are fused in the MRF framework. Image reproduced from (U. Wong, Lumenenhancement: Exploiting Appearance for Planetary Modeling 2012).

A flow chart overview of the technique presented is shown in Figure 45. Raw LIDAR data is first projected into the space of the image, and resampled (interpolated) to form a co-registered range image at the resolution of the color image. Then, HDR color images taken under controlled illumination are transformed into a purely diffuse intensity image using the SUV transformation and knowledge of the spectrum of the light source. The diffuse image along with image features like saturation, illumination and albedo estimates are utilized in a shape-from-shading approach to generate surface normal estimates at every pixel. An MRF fuses the range image, surface normals and uncertainty map into a single high-resolution depth map.




Figure 46: Intermediate Representations of Data for Fusion. (1) Raw fish-eye Image of an underground tunnel. (2) Range Image (depthmap) from raw LIDAR readings. Depthmap shown is warped to the space of the fish-eye image. (3) Ground truth depthmap. (4) Irradiance compensated color image, clipped to the boundaries of the LIDAR data. (5) Intepolation Uncertainty map. White values indicate scan points, while varying degrees of gray indicate increasing interpolation distance between scan points. (6) Specular-Invariant image after SUV transform. (7) Shading estimate from intensity image. (8) Surface Normal map from shading estimate utilized in MRF. RGB channels correspond to magnitude in XYZ Cartesian coordinates of unit normal vector. (9) Super-Resolution point cloud generated using MRF technique, showing detail of roof supports (right side of image). Image reproduced from (U. Wong, Lumenenhancement: Exploiting Appearance for Planetary Modeling 2012).

Some examples of intermediate data representations in the fusion process are shown in Figure 46. This data is from an underground mine scene with a mine curtain on the left, roof supports on the right and mesh, which stabilizes the exposed rock of the ceiling.

8.3.3.4 Projections of Range Data on Planar Manifolds

In the previous section, a range image (a 2D matrix structure) is utilized to fuse multi-sensor data. Though point cloud data is 3-dimensional, a low dimensional embedding is utilized for computational tractability and simplicity of neighbor associations between scattered points for interpolation. This method must be modified, however, when dealing with substantially 3D point clouds.




Figure 47: Point clouds are projected onto two-dimensional range images for processing. This introduces artifacts if the point cloud has substantial 3D information. In the left image, a point cloud of a cratered surface is stored as a range image. White areas are high frequency (highly non-planar) information that would normally produce errors. However, by storing the 3-channel Cartesian coordinates with each pixel (center image) and adaptively spacing where measurements are stored, the final reconstruction shows low error.

A *substantially 3D* point cloud features complex concavities and disjoint objects that cannot be represented as a lower dimensional manifold. Most surface terrain, for example, is well-approximated as a flat plane with small variations of height in the z-axis. However, representing a pit in the ground with this single plane approach would create vast distortions and erase any details on walls, which are perpendicular to the plane. Precisely, the projection of a point cloud x onto a range image is only valid if there exists a bijection $f: \mathbb{O}^3 \rightarrow \mathbb{O}^3$ such that the covariance matrix of $f(x)$ has ordered eigenvalues $\{\lambda_1, \lambda_2, \lambda_3\}$ where $\lambda_1, \lambda_2 \gg \lambda_3$.

With this *planarity* condition, it is possible to implement a simple algorithm which enables the super-resolution approach to be extended to a variety of substantially 3D inputs. Firstly, an eigenvalue test is applied by projecting datapoints using sets of predefined basis functions (e.g. Cartesian, cylindrical, hemispherical, etc). Each of these projections results in a planarity score, with the lowest score indicating the preferred transformation if it passes a user defined threshold. Using planarity as a metric, a large, complex map can be carved into small, locally planar pieces using an octree. Any sub-piece failing the test, is recursively subdivided into smaller pieces and retested. Subdivision terminates when all pieces are planar-projectable. Finally, all of the constituent pieces are separately processed and then recombined into a large super-resolution map.

Local distortions that persist are minimized by preserving Cartesian coordinate data when projecting to a range image (as a 3-channel Cartesian map), and enabling adaptive distortion in the projection (to space out areas of high frequency detail.).

9 CHARACTERIZATION OF EXPLORATION METHODS

This section describes the investigations performed to (1) characterize the pit and cave exploration elements; and (2) robot and sensing configurations developed by this research.

9.1 ORBITAL RECONNAISSANCE & MODELING

Orbital reconnaissance and modeling is critical to the identification of pits to target with landed missions. Orbit data collected over a wide area can identify many candidate pits, but target pits must then be imaged from multiple angles and at multiple lighting conditions for indications of void extent. Research focused on utilization of existing satellite data, such as that from LRO and MRO, to produce models of candidate pits.

9.1.1 Stereo for Modeling Pit Surroundings

Reconnaissance in the form of a pre-existing surface map is invaluable for identifying potential pits and providing pre-maps for precision landing and surface rover mission planning. Using the Moon's Lacus Mortis pit as a mission case study, a stereo digital elevation model (DEM) was constructed from two LROC Narrow Angle Camera (NAC) images of the pit and surrounding region. The LROC NAC images selected were M1105737674L and M1105759104L (see Figure 49). For stereo modeling from orbital images, there are two important considerations when selecting image pairs; namely, the difference in angle from which the terrain is viewed, or emission angle (the target is 20°), and the difference in lighting angle, which includes incidence angle and sub-solar azimuth angle (the difference in these angles should be as




Figure 49: The pit of interest as seen in LROC NAC images M1105737674L and M1105759104L. These images were used to build a stereo digital elevation model.




Figure 48: DEM of a part of the Moon's Lacus Mortis region that contains a pit. Height values are in meters.

small as possible, but definitely less than 40°). This pair of images was deliberately taken as a stereo pair, with good relative view and lighting angles. The USGS tool ISIS (Integrated Software for Imagers and Spectrometers) was used for calibration, re-projection, and bundle adjustment (Gaddis, Anderson, et al. 1997). Tie points for bundle adjustment were manually matched using the ISIS tool qnet. Ames Stereo Pipeline (ASP) was used to create the elevation map (Broxton and Edwards 2008) (Moratto, et al. 2010). Figure 48 shows the full DEM. The DEM resolution is 5m/pixel, or approximately three times the pixel size of the lower resolution of the two images. While the ASP stereo software could compute a height value for each pixel of the input images, the source images may not have enough texture to compute valid heights at that resolution. Thus it is suggested that the pixel size in the DEM be at least three times the resolution of the input images (Intelligent Robotics Group 2013).

This DEM was evaluated against a 100m/pixel DEM created from LRO wide-angle camera (WAC) images, using CloudCompare software (CloudCompare (version 2.5.5) [GPL software] 2014) (M. Robinson 2011) (Scholten, et al. 2012). Most errors were in the range of +/- 100m and, after iterative closest point (ICP) alignment between a mesh created from the WAC DEM and a point cloud of the new DEM, most errors were in the range of +/- 50m (see Figure 50).




Figure 50: Histogram of distances between created DEM and WAC DEM

Figure 51 shows a close-up of a 4km by 4km section of the DEM centered around the pit; (a) shows the DEM, (b) shows the distance from the WAC DEM in this region (after ICP), (c) shows a slope map for this region, and (d) shows the intersection error computed during the stereo processing.

The ground resolution of the WAC DEM, while nominally 100m/pixel, is actually closer to 300m/pixel. The height accuracy of the WAC DEM is estimated to be 10m to 20m. Note that in many cases, areas of with larger distances between the created DEM and the WAC DEM align with areas of high slope or

small terrain features. This makes sense, because the resolution of the newly created DEM (5m/pixel) is much higher than the WAC DEM, thus small, high-slope features will not appear in the WAC DEM. The pit is one feature with a high difference between the newly created DEM and the WAC DEM, since it is a small feature. There is also high intersection error in several areas of the pit, indicating that the pit shape was not well modeled with two-view stereo.




Figure 51: (a) Close up of 4km x 4km section of created DEM centered around the pit of interest, height scale in meters. (b) Evaluation of distance from the LRO WAC DEM, distances in meters. (c). Slope map, scale in degrees. (d). Stereo intersection error from creation of NAC DEM, values in meters.

9.1.2 Multiview Stereo Modeling of a Pit from Orbit

Stereo modeling methods designed to operate on pairs of orbital images do not necessarily do a good job handling the steep slopes and vertical cliffs of pit walls. This was observed in the experiment conducted in Section 9.1.1. Incorporating more than two views can greatly improve the model for substantially 3D terrain such as pits. An experiment in multi-view stereo for orbital modeling of pits was conducted using the Moon's Lacus Mortis pit.

The experiment used five LROC NAC images with different view angles. These included the stereo pair used in Section 9.1.1, which had emission angles (the angle between the satellite's view and a straight down, or nadir view) of 14° looking east to west (M1105737674L) and 12° looking west to east (M1105759104L), a view 47° looking east to west (M1105701957R), a view close to straight down at 3°

looking east to west (M1121075381R), and a view 33° looking south to north (M1136377273R). All of these images had similar illumination angles corresponding to mid-morning local time.

Each image was trimmed to a small region around the pit. Tie points were manually identified in the images and bundle adjustment was performed using USGS’s ISIS tool (Gaddis, Anderson, et al. 1997). Ames Stereo Pipeline was used to do the multi-view stereo reconstruction on the bundle-adjusted images (Broxton and Edwards 2008) (Moratto, et al. 2010).

Figure 52 shows a view of the pit in one of the images used for the reconstruction and a similar view of the reconstructed model. Figure 53 shows three views of the point cloud model. Note that the region that is shadowed in the images is not well modeled.




Figure 52: Map projection of an orbital image with a straight-down view (M1121075381R) of Lacus Mortis pit (left) and top view of 3D point cloud model built using multi-view stereo (right). Note that the shadowed region in the image corresponds to a region with more holes in the point cloud.




Figure 53: Views of a 3D point cloud model of Lacus Mortis pit built using multi-view stereo. The ramp observed in the images can also be seen in the model, at least for the un-shadowed side of the pit. In the shadowed area, because there is little information available to stereo, the model is quite noisy (note noisy black points at the bottom of the pit in the lower image). The top left image shows gaps along the north wall. This could be due to an overhang, or it could be due to a vertical wall that extends to the ground but is not well modeled with only one oblique view that looks north.

There is a discontinuity in the model between the southwest and northwest regions of the floor. This is not consistent with the appearance of the pit in images. Looking at the stereo intersection error computed by the stereo reconstruction software sheds some light on the issue (see Figure 55). Error is much higher in the southwest region. Also, note that the southwest region borders the shadow boundary, which shifts slightly between the different images used in this experiment. This sharp shadow boundary may have been incorrectly identified as a feature by the stereo modeling software, causing errors in the model.




Figure 55: Stereo intersection error computed by Ames Stereo Pipeline for the pit model (top) and corresponding colored view (bottom). Note that the southwest floor region that appears disconnected from adjacent patches of floor has very high error. Errors are given in meters.

To make the model easier to view, areas of high error were smoothed and the point cloud was converted to a mesh using Poisson Reconstruction (Kazhdan and Hoppe 2013) (Kazhdan, Bolitho and Hoppe 2006). This model can be seen in Figure 54.

This experiment illustrated that multi-view stereo can be successfully used to construct pit models from orbital imagery. It also showed that shadows and shadow boundaries are not well handled by this method. Improvements to the model could likely be made with approaches that explicitly consider illumination, such as shape from shadow or shape from shading methods.




Figure 54: Mesh model of Lacus Mortis pit created using multiview stereo

9.2 FLYOVER RECONNAISSANCE & MODELING

Lander flyover can capture both detailed overview data and perspectives that cannot be observed from a rover viewpoint, such as bird's-eye views. It can use LIDAR, an active sensor, to peer into shadows in a way that orbiting satellites cannot, since the LIDAR spot size at orbital distances is too large. Since lingering over a target of interest spends costly lander fuel, this work assumes that lander flyover occurs along a one-pass, always-moving trajectory as the vehicle lands.

Stitching detailed lander flyover data into a 3D model can provide an invaluable tool for planning traverses and views for a surface reconnaissance rover. It can also be used to plan anchor points for a Tyrolean robot or routes for a ramp-descending crawler. Combining lander data with data from subsequent robot explorers enables autonomous construction of high-quality 3D models of skylights that are not possible from any one platform alone. Onboard sensing for a lander conducting a skylight or pit reconnaissance is assumed to include at least LIDAR and camera. These sensors may be dual-purposed landing sensors or a purpose-built modeling package. Cameras can quickly capture data at higher resolution than LIDARs but, as a passive sensing method, cameras are of limited use for modeling shadowed regions. While a rover can take images at different times of day to avoid shadow, a lander cannot. Thus, an active sensor like a LIDAR is particularly useful.

In order for flyover modeling to succeed, the lander must be able to reliably fly over the pit during descent, and do so close enough to landing that its altitude over the pit is not too great. Identified lunar pits are mostly 100m or less in diameter. There are similarly small, identified features on Mars. To reliably fly over such a small target requires improvements in precision landing technology. To safely land very close to a pit, which may not be located in statistically safe terrain, requires improvements in hazard avoidance.

9.2.1 Precision Landing and Hazard Avoidance

Advances in terrain-relative navigation make it possible to precisely fly and land by matching lander camera images with prior satellite imagery of a planetary destination. Through this technique, landers can construct trajectories to precisely fly over features of interest (e.g., skylights) during final descent to the surface. With this technology, landers can 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, enabling collection of the data necessary to model the inner walls and detail of these skylights. Additionally, the terrain around pits can be rocky and uneven, creating numerous hazards to a safe landing. 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. Hazard detection technology at Astrobotic currently detects hazards greater than 30 cm.




Figure 56: Gimbaled LIDAR scans landing zone 80 seconds before touchdown to detect hazards and to map terrain features of interest.

9.2.2 Flyover Modeling

A 2012 paper by this NIAC team explored flyover modeling as part of an investigation of the Flyover and Surface Reconnaissance mission concept (Jones, Wong, et al. 2012). Simulated LIDAR and imagery data for flyover were generated, then combined into a terrain map and a 3D model of the pit (see Figure 57). The views taken along the trajectory were chosen to look forward into a potential cave entrance and to get good coverage of the floor. Optimization of lander trajectory and views for pit coverage were not attempted in this work, but the constraint that flyover modeling happens shortly before landing puts limits on what a lander can do. Flying above the surface and having a minimum LIDAR range, the ability of a lander to view the area beneath an overhang is limited by the aperture of the pit. Similarly, the lander's views of the pit walls will be mostly oblique (not ideal for modeling). Although the sensors were swept over the entire pit in the cross-track direction in the simulated data, they could not get good coverage of the walls.




Figure 57: Simulated pit flyover images stitched together to create a map (left), and simulated pit flyover LIDAR data combined into a 3D model and colored by height (right)

Inspired by the work on this NIAC project, another project at Carnegie Mellon is currently investigating flyover modeling of planetary pits. Through the NASA Undergraduate Flight Opportunities program, this project (led by NIAC PI Red Whittaker) has a flight planned for early 2015 on a Masten Space Systems' propulsive lander to field-test their work. Their sensor package (Figure 58) and planned flight trajectory (Figure 59) are shown below. A pit analog will be constructed to enable testing of flyover pit modeling.




Figure 58: Assembled sensor package showing stereo cameras and gimbaled LIDAR and color camera (left). Side view showing stereo camera, motor, and KVH IMU (right).




Figure 59: Planned flight trajectory dynamics simulated and visualized over test site in Mojave, CA.

9.3 SURFACE RECONNAISSANCE AND MODELING

Unlike flyover, which happens in minutes or even seconds, robots operating on the surface can spend hours or days exploring and modeling a pit. Surface reconnaissance and modeling is a good vantage point for examining pit walls, since the viewing angle from the surface is much better than from flyover. No one position will enable a rover to model the entire pit, so the rover must move to different positions around the pit. There is risk to traversing the potentially rough or unstable ground around a pit, so surface reconnaissance and modeling should be carefully planned to get the best model possible without travelling more than necessary.

Cameras can capture the highest-resolution data of pit walls and floors, though the resolution is highly dependent on the FOV. For a given camera resolution, narrower FOV lenses will produce higher resolution data on an imaged pit wall. Narrower FOV lenses also have a smaller footprint, so more images are needed to cover the entire visible area of the pit. Using a camera as a purely passive sensor is very energy efficient, but it makes seeing into shadowed regions difficult. Some light may reflect off of a pit's lit surfaces into shadow, enabling images taken with long exposure times to gather some information in shadows. This reflection is dependent on pit geometry, however, and some shadowed regions may remain impenetrable. Imaging both light and shadowed regions can be done effectively with HDR.

Surface robots could also carry a number of other sensors besides cameras. LIDAR would enable precise range measurements, even in deep shadows, though likely at lower resolution (in terms of points per meter) than cameras. Spectrometers could be used to get data about material type at points selected by scientists. Gravimeters and GPR could be used to survey the region around a pit to look for evidence of subsurface voids that may not be visible from the pit. The robot would also likely carry an IMU to assist with navigation.

In this work, study of surface reconnaissance and modeling focused on camera sensing. Section 9.3.1 discusses building 3D models from surface rover images. Section 9.3.2 investigates HDR imaging from the rim of a pit. LIDAR data was also collected from a pit rim. With the help of careful surveying, this LIDAR data was used to construct ground truth models against which models built from camera data could be compared (see

Appendix B: Craters of the Moon Field Demonstration and Surveying).

9.3.1 Building 3D Pit Models from Surface Rover Images

One concept for surface reconnaissance and modeling is to tour the rim with a surface rover and build high-fidelity models by cross-pit imaging. Camera-only methods are power and weight efficient, though models have scale ambiguities and perspective artifacts. The quality of these models is particularly dependent on the positions of imaging and natural illumination. In a planetary environment, the illumination angle is determined by the motion of the Sun over time, which is well understood. View angles are determined by rover position. The set of views - and the trajectories between them - must be carefully planned in order to effectively and efficiently model the pit. Heather Jones' work (mutually supported by a NASA Space Technology Research Fellowship) is developing techniques to determine where to view from and when to capture images that meet illumination and perspective constraints in order to optimize model quality.

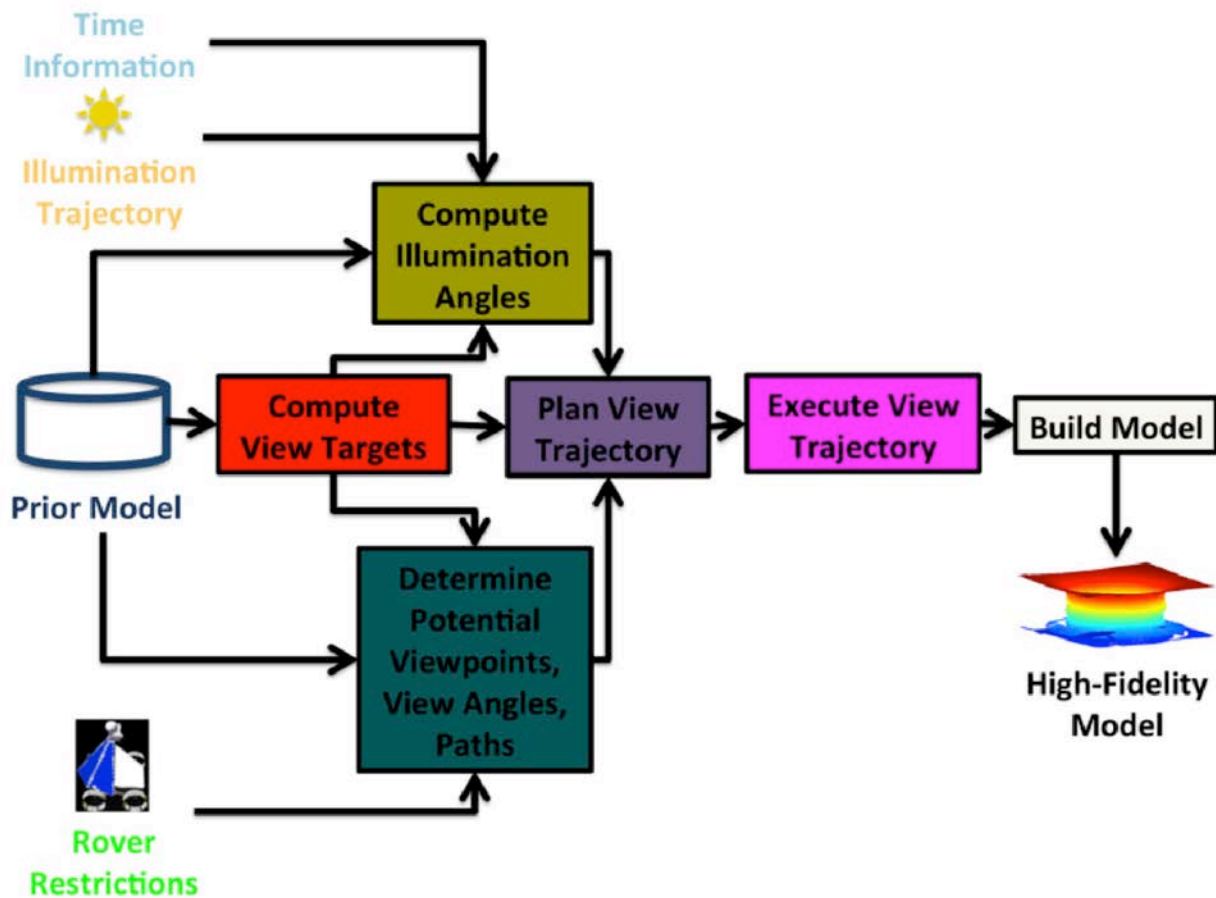


Figure 60: Flow diagram for creating a high-fidelity pit model from a surface robot

Figure 60 shows the process for planning and executing view trajectories to build detailed models of pits from surface rovers (Jones, Tabib and Whittaker n.d.). The planning process begins with an initial coarse model of a pit, which may be as simple as a cylinder of appropriate diameter, or may be a more complex

model built from satellite imagery or lander flyover. Information about time (when the mission occurs, and how long it lasts), illumination (how the solar illumination angle changes during a mission), and rover restrictions (maximum speed and mobility parameters) is also incorporated. Pre-processing steps compute a set of pit surface patches (see Figure 61), determine the illumination angles on these patches over the course of a mission, and find potential rover viewpoints. This information is then used to plan a view trajectory for a rover's camera (see Figure 62). The view trajectory includes the images to be taken, rover position when imaging, and image capture time. The view trajectory is executed to capture the desired images. Finally, the captured images are used to build a high-fidelity model of the pit.

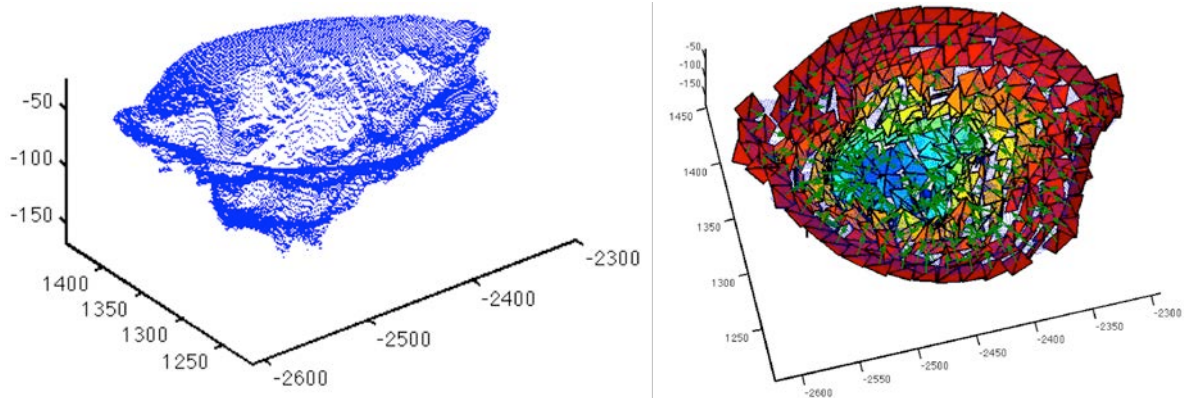


Figure 61: A coarse model of a pit from satellite stereo imagery is used to generate a set of target patches for a rover to view. The average normal vector for each patch is determined.




Figure 62: Example of a view trajectory as shown in a simple pit simulation. A coarse pit model (a cylinder) has been discretized into six target patches. Lighting has been simulated over the course of one day (light patches are illuminated, dark patches are not). Green circles illustrate potential rover positions. Asterisks show the rover position at each time, and cyan lines show the camera views selected at that time.

9.3.1.1 Field Testing of Surface Rover Modeling at Indian Tunnel

As a feasibility study for 3D pit modeling from a surface rover perspective using only camera images, a field test was conducted at the Indian Tunnel analog site (see Appendix A:

Indian Tunnel). The large skylight near the southern end of Indian tunnel (the south-most extent of underground work) was the focus for this test. Images were collected from tripod-mounted DSLR cameras to simulate what might be visible from a surface rover approaching the pit. Two Canon EOS Rebel T3 DSLR cameras were used. Both had manually adjustable zoom lenses, and both were set to 18mm focal length, yielding a field of view of approximately 45° vertical and 63° horizontal.

A set of 205 camera views were collected from 14 stations. Tripod stations were spaced with a target distance of roughly 4.5m ground distance, though stations were adjusted as needed to accommodate hazardous terrain and precise measurements of station positions were not done. At each station, the tripod was manually leveled by eye using an integrated bubble level. The tripod was adjusted such that the camera tilt was zero. The camera was panned to the left-most edge of the pit (in the current view). Additional views were taken by tilting down by 22.5° to 45° and panning by 30° until the right-most edge of the pit was reached. The test was conducted over less than 1.5 hours on an overcast day, so the illumination did not vary significantly between images.

Each image was taken with auto-exposure using center-weighted average metering. Other camera parameters were fixed for all images, including ISO 100, aperture F8 and white balance of 5200K (daylight).

A 3D model of the pit was constructed (Jones, Tabib and Whittaker n.d.) using structure from motion (SIFT for feature detection (Lowe 1999), Bundler for bundle adjustment (Snavely, Seitz and Szeliski 2007) (Snavely, Seitz and Szeliski 2006), and CMVS/PMVS2 for dense reconstruction (Furukawa and Ponce 2007) (Furukawa and Ponce 2010)). Figure 63 and Figure 64 show views of the colored point cloud created using this method.



Figure 63: Side view of colored point cloud model created from dense reconstruction with CMVS/PMVS2. The skylight pit is in the center, and segments of lava tube cave can be seen extending on either side.




Figure 64: Top view of colored point cloud model created from dense reconstruction with CMVS/PMVS2. The skylight pit is in the center, rubble breakdown on the floor. Accessible tube tunnel segments extend on either side, and a smaller lava tube branch, at a higher level than the main tunnel, extends from the bottom in this view.

To evaluate the success of the reconstruction, the point cloud built from camera images was compared against a ground-truth point cloud built using a FARO laser. Because structure from motion does not preserve scale, the camera reconstruction first had to be scaled. The data also had to be rotated and translated to match the frame of the ground truth model. Initial alignment was done by manually identifying a small set of features in both models and using these to compute scale, rotation, and translation. In a mission scenario, scale could be obtained from knowledge of a robot's relative position between pairs of images, or from comparing a pit model to orbital images.

After initial alignment, the camera reconstruction was cropped to the same bounds as the ground truth model. While the ground truth model has extensive surface coverage, the focus of the camera modeling was on the pit walls and floor (the surface terrain is only covered in a few sparse patches), so surface terrain was cropped out of both models before comparison.

After cropping, both models were loaded with CloudCompare software (CloudCompare (version 2.5.5) [GPL software] 2014). Fine alignment and scale was done using the iterative closest point method (Zhang 1994), then distance between the two point clouds was calculated. A maximum distance of 8.1m, a mean distance of 0.042m, and a standard deviation of 0.16m were computed. Figure 65 and Figure 66 show a more detailed view of the point cloud distance error. As can be seen in Figure 65, distance errors above 1m (including the 8.1m maximum distance) tend to be outliers, which could easily be removed by filtering. The vast majority of the points have distance errors in the single-digit centimeter range (see histograms in Figure 67 and Figure 68). Figure 69 shows detail for the error evaluation over the range 0cm to 10cm. This indicates that error tends to be lower on the pit walls, and higher at the extremes of the visible tunnel region. This is not surprising given the more oblique viewing angles and greater distance from the camera for the tunnel segments. There also appears to be high error in some areas of the pit floor. While some of this appears to be noise, other high-error patches are due to gaps in the ground-truth data, likely caused by boulders occluding the laser's view.

Figure 70 shows the ground-truth model overlaid on the distance error evaluation to give an indication of model coverage. The lower corner of the lava tube segment on the left and the lower portion of the small tributary lava tube on the bottom, for example, are not included in the camera model. The ground-truth model is also denser than the camera model. Note that in some cases, as illustrated in Figure 71, apparent error in the camera point cloud results from gaps in the ground-truth data.

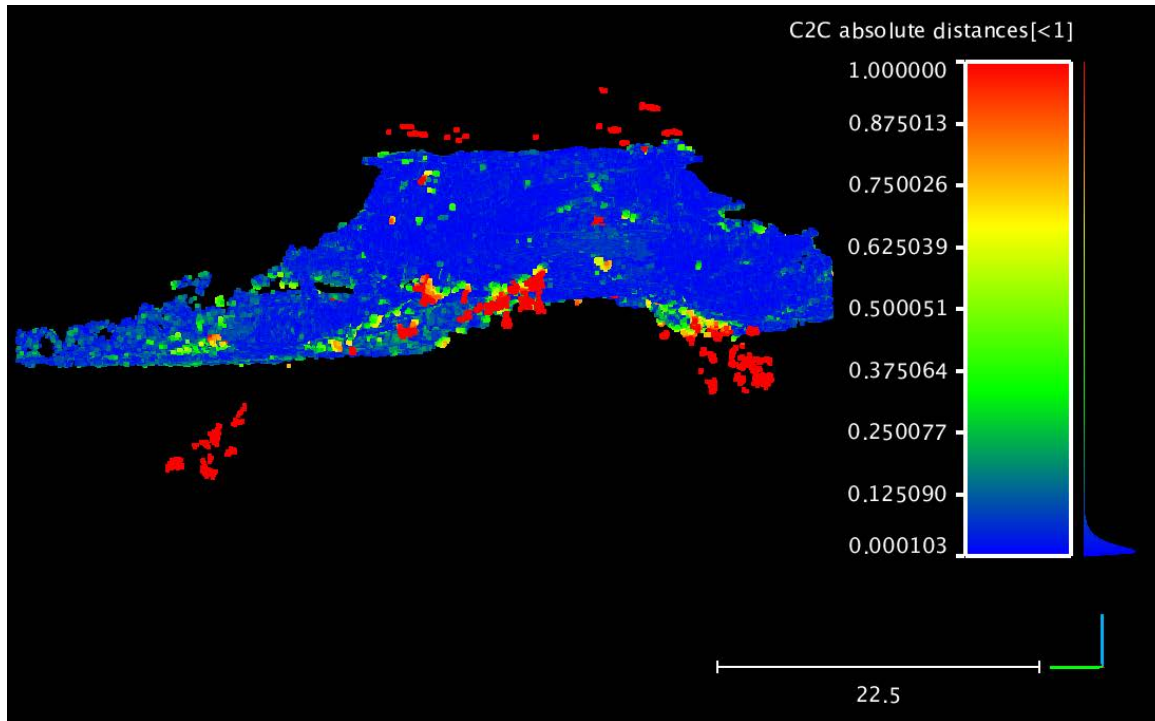


Figure 65: Error evaluation for point cloud model, side view. Distances are in meters.

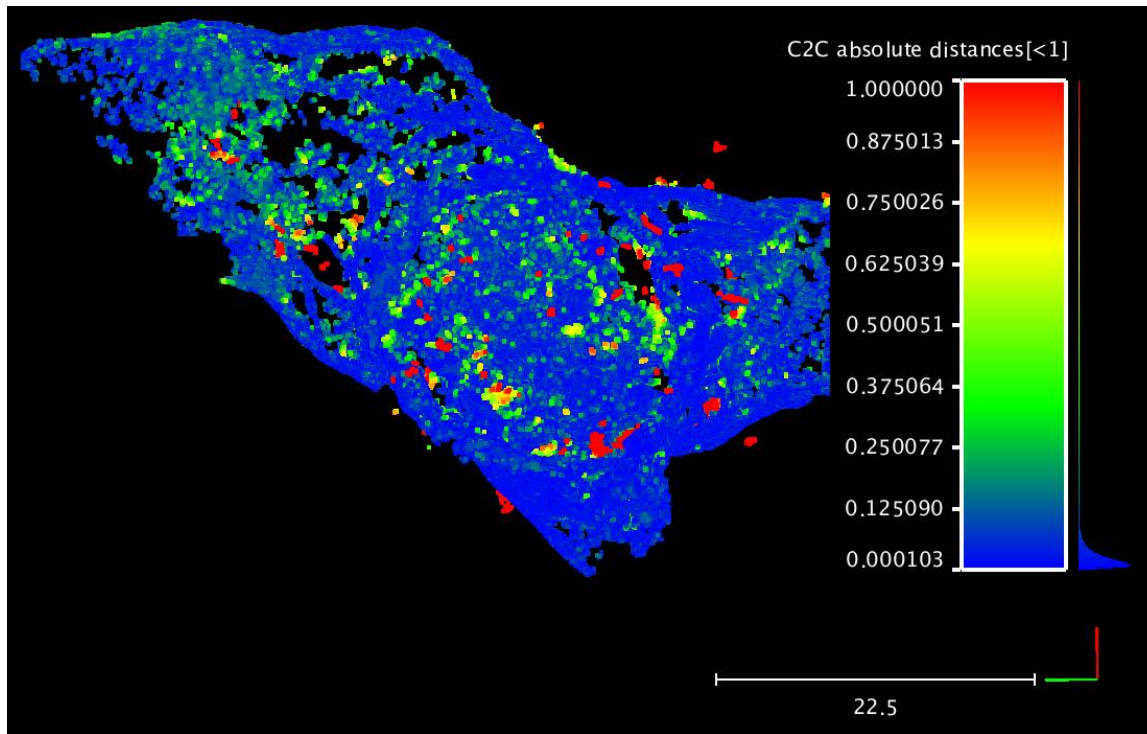


Figure 66: Error evaluation for point cloud model, top view. Distances are in meters.




Figure 67: Histogram of error distances for camera reconstruction model over the range of 0m to 1m

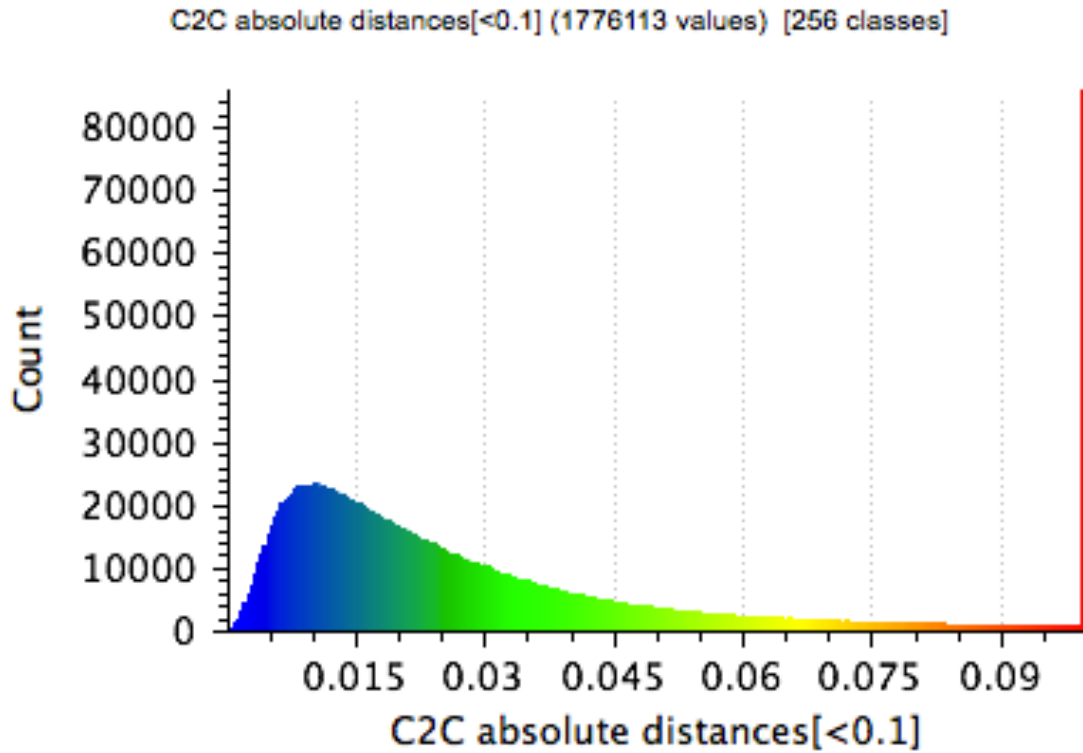


Figure 68: Histogram of error distances for camera reconstruction model over the range of 0 to 10 cm




Figure 69: Error evaluation for point cloud model, top view, showing error range 0 to 10cm. Distances are in meters.

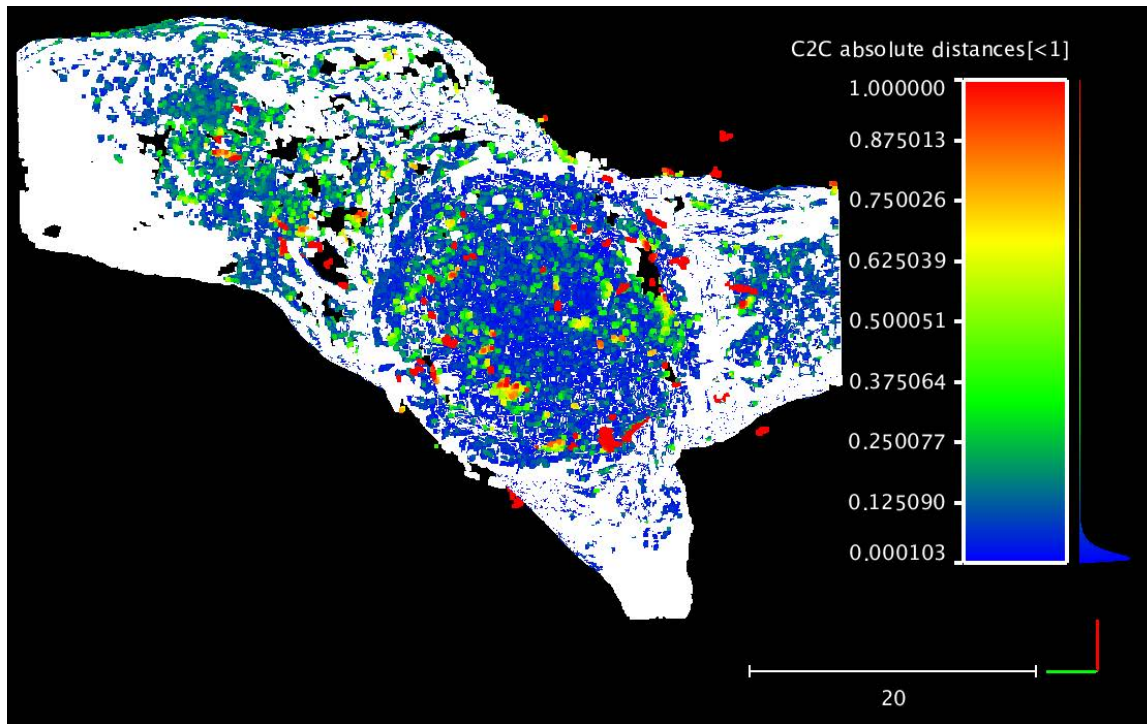


Figure 70: Overlay of ground-truth data (white) on error evaluation (same view as Figure 66) to show coverage.




Figure 71: A close-up of model with white ground-truth points and error-colored camera model points shows that in some cases, like the green patch in the center of this image, the apparent error is likely due to a gap in the ground-truth data. Distances are in meters.

9.3.2 High Dynamic Range Imagery of a Pit from a Surface Rover

Testing at King’s Bowl pit (See Appendix A: King’s Bowl Pit) investigated HDR imaging of a pit from a surface rover. Panoramas of the Rift Pit were created by first combining multiple exposures of a single image into a High Dynamic Range (HDR) image. This process is illustrated in Figure 72. Google’s HDR Efx Pro 2 software was used to combine multiple exposures into a single HDR image. Next, the HDR images were stitched together into a high-resolution panorama using Autostitch Panorama Stitcher. The resulting panorama can be seen in Figure 73.




Figure 72: Multiple exposures of the same image are combined into a high dynamic range (HDR) image.




Figure 73: HDR panorama taken from a tripod placed at W14.

To highlight the effectiveness of the HDR at conveying contrast, the HDR and LDR images were compared using a contrast metric (shown in Figure 74). The contrast metric gives a score for each pixel in the image based on how much contrast there exists between it and the pixels around it. Because an HDR image is composed of 3 images taken at different exposures, it should have greater contrast than a single image taken with a uniform exposure. In order to validate this theory, the LDR contrast values are subtracted from the HDR contrast values and the results are displayed in Figure 75. The colormap varies from black to white, with black meaning the contrast values are nearly identical for the LDR and HDR images. The white values imply that the HDR image has significantly higher contrast at that location.




Figure 74: LDR panorama is shown on the left and HDR panorama on the right.




Figure 75: Colormap of contrast between the HDR and LDR image. Darker colors (minimum value = 0) means lower contrast between the HDR contrast value is close to the LDR contrast value. Lighter colors (maximum value = 1) means the HDR image has greater contrast than the LDR image.

9.4 PIT ACCESS AND MODELING

While pit descent is an extremely challenging mobility problem, it provides unique views and close-ups of pit walls and floors and views into void spaces. Cameras, LIDARs, IMUs and spectrometers are useful sensors during pit descent. Descent can also enable sampling on pit walls and floors if a robot carries instruments to collect (and perhaps analyze) samples. For most other pit and cave exploration vantage points, the primary task is sensing and modeling. For pit descent, however, transport is the primary task, and infrastructure link may also be quite important.

Two robot configurations for pit descent were selected for examination in this work (see Section 6, Section 11). For the Krawler configuration, the process of modeling is similar to what would be done from a surface rover. For the Tyrobot, sensors are suspended by cables. This unique mobility mode necessitated further study of modeling from this platform.

Work presented in this section discusses demonstration of pit access by robotic descent of a scree slope and descent from a Tyrolean line (Section 9.4.1). Field tests of modeling from a Tyrobot are presented in Sections 9.4.2 & 9.4.3). Mine locations utilized by this research are described in Appendix A: Mine Pits.

9.4.1 Pit Access Mission Simulations

In order to validate the two access methods selected for investigation, a mission simulation field test was conducted at a strip mine. Figure 76 shows a model of the mine pit, with descent locations and Tyrolean line placement indicated.

For deployment of a crawler by a Tyrolean robot, the test demonstrated descent into the pit, deployment and floor exploration by the crawler, re-docking of the crawler to the Tyrolean robot, and transport of the crawler after re-docking. Figure 77 shows the deployment. Figure 78 shows exploration of the pit floor by the crawler. Figure 79 shows the re-docking.




Figure 76: Diagram of mission simulation tests conducted shown on a 3D model of the mine pit. Scree slope descent was performed by crawler in two locations (indicated by red arrows). Deployment of Krawler from Tyrobot was performed with the Tyrolean line located as shown in yellow.



Figure 77: Deployment of Krawler into pit by Tyrobot



Figure 78: Exploration of pit floor by Krawler



Figure 79: Krawler approaches and re-docks with Tyrobot, facilitating re-charge or transport to another location in the pit.

For pit access by crawler descent, descent of two different slopes was performed. Figure 80 shows slope descent by the Krawler robot. Slope ascent was also demonstrated.



Figure 80: Rock crawler descent was demonstrated for two different slopes at the mine.

This test successfully demonstrated both pit access methods proposed in this work using the developed robot platforms. Pit access by scree slope descent with a crawler is the simplest method – it requires only one robot and works for pit access on scree slopes ranging from gentle ramps to near angle-of-repose slopes. Pit access by deployment of a crawler with a Tyrolean robot can work even for pits with steep walls, and allows docking and repositioning of the crawler by the Tyrolean robot.

9.4.2 Tyrobot Modeling Field Test – Scan Patterns and Speeds

A test of the modeling capability of the Tyrobot configuration was conducted at a mine pit site (see Appendix A: Analog Field Testing Sites). The objectives of the modeling investigation were to compare scan trajectories and determine if the Tyrobot configuration could provide adequate views, pointing stability, and damping for quality sensing. To quantify this, the sensing payload featured a reflective survey prism tracked by a robotic total station to provide continuous feedback on the three-dimensional position of the payload. This configuration demonstrates the possibility for a two-robot modeling scenario, where a rim survey robot could track and register a SkyCam payload for significant greater accuracy than a single-robot system. This evaluation of modeling capability comprised two trials, each testing a different sensor trajectory for modeling the skylight walls while recording comprehensive statistics on the position, settling time, and view angles.




Figure 81: The skylight wall-modeling payload was surveyed with an automated total station. The total station provided valuable data on the stability of the Tyrolean system for sensing and ground truth locations for modeling viewpoints. This configuration also illustrates a possible 2-robot scenario for modeling the skylight. A survey robot on the rim could provide line-of-sight tracking for the SkyCam modeler to greatly increase accuracy over internal-only sensing.

During the first modeling experiment, the SkyCam was moved back and forth to scan the wall in horizontal stripes at decreasing heights. In a second test, vertical stripes progress from left to right. Figure 82 shows the position of the SkyCam during the two experiments, as measured from the survey station. Yellow caution tape was used to provide a visual marker for the size of the scan area in comparison to the wall. Preliminary analysis shows that the horizontal stripes provided a smoother trajectory with less undesirable ego-motion (pendulum, vibration, or oscillation) than the vertical stripes, which may be the result of acceleration required to transition between vertical stripes introducing a pendulum effect on the payload. The swing requires significant settling time, which was not observed in the trial. This effect would be magnified in lunar scenarios where there is no damping from atmospheric drag.




Figure 82: Tracked position of SkyCam during horizontal stripe scan trajectory (blue) and modeling experiment 2 (red).

GoPro camera images were stitched to create quasi-orthographic photo-panoramas of the high wall. These panoramas are generated from linear motion, and are devoid of many of the perspective artifacts in common rotation-only panoramas used in hobby imaging. Moreover, the process of stitching generates estimates of camera views that can be utilized for a position or orientation estimate. Figure 83 shows panoramas constructed from SkyCam data using horizontal-stripe trajectories. The data shown is from two stripes (topmost and bottommost). Each linear panorama comprises 13 photos. The only visible artifact is a discontinuity from movement of the caution tape due to wind.




Figure 83: Stitched photo-panorama of high wall (left), and view of high wall stripes overlaid with model constructed from SkyCam data (right). A true-color stitched photo from 13 horizontal views illustrates high-detail sensing capability from this method (bottom).

9.4.2.1 Ground Truth with Survey

To evaluate Tyrobot mapping capability, a ground truth model of the pit was by stitching 4 static survey scans at each corner of the pit into a single model. The ground truth model features an areal density of about 1 measurement per sq. cm.




Figure 84: A 1cm ground truth model was constructed from four static surveyed LIDAR scans (green arrows). This point cloud is shaded with the reflectivity of surfaces. A crushed coal pile in the center of the pit is noticeably less reflective than the walls.

9.4.3 Tyrolean Modeling Performance

The dynamic nature of Tyrobot (exaggerated by cable suspension) presents significant challenges to LIDAR modeling. Pendulum motion and vibrations corrupt data and result in scans of varying density and missed areas. At its core, the LIDAR is a single beam sensor that is scanned over the scene utilizing the motion of mirrors and motors. When the sensor experiences ego-motion, the encoding of the scan axes no longer produce a valid transform into an absolute reference frame. Other sensors, like inertial measurement units and cameras (which track pose), must be utilized to reverse the effects of the ego-motion to produce an accurate scan. This process, along with correcting calibration errors, is called “dewarping” the scan. Inertial measurement dewarping has been implemented to correct the Tyrobot models by rotating planar slices such that the measured gravity vector aligns with the vertical axis. In the future, vision-based tracking will be implemented for greater accuracy.

9.4.3.1 LIDAR Calibration

Tyrobot uses a Hokuyo LIDAR to generate point clouds of its surroundings. The Hokuyo measures 270° in a single plane at $.25^\circ$ increments. A motor spins the LIDAR 360° in order to construct a full point cloud, which only works correctly if the exact orientation of the LIDAR is known with respect to the axis of rotation. Misalignment errors are common in robotics, and are the result of as-built inaccuracies and sensor placement decisions. Even small misalignments can have major effects in the quality of sensing. For example, a misalignment of 1° in a rotation axis may show up as half a meter of linear error over the long distances that LIDAR is expected to perform.



Figure 85: A “dewarped” geometric point cloud of mine test site (left) and an image of the same scene (right).

In an effort to reduce error, a comprehensive calibration routine specific to rotary scanners was developed and implemented for this work. Because the Tyrobot LIDAR measures 270° in a plane by spinning it more than 180° around the axis of rotation, there is overlap in the point cloud. If the LIDAR is aligned perfectly, the overlapping slices of the point cloud will match each other. Using the Iterative Closest Point algorithm (ICP) and comparing corresponding overlapping slices, calibration parameters ranging from angular misalignment of the sensor body to shift of the LIDAR’s internal optics from ideal can be estimated. The calibration routine is essentially an optimization that finds the parameters that “best fit” large amounts of observed scan data. Once the parameters are found, they are stored and applied to every subsequent LIDAR scan through the forward transforms that bring range measurements into a global Cartesian (x,y,z) coordinate frame. This process is known as “dewarping”, as LIDAR scans frequently appear as fuzzy or warped versions of the scene before correction.




Figure 86: Side view of Tyrobot test scans in a tunnel before (top) and after (bottom) dewarping. Because of the rotary nature of the device, a misalignment in pitch shows up as multiple copies of the floor and walls far away from the sensor origin. After calibration, the double-walling effect is eliminated.

The calibration algorithm found linear offsets of 2mm and 11mm as well angular offsets of 0.02 radians and 0.002 radians from ideal. The results of the calibration are shown in Figure 86. Both point clouds are generated from the same raw data. The red point cloud is pre-calibration, and the blue is post-calibration. The “double-walling” effect evident on the lower right of the raw point cloud is clearly fixed by the calibration. This same calibration is applied to the point clouds from experiments in terrestrial pit analogs. An example of five aligned scans is shown in Figure 85. Performing this calibration ensures the most accurate data possible from the Tyrobot-based modeling and informs the expected accuracy of those data.

9.4.3.2 Alignment of Camera and LIDAR

LIDAR alone provides only geometric structure, not color information. The point cloud is colorized by integrating the imagery from the fisheye camera mounted with the LIDAR on the sensor head. To colorize the point cloud, the fisheye camera was first calibrated, then the transformation between the camera and the laser was found. To calibrate the camera, an omnidirectional camera calibration toolbox was used in order to calibrate the fisheye lens effectively (Mei and Rives 2007). Calibration required taking many images of a known planar grid, which the algorithm then uses to find the best-fit camera model that minimizes re-projection error of the grid corners. An example calibration image and post-calibration grid re-projections are shown in Figure 87.

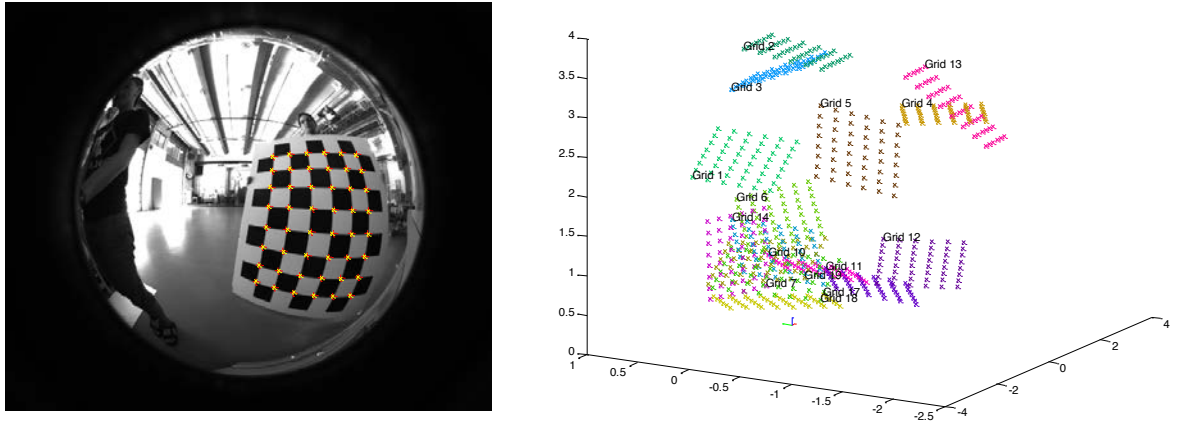


Figure 87: An example calibration image of a known planar grid (left) and re-projection of the grids after calibration (right).

To align the camera image and LIDAR point clouds, calibration images and point clouds were taken in a static environment with a 3D calibration target. An initial guess of the transformation between the camera and LIDAR was used to project the LIDAR into an image using the calibrated camera model. The LIDAR image was then colorized by depth. Matching points in the LIDAR image and the camera image were selected manually. An example showing corresponding points is shown in Figure 88. Finally, an optimization routine was run to minimize the distance between sets of corresponding points by adjusting the six-degree-of-freedom transformation between the LIDAR and camera reference frames.




Figure 88: Corresponding points selected between fisheye camera image (left) and projection of LIDAR point cloud colorized by depth (right).

The resulting best-fit transformation produced an average error of 16 pixels between the camera and projected LIDAR images given the manually selected matching points. A plot of the resulting points is shown in Figure 89. There is lower error in the center of the image and more error as points get further away from the center. This is likely caused by errors in the camera calibration, which is more accurate towards the center of the image for the fisheye camera. There is also human error inherent in the

manual selection of points. The matching points selected in the camera and projected LIDAR image were not exactly the same to begin with, which made it impossible for the optimization to find a perfect transformation that resulted in exact matches of all the points. An example colored scan is shown in Figure 90 from a mine pit field test.




Figure 89: Matching points between the LIDAR (blue) and camera (red) plotted after optimization.




Figure 90: A colored scan from a Mine Pit field test.

9.4.3.3 Scan Alignment

During data collection, Tyrobot takes sequences of overlapping 360° scans. Each scan contains information about one part of its environment. A better model with more coverage is produced by integrating the scans. Initial guesses at transformations between scans were determined using odometry from Tyrobot that measured absolute distance in traverse on the Tyrolean line, as well as the distance of the sensor payload from the robot. Error in these guesses is due to the slack in the traverse line, line stretch, and slip. The iterative closest point (ICP) algorithm was used to refine the initial guesses of transformations and produce a final point cloud using multiple scans performed during the test.

9.4.3.4 Modeling Performance

Modeling performance was estimated using the data from the three experiments during one field test at a strip mine. Each of the three models was individually matched to the ground truth FARO model, first using a manual estimate then refined with the ICP algorithm. The aligned models are shown with the ground truth data in Figure 91 through Figure 94.



Figure 91: First of four Robots scan (saturated section) overlaid with FARO ground truth measurements



Figure 92: Second of four robot scans (saturated section) overlaid with FARO ground truth measurements



Figure 93: Third of four robot scans (saturated section) overlaid with FARO ground truth measurements



Figure 94: Fourth of four robot scans (saturated section) overlaid with FARO ground truth measurements

The modeling error was evaluated by first combining the models from the three experiments into one large model. That model was then down sampled so that the points in the resulting point cloud were evenly spaced. Error was determined by matching points to their closest point in the ground truth point cloud using ICP and then measuring error in meters. The average error was .47m and the median error was .31m. A plot of the error over the entire model is given in Figure 95. This error is low-moderate and predictable for this type of device.

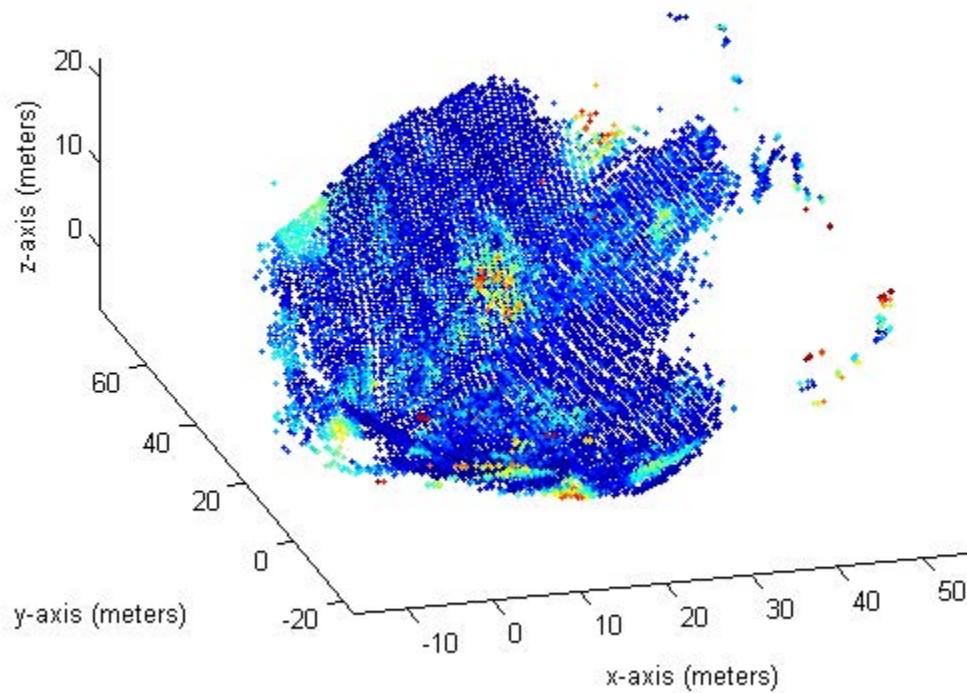


Figure 95: A plot of the modeling error in the Sheep's Ridge Mine field experiment. Blue corresponds to near zero error and red corresponds to around 3m error. Some outliers were removed so that they not saturate the color map.

Coverage of the model was approximated by counting the number of points in each 1m voxel. Voxels with fewer than 5 points were considered outliers and subsequently ignored. The average number of points in a 1m voxel was 5170 and the median was 773. The average was much higher than the median because there was very high point density at points where the scanner got very close to the wall or ground. Overall, coverage was very good, which was expected, given the varying distances from the surface of the mine.




Figure 96: A histogram showing the number of points per 1m voxel in the model. The histogram was cut off after 5000 points per voxel to show more detail.

9.5 SUBSURFACE EXPLORATION

For vantage points from orbit through pit descent, solar illumination is available for most areas modeled. For subsurface exploration, a robot must operate in the dark. Any optical sensing must be active, and cameras and spectrometers must be accompanied by an appropriate light source. While all sensors discussed for use from other vantage points are potentially useful in subsurface exploration, available power for sensors is most limited in the subsurface. Subsurface exploration robots will likely carry sensors specialized for destination-specific investigations (e.g., to search for life on Mars, to search for volatiles on the Moon).

This work investigated four modeling paradigms for subsurface exploration: (a) using structure from motion techniques with only a camera and a light source inside the cave (Section 9.5.1); (b) using structure from motion with robot-mounted cameras (Section 9.5.2); (c) using a robot-mounted flash LIDAR (Section 9.5.3); and (d) using Lumenhancement techniques for high-resolution modeling with LIDAR, camera, and a light source (Section 9.5.4).

9.5.1 Cave Modeling by Structure from Motion with a Hand-Held Sensor

The test site was Pluto’s Cave, which is a lava tube cave with a skylight. (See Appendix A: Pluto’s Cave). In this test, a dataset of over 1,700 images was collected using a DSLR camera with an LED light array. Images were taken of various features, including tunnels, skylights, and a natural bridge.

Reconstruction of 3D geometry from some of these images was performed using structure from motion techniques. Features were detected and matched across images using SIFT (Lowe 1999), bundle adjustment was done using Bundler (Snavely, Seitz and Szeliski 2007) (Snavely, Seitz and Szeliski 2006), and dense reconstruction was done using the CMVS/PMVS2 stereo package (Furukawa and Ponce 2007) (Furukawa and Ponce 2010). This step produced colorized 3D point clouds with a normal vector for each point. Images of cave features and their 3D reconstructions are shown in Figure 97, Figure 98, Figure 99, and Figure 100.




Figure 97: Image and reconstructed point-cloud model of a lava tube tunnel. 280 source images.




Figure 98: Image and reconstructed point-cloud model of a natural bridge. 238 source images.




Figure 99: Image and reconstructed point-cloud model of a skylight with ramp entrance. 192 source images.




Figure 100: Image and reconstructed point-cloud model of a steep-walled skylight. 191 source images.

This test demonstrated qualitatively that modeling using only a camera, a light source, and structure from motion techniques can be used successfully for planetary analog caves. However, since no ground-truth for the cave shape was available, quantitative evaluations of model quality are not possible from this test.

9.5.2 Cave Modeling by Structure from Motion with Robot-Mounted Sensors

Structure from motion was used to reconstruct the terrain seen by the robot as it drove through Indian Tunnel cave, using the sparse reconstruction step of VisualSFM (Wu 2013).

Results of structure from motion for the left camera, the right camera, and both cameras together were computed. These sparse reconstructions were then compared against a ground truth model of the cave generated by a FARO LIDAR. Each SFM result was scaled, translated, and rotated manually to get a rough alignment with the ground truth model. The roughly aligned SFM result was cropped to the bounds of the ground truth model and then fine alignment was done using the iterative closest point method as implemented in CloudCompare (CloudCompare (version 2.5.5) [GPL software] 2014). CloudCompare was also used to compute the distance of the SFM results from the ground truth point cloud. Figure 103 shows the robot traverse point cloud overlaid on the ground truth model. Figure 104 through Figure 106 show visualizations of the distance error for each model (over the range [0,1] meters) in a top-down view, as well as histograms of the errors for all points.




Figure 101: Krawler during field test in Indian Tunnel




Figure 102: Side view of sparse reconstruction of robot route through Indian Tunnel Cave. Note the arched shape of the cave walls. Camera positions are indicated in green.




Figure 103: Point cloud model reconstructed from the robot's left camera (blue, green, yellow, and red points) overlaid on a the ground truth model of the cave segment (true color). Note that the red/yellow/green clusters in the center of the image are likely due to reconstruction of people videotaping the robot. These people were not present in the ground-truth model.

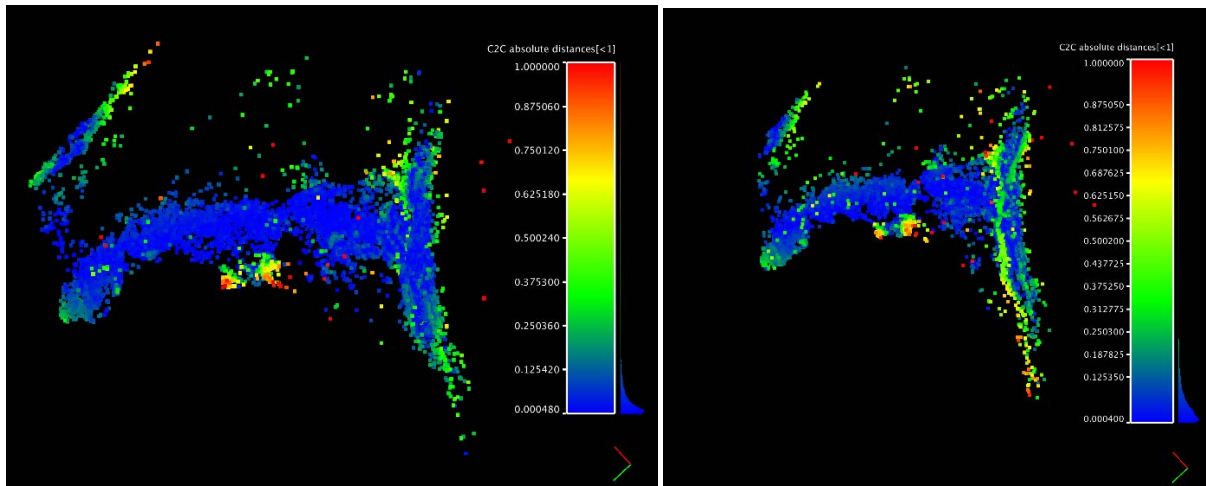


Figure 104: Visualization of distance errors for left camera model and right camera model.

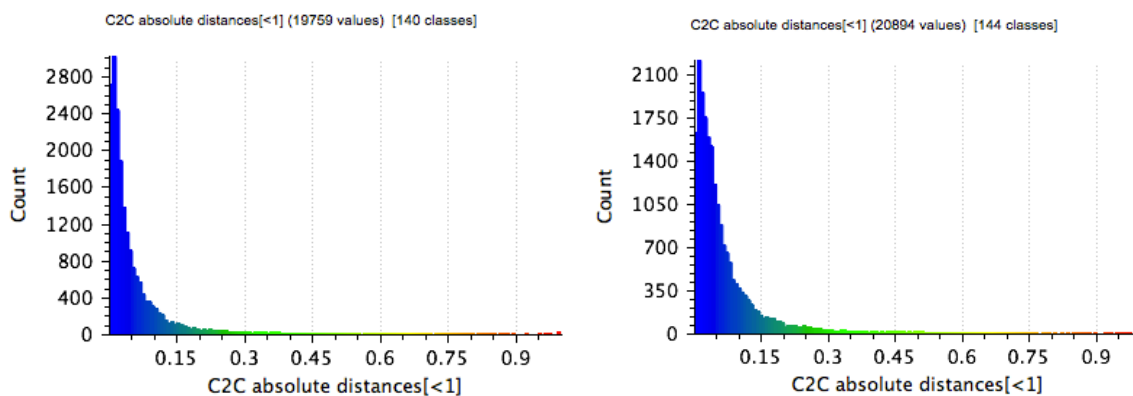


Figure 105: Histogram of distance errors for left camera model and right camera model. Distances in meters.

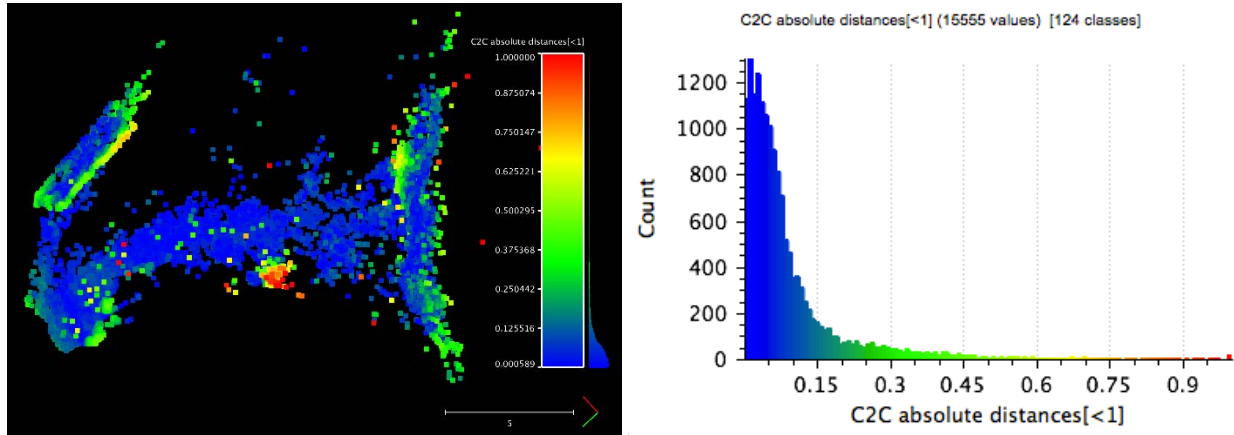


Figure 106: Visualization and histogram of distance errors for model from both cameras..

This test demonstrated that a robot-mounted camera and light source can be used with structure from motion techniques to model a planetary analog cave. While structure from motion only provides a reconstruction up to scale, after setting the scale using the ground-truth model, most points in the reconstructed model were less than 0.15m away from the ground truth.

9.5.3 Cave Modeling with Solid State Flash LIDAR

Testing of commodity flash LIDAR (also known as a ToF Camera) technology for mobile mapping was also performed at the Indian Tunnel cave site with a SwissRanger 4000 (SR4000). Solid-state sensors like the SR4000 are of particular interest for agile robotics, as many range readings within the field of view are simultaneous captured. This prevents blurring of data during quick motion and preserves gross geometry of the map. The lack of moving parts, such as those required to raster a beam, mean that the sensor can be greatly miniaturized for small, low-power planetary explorers. The SR4000 utilized here is approximately the size of a Rubik’s cube and could be carried by a robot slightly larger than a radio-controlled car.

Many of the known disadvantages of flash LIDAR technology stem from the need to illuminate the scene with sufficient energy for clean range detection. As with a photographer’s flash, the illumination must be powerful and energy-consumptive with successive use, very narrow in field of view, or deliver a grainy exposure (and low resolution). For commodity terrestrial sensors, the last option is typically chosen and mitigated with generous filtering to clean the data. Dark tunnels are an advantageous environment for these sensors due to lack of natural illumination, which increases the signal to noise ratio. Testing at Indian Tunnel investigated the above assumptions while characterizing the modeling performance of the SR4000.




Figure 107: Two consecutive frames of Infrared intensity maps generated from solid-state flash LIDAR on a moving Krawler robot (Top). Keypoint correspondences in the depth data provide 3D rigid transformation and pose of the robot between the frames (Bottom). The scene consists of a relatively flat terrain with a large boulder and people in the distance. Significant noise is apparent at range.

Figure 107 illustrates typical data from this sensor in the dark cave. The figure shows two consecutive frames of depth and intensity data captured from a moving robot at native resolution. The scene is largely a planar rubble field, with some much larger rocks and test personnel in the background. Intensity is a function of the natural reflectivity and distance of an object, which obeys an inverse square law. The sensor applies a distance correction internally so that the objects are normalized in reflectivity while they are above a threshold and effectively zero otherwise. It is possible to see this transition occur at approximately 8m, given the reflectivity of this particular environment – many of the background objects do not show up in the intensity map. Range data is invariant to true distance; however, the noise increases as a function of distance and reflectivity. Utilizing a neighborhood variance filter, if the local noise exceeds a threshold, the readings are filtered out. This results in the “static” like coverage of the background in the lower images. Though visually jarring, the missing data is preferable to a noisy data in building a map.




Figure 108: Many registered frames of flash LIDAR data from a Krawler exploration in Indian Tunnel. Each frame is a unique color. The robot trajectory is recovered from the displacements between frames.

The SR4000 sensor captures at 15 frames per second, so the result of an exploration trajectory is a continuous “movie” of depth and intensity values. With knowledge of the lens transformation, a point cloud is generated for each 2D frame using the equation $x_i = d_i \cdot \mathbf{h}_i^w$, where x_i is the 3D coordinates, d_i is the depth reading, and \mathbf{h}_i^w is the 3-element normal ray from calibration of pixel i respectively. To register all of the point clouds and build a map, we use a progressive keyframe ICP algorithm inspired by (Kerl, Sturm and Cremers 2013). Corresponding 2D features are detected in consecutive intensity and range frames. These features are then projected into 3D coordinates using the lens transformation. An initial fit using only these features is performed to align the two clouds with a rigid body transformation. Then, a range/intensity ICP minimization is performed on all the points for fine registration. A history of frames is built, starting with the *keyframe* (first) and each new frame of data is registered against all of the past frames. If the overlap between the current frame and the keyframe falls below a certain degree, a new keyframe closer to the current time is selected and the process is restarted with a new history window. This scheme greatly reduces drift from simple neighbor registrations.




Figure 109: Despite high frame rate and capability for agile ego-motion, the low density and high noise of the ToF Camera sensor produces maps with many artifacts including misalignment and holes from low reflectivity surfaces (left). These artifacts are magnified when registering multiple frames to create a single model. However, many of these issues can be mitigated with intelligent filtering. This particular scene features a dusty path for the robot to traverse between rocks (right).

An example of the detail modeling capability of the SR4000 is shown in Figure 109. This scene shows a clear path in front of the robot between rocky areas, built up by registering many scans as the robot approaches. The left image illustrates the raw map (displayed as a height field) created from the keyframe ICP algorithm. Due to the high noise and low resolution of the sensor, point-based registration methods are prone to ill fitting. Utilizing intensity data (for correspondences across frames) improves the performance slightly. Additionally, there are holes in the map from erroneous returns that show up as rectangular pitting. However, with this approach, the modeled area is minimized in scale and stretching error.

It is possible to emphasize detail in these models by transforming or distorting the data locally. For example, planetary geologists might find the surfaces of rocks of interest. The location of rocks may be as important as accurate modeling of pitting on surfaces. The approach is illustrated in this dataset by making the assumption that the terrain exists as small deviations in a ground plane. The individual frames of the map can be projected onto this imaginary plane by rigid rotation after they have been first registered by ICP. A 2-dimensional feature match and alignment is then performed on the data in this plane. It is additionally possible to allow affine warping in alignment for a tight but distorted fit in successive steps. This process lines up all the features, but they may no longer be metrically accurate to assumption of underlying geometry. The right image in Figure 109 illustrates the results of this approach, which greatly improves the perception of rocks and S-shape of the path.

9.5.4 High Resolution Modeling of Cave Walls

The Lumenenhancement super-resolution process, described in Section 8.3.2, was demonstrated on images of a lava tube wall collected from Indian Tunnel (site described in Appendix A:

Indian Tunnel). A target patch of approximately 10ftx10ft was chosen for modeling. HDR photos of this patch were taken with DSLR cameras at (-45°, 0°, 45°) incidence angles from the center of the patch at 8ft distance. LIDAR data from a FARO X330 scanner was collected from the same locations as photography, with the low-resolution on-board camera providing a means of accurate registration between HDR and LIDAR data. Illumination was provided by banded LED illumination mounted on either side of the sensor. This configuration provides adequate illumination at range, and reduces cast shadows that reduce processed data quality.




Figure 110: Experimental setup for super-resolution wall modeling experiment.



Figure 111: Experimental setup for lava tube wall modeling. An intern adjusts a camera next to the LIDAR scanner in Indian Tunnel (left). The banded LED illumination provides a wide, diffuse area of illumination. Lateral symmetry reduces harsh cast shadows (right).

The HDR photos were first stitched into panoramas from each of the viewpoints. Each of these panoramas was then aligned to corresponding LIDAR scan data by projecting the colored point cloud to a 2D color view from a similar perspective. The colored point cloud produced by the LIDAR is low resolution and noisy, but suffices to align the HDR images, which are primarily used by the algorithm. Low-resolution color data is blended with the high-resolution HDR imagery in a contrast-weighted manner, and used to fill gaps in the HDR coverage. Color data from the X330 is histogram equalized with the HDR images for a smooth transition in the gap areas. Next, the three LIDAR scans are combined into a single point cloud by rigid registration. Registration provides global transformations for all of the photography as well.

The combined point model is projected to an orthographic range image at the resolution of the HDR data and gridded. This range image is approximately fronto-parallel with the plane of the wall. There are 9 channels of color data per pixel associated with this range image, one for each of the view/illumination directions.




Figure 112: Estimation of shape from appearance in the form of high-resolution surface normals. The depth data, which is directly measured by LIDAR scanning, is comparatively low resolution and used as a corrective observation.

Utilizing an iterative photometric stereo approach, which is a type of Shape from Shading (U. Wong 2012), the albedos and high-resolution surface normals of the scene are simultaneously estimated from the appearance (color) and interpolated depth data. Surface normals do not exhibit correct gross geometry and an attempt to integrate them directly will produce a distorted map. However, they more accurately represent the local texture of the scene than depth measurements, a fact which is exploited by the MRF fusion process. Surface normals are integrated into depth values by *wrapping* around the

base geometry measured by LIDAR. The final result is a high-resolution depth map, which is then re-transformed into Cartesian coordinates as a super-resolution mesh model.




Figure 113: Final result of super-resolution modeling of Indian Tunnel wall patch.

10 VISUALIZATION OF MISSION DATA

This section presents techniques for immersive visualization of pits and caves. This includes methods for model generation from robot data and methods for model display to users. The techniques were tested using data captured during field tests described in Section 9. Model generation techniques include point cloud rendering and mesh modeling. User display techniques include high contrast illumination for increased understanding, 3D printing of pit models, and anaglyph stereo visualization of models.

10.1 POINT RENDERING FOR CLOUD DISPLAYS

Point rendering methods for cloud displays were adapted from state-of-art and applied on field test data. These point-cloud methods produce immersive, accurate, continuous 3D models without the computing overhead of mesh modeling techniques. The discrete modeling method renders a watertight model from a point cloud through repeated copying, growing of points, and blurring. It then combines the images produced at each blurring increment to form a solid geometric model. This can be performed in real time on modern graphics processors. A detailed explanation of the approaches utilized and their relationship to prior work is summarized below, followed by images of a resulting model.




Figure 114: Point Rendering methods are used to produce an immersive, continuous 3D model from discrete point cloud models of the skylight floor and rubble pile. Unlike primitive “hole-filling” methods, point rendering produces a watertight, geometrically coherent model with correct occlusions. The process involves repeated copying, growing of points and blurring. The sequence of images is combined to form a solid geometric model of the scene. The whole process can be accomplished in real-time on modern graphics processors.

While robots natively generate and reason with point cloud data, this type of model is not effective for conveying information to humans. Point clouds are disorienting from a single view and do not convey an intuitive sense of things like “is this a wall?” or “where are the rocks?” The human viewer does not gain the immersive sense of being surrounded by 50m high walls and giant boulders. Therefore, mesh models (consisting of millions of triangles) are the common way of displaying 3D models. However, the process of “connecting the dots” to make solid triangles is very slow and introduces a different set of errors and artifacts.

If point cloud models are dense and uniform, however, there are excellent alternate approaches for visualizing the data. Points with color, radius and normal information (surfels) form the backbone of many of these alternate methods (Pfister, et al. 2000). As these properties are not sensed with LIDAR, they must be estimated from the data in conjunction with another sensor. With the advent of general pixel shader hardware and high-throughput measurement techniques, *point rendering* has become an alternative to meshing, especially when real-time structural updates are necessary.

The prior work of (U. Wong 2009) explored a middle ground between point clouds and meshes for display of planetary 3D models. These point rendering methods build on the natural point cloud data of robots, but make no attempt at intensive pre-processing or pre-filtering of the data as in meshing. Instead geometry is interpolated while rendering to the screen by filling holes between points and inferring occlusions with a series of image-space Gaussian masks. The key requirement is the availability of dense camera images in addition to the LIDAR scan. All of these steps can be performed in real-time on graphics processing units (GPUs).

Many results presented here are displayed using a hole-filling method similar to the multi-scale push-pull technique in (Grossman and Dally 1998). This display system is adapted to benefit from high-density clouds generated using super-resolution methods. Point clouds are rendered with push-pull interpolation in image space by generating a pyramid with successive radial blurring. The key idea is to factor the estimated point radii and surface normal (appearance) into the shape and weight of the blur mask. A min-depth check and kernel density estimator are used to resolve edge discontinuities and remove occluded background measurements. The utilization of texture in-painting for both color interpolation and depth reconstruction provides the viewer with graphical continuity as well as proper occlusions, which standard point displays lack.

This produces a “good enough” immersive model that can be generated in real-time and streamed from a robot (given appropriate bandwidth). Many artifacts of point clouds – like noise and uneven measurement density - are simply covered over by the process without need for explicit filtering. Given other parameters that can be estimated from image information (such as the illumination field), Wong showed that these techniques could approach state-of-art mesh models in fidelity and appearance, while being computationally efficient.

This research produced models based on data gathered in experiments at King’s Bowl Pit (See Appendix A: King’s Bowl Pit).




Figure 115: Overview of point cloud model of King's Bowl pit, rendered using the point-based rendering techniques described in this section




Figure 116: View from inside point cloud model of King's Bowl pit, rendered using the point-based rendering techniques described in this section

10.2 MESH MODELS

While the point-cloud rendering described in Section 10.1 is a great way to quickly visualize the results of robot modeling, for some applications a mesh model is required. Mesh rendering is better supported by general-purpose renderers than point cloud rendering, so it is likely a better choice if a 3D model is needed by an Education and Public Outreach representative to create educational animations, or if a 3D model from a reconnaissance mission is needed for those designing pit access robots. 3D printing, as described in Section 10.4, also requires mesh models. This section presents mesh modeling techniques utilized in this research along with images of models generated.

A mesh model was created of the Indian Tunnel skylight point cloud (see Section 10.1). To obtain the mesh model, visually apparent outlier points were manually removed from the point cloud and then Poisson reconstruction was applied (Kazhdan and Hoppe, Screened Poisson Surface Reconstruction 2013) (Kazhdan, Bolitho and Hoppe, Poisson Surface Reconstruction 2006). The mesh created using this method is shown in Figure 117. Poisson reconstruction tends to create watertight meshes, as seen in the domed cap created over the pit. The mesh was trimmed, and color data from the point cloud was transferred to the mesh (Cignoni, Corsini and Ranzuglia 2008). Figure 118 shows a close-up view of the colorized mesh of Indian Tunnel.




Figure 117: Side view of mesh model created with Poisson reconstruction. In an effort to create a watertight mesh, the reconstruction method has created a domed cap over the skylight pit.



Figure 118: Close-up view of the colorized mesh, showing skylight wall descending to rubble floor.

For the steep-walled skylight feature in Pluto’s Cave, a mesh model was created from the 3D point cloud generated during the experiment described in Section 9.5.1. The mesh was created utilizing Poisson surface reconstruction (smooth local area), which resulted in 500k triangles. The mesh model was then rendered using a Phong-based non-photorealistic highlighting to emphasize areas of high texture (see Figure 119).




Figure 119: Phong re-rendering of the lava tube mesh to emphasize surface texture and smoothed triangular discontinuities.

10.3 ENHANCING USER UNDERSTANDING WITH HIGH CONTRAST ILLUMINATION IN RENDERING

This research innovated an approach for photorealistic visualization of a model to be generated with angles and lighting conditions that differ from the conditions in which the imagery was captured. The technique takes advantage of material and lighting assumptions that are unique to planetary (non-terrestrial) situations. The result is a dramatic enhancement in user experience and information that can be understood quickly by viewing a rendered model. This section explains the technique and shows examples from models generated on field test data.




Figure 120: Principled visualizations of 3D data have great effect on the perception of geometry. Raw textured 3D data (left) appears flat and unconvincing. The appearance of the model is entirely the result of imaging and illumination conditions during capture. However, by estimating the fundamental components of albedo, reflectance and geometry, novel views can be rendered that enhance the user experience while maintaining photorealism (right).

Robot models are often visualized on two-dimensional surface displays (e.g. monitors, projectors, etc). This lossy projection of 3D geometry results in shape ambiguity and difficulty understanding the data. Even with 3D digital displays, like the Oculus Rift or volumetric displays, the model is often limited to key perspectives where pre-captured appearance information and shape are coherent. For example, if an image containing cast shadows is naively painted on a surface mesh, only the viewpoint and lighting that produced the image would generate an accurate rendering. Rendering the 3D model from any other perspective would introduce visual ambiguities and add “phantom” features.

Properties like perspective and appearance are the result of how data is collected. As robot sensing in planetary missions is a byproduct of other science objectives, the metric quality and photometric completeness of a model are rarely optimized for visualization. Given the importance of multi-perspective viewing to planetary science, there is a compelling need for post-processing techniques that enhance the way imperfect 3D models are displayed.

This research developed a novel approach to automatically select illumination that maximizes contrast while maintaining photorealism for arbitrary scenes. The idea is very simple: a search for the “best” image by rendering the space of possible illumination conditions for each viewpoint. The efficiency of

the approach (and quality of result) is drastically improved by applying a theory of illumination made possible by planetary assumptions. The method is practical for real time visualization.

Cameras record the appearance of a scene as a composition of surface geometry, material (reflectance functions) and illumination. Any single camera image encodes a combination of these parameters, but not in a way that allows extracting that unique combination from the image. Our approach is to use appearance assumptions and multiple observations to factor the scene into these fundamental building blocks. Once the image parameters are disentangled, it is possible to virtually recombine them with a custom illumination field that is optimized for the geometry and materials in the scene as well as the desired viewpoints. Even the geometry or materials can be virtually modified to enhance chosen viewing criteria like texture, though the result will no longer be photorealistic.

First, the effects of illumination during imaging – cast shadows, attached shadows and specularities – are removed from model data. Cast shadows are minimized by light source tracing on surface geometry and discarding of shadowed regions in computation of color data. Attached shadows are utilized in computation. Specularities are removed by projection into the SUV space. Second, we observe that the super-resolution method discussed in Section 8.3.3 disentangles material from geometry with either multi-image photometric stereo or single-image albedo estimation by assuming known illumination. For purely Lambertian scenes, albedo is the only unique attribute of a material and therefore albedo and material are equivalent. Thus, surface geometry is given in the form of a triangulated mesh from LIDAR and sensor fusion, with material estimates for each polygon in the form of an albedo map and assumed Lambertian reflectance. All contributions of illumination are hence removed from this *intrinsic model*, which is readily rendered to novel views.

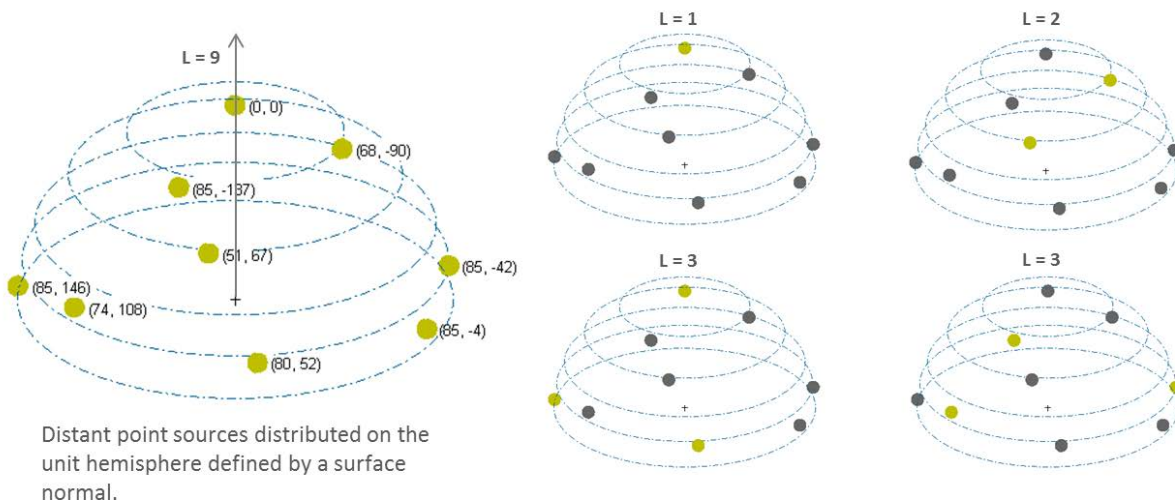


Figure 121: Key point illumination sources are identified that span the space of convex surfaces (left), which is a nine-dimensional space. A linear combination of these sources is sufficient to approximate appearance under all lighting conditions. Some illumination conditions are sampled from this set to render images of the scene for comparison (right).

The problem of presenting the “best” illumination for a scene can be considered in a brute force manner by rendering all possible lighting conditions and selecting the image that maximizes some quality metric. As there exists an infinite number of lighting conditions (a continuous domain), this approach is generally infeasible. However, by relaxing requirement for global optimum, it is possible to select the best image from a sample set of renders of arbitrary size. In order for this result to have any significance, these rendered samples should be key images that span the space of surface appearance under illumination³.

The groundbreaking work of (Lee, Ho and Kriegman 2001) identified two results that apply here: (1) there exists a nine-dimensional subspace in which all illumination conditions of simple Lambertian scenes can be represented and (2) this subspace can be approximated from nine real images of the scene under point illumination. The nine-dimensional subspace arises from spherical harmonics theory, where the nine bases of degree $L < 3$ capture 99% of the reflectance energy. A linear combination of these spherical harmonics is sufficient to approximate all illumination at a surface. As harmonic images are not real (they contain negative illumination values), the great result of the work was to identify physically plausible conditions that can be used to approximate this space. Amazingly, it turns out that images collected under point illumination with certain fixed positions are sufficient – conditions that are conducive to imaging or rendering. For example, rendering a scene with a linear or area light source is computationally intensive. With this theory, the scene could be instead rendered several times with only well-placed point sources. The final appearance with complex illumination effects is composited from the point source images in a much more efficient linear manner.




Figure 122: Rendering models with specific illumination and view angles maximizes understanding of surface geometry. These parameters are dependent on the reflectance of materials and geometry contained in the scene. Two perspectives of the Indian Tunnel wall model are shown with optimal illumination determined by our algorithm.

Though Lee primarily establishes theory for use of this nine-point subspace for object recognition, the effectiveness of point sources for modeling of novel illumination effects is shown in similar work by (Koudelka, et al. 2001). Koudelka chooses a larger number of images to enhance fidelity of interpolation. For the objective here of simply generating an image and scoring it, the accuracy of interpolated illumination is not as important as sampling orthogonal key images. The intuition is that these “characteristic” images capture the most important illumination effects for compositing, such as cast shadows from opposing angles. Thus, the process here greatly simplifies many steps for efficiency.

³ known as the *illumination cone* in technical literature

First, the user selects a viewpoint to display the three dimensional model. The nine key images are rendered from this view using global illumination and Lambertian materials estimated during the model creation process. These nine images are then used to assemble a fixed set of composited images by the following equation:

$$I_{composite} = \sum_{k \in N} \alpha_k I_k \quad (0.13)$$

where α_k are weights and I_k is an image from the set (N) of 9 key images. This is done by choosing random subsets of the key images (utilizing 3- and 4-image sets) and applying a nonlinear optimization to determine a combination of weights α_k that would result in the highest-contrast composite. This highest-contrast composite is displayed to the user. This process is very fast because it is based on weighted addition of images.




Figure 123: An image pyramid for number of surface samples. The metric for contrast here calculates a score at each scale, and the final score is a weighted sum of all the scale scores. This places importance on both clarity of small details as well as regional contrast.

The contrast metric has profound effect on the appearance of the final result. We chose to build our metric around one of the canonical definitions: standard deviation of pixel values in a neighborhood. This metric is calculated multiple times for all levels (L) in a subsampling pyramid such that every pixel has scores $\sigma_1 \dots \sigma_L$. This emphasizes both local (details) and regional (big picture) contrast equally (see Figure 123). Each pixel receives a contrast score from the weighted sum (ω_i) of the level scores, and the entire image receives a score that is a weighted sum of all pixels:

$$Contrast = \sum_{k \in Pixels} \beta_i \sum_{i \in L} \omega_{k,i} \sigma_{k,i} \quad (0.13)$$

where β_i is a regional weighting function that might place greater emphasis on high contrast at the center of the image compared to the edges.

10.4 3D PRINTING

This research innovated approaches for enhancing user experience of models of terrain through physical manifestation. Physically interacting with models can provide unique tactile and visual understanding of an environment sensed by a robot on another planet, far beyond what is possible through computer display and rendering alone. This has great advantages both for scientists and for public engagement. This section describes the techniques and models generated in this research.



Figure 124: 3D printed model of a wall segment in an analog pit

A 3D printed model was generated of a mine pit high wall following model enhancement through Lumenhancement as described in Section 8.3.2. The surface model of the pit was extruded in a CAD program to produce a 3D model with flat sides and back to enable 3D printing. The 3D printing of this robotically created geological model provides a tactile artifact that scientists can use to better understand the modeled surface as shown in Figure 124. This model was printed in a sand-based material with a colorized top layer.

This research also demonstrated 3D printing of expansive, multi-view pit “shell” models. This presents an additional challenge because robotic-generated models of complex scenes with features like concavity and extreme vertical relief (e.g., entire skylights) are often incomplete. Unlike human-made CAD models that print without problems, robots are limited in their access and viewpoints while moving sensors through a scene. Often, there is not enough geometric information from multiple perspectives that a coherent volume can be created without glaring artifacts like holes, sharp discontinuities and phantom walls along edges. The ability to print the mine pit (mapped by Tyrobot and survey scanning) as a shell model demonstrates the efficacy of modeling techniques developed in this project (as described in Section 10.2).



Figure 125: Computer rendering of mine pit 3D model created by robots (left). Top view of 3D printed shell model looking into the pit (center) and reverse view (right).

Figure 125 illustrates the shell model that was created. The model is approximately 25 x 15 x 5cm representing a pit of over 100m along the long dimension. It was made from a mesh of 1 million polygons with details as small as 10cm visible. The interior walls are an accurate scaling of the actual site, while the outer walls are minimally distorted in some locations to accommodate the thickness of wall that could be printed. The model was made of a strong plastic material that supports only a single color. This material is much stronger than the sandy, full-color material used to print the previous wall model, which was necessary for strong, thin walls in a complex shape. A hole exists in the center of the “bowl”

due to lack of LIDAR readings from a puddle of water. This could have easily been filled in for aesthetic reasons, but the research team wanted the model to be a 100% representation of the data that was collected.

10.5 ANAGLYPH STEREO

Anaglyph stereo enables people to perceive 3D shape from a single image using stereo glasses with different color filters (in Figure 126, the filters needed for 3D viewing are red and cyan). The filters on the glasses transmit certain colors more strongly than others such that each eye sees a different version of the image, as if the 3D object was viewed from two slightly different perspectives. Although the film industry has gone with other technologies for 3D movies, such as projecting polarized light or using active shutter glasses, anaglyph stereo is a method that can work in printed materials and with very inexpensive viewing hardware such as cardboard glasses with red and blue gel lenses.

This research utilize standard techniques for generating anaglyphs from rendered models to provide accessible 3D visualization of models and of 3D videos captured by robots during field experiments.




Figure 126: Anaglyph stereo representation of a high wall model from a mine pit field test.

11 ROBOT DEVELOPMENT

This research conceived, developed, and demonstrated two robots for exploration and modeling of planetary pits and caves. A Tyrolean robot (*Tyrobot*) and scree slope descender (*Krawler*) were selected based on the feasibility study presented in Section 6.2.6. Mine locations utilized for testing *Tyrobot* and *Krawler* are described in Appendix A: Mine Pits.

Tyrobot can access pit floors regardless of taxonomy, minimizing the impact of pit aperture on communications and power supply to sensors and dependent robots on the pit floor. Its dual-point anchor system significantly reduces sensor oscillation during descent and improves motion control for camera perspectives over the single-point anchor *Ferret* borehole sensor. While the configuration introduces an additional robot to the cave exploration mission, that robot remains useful and acts as the Pit Infrastructure Link.

Krawler utilizes a single robot for both pit descent and subsurface exploration, reducing mission complexity. A low center of gravity and conformable suspension enables descent of dramatic, steep slopes. *Krawler* was deployed by *Tyrobot* in field experiments.

11.1 TYROBOT

Tyrobot traverses along a Tyrolean line, lowering and raising a payload on a winch. It can lower a sensing payload to view cave walls and peer into subsurface caves beyond line-of-sight from the surface. It can act as an elevator, emplacing subsurface robots and serving as their communications relay and recharging station.

Section 11.1.1 describes *Tyrobot* configuration development and the prototypes constructed and tested to validate the approach. The configuration was tested to validate: the robot’s drive system (Section 11.1.2); the effects of tension in the Tyrolean line (Section 11.1.3); anchoring the ends of the Tyrolean line (Section 11.1.4); and recharging of subsurface robots from *Tyrobot* (Section 11.1.5). Section 11.1.6 describes the configuration of a sensing payload integrated with *Tyrobot* for field tests. Field tests results of *Tyrobot*’s modeling performance are covered in Section 9.4. Field test results demonstrating *Tyrobot*’s mobility performance are covered in in Section 9.4.1.




Figure 127: *Tyrobot* modeling a pit in a terrestrial analog field test

11.1.1 Configuration & Robot Development

A preliminary prototype was constructed to evaluate modeling from a Tyrolean line. This first prototype could raise and lower a payload, but traverse along the Tyrolean line was accomplished manually by pulling traverse ropes from the pit rim. A GoPro camera was attached to the carriage for modeling, along with a survey prism for ground-truth payload motion by tracking from a survey instrument. This prototype was tested at a pit mine.

The next-generation prototype of Tyrobot added capability to the trolley, including (see Figure 130). computing to control actions, monitor sensors and log data; and motorized actuators for self-powered traverse. The top motor provides lateral actuation along the Tyrolean line while the bottom motor raises and lowers the sensor head. Encoders on both motors provide feedback along those axes. A cantilevered pulley and weight system keep the tether taut and prevent it from dangling into the sensor’s field of view. 802.11g wireless communications enables operators to monitor the robot remotely.




Figure 128: Diagram of first prototype Tyrolean modeler. A camera and reflective prism for survey tracking are rigidly mounted on a weighted frame. The frame is lowered from the main Tyrolean line by a two-cable winch and pulley system. The two lines provide lateral stability.



Figure 129: The first Tyrobot prototype included motorized raise-lower for a payload, but Tyrolean traverse was done manually, by pulling on cables from the pit rim. Command of the raise-lower mechanism happened through a wire (orange cable in image) strung along the Tyrolean line.




Figure 130: In the second Tyrobot prototype, the trolley provides positioning for the modeling sensors vertically and along the Tyrolean line.

A prototype attachment system for Tyrobot was fabricated to investigate rapid deployment by securing to a cable that is already strung. This device is placed on an elevated cable by a robot that collides with the line at a shallow angle. Two pivot arms, indicated by the red arrows below, rotate into place to create a wrap angle of approximately 180 degrees around the drive pulley. This gives the system the necessary friction to traverse the cable under a wide range of cable tensions. Cable slip was minimal in all cases; predicted travel based on encoder distance was within 5% of actual travel. While this device was not implemented in the final configuration of Tyrobot, it was demonstrated as a potential add-on for relevant mission concept of operations.




Figure 131: Prototype Tyrolean trolley that investigated a two-arm clamping mechanism that secures to the cable. This mechanism also increases wrap angle on the drive pulley, enabling steeper climbs.




Figure 132: CAD Design of Tyrobot mechanism. Instead of attaching the winch cables to the payload directly, the cables attach to a platform that stores the electronics and provides a more stable base. When lowered, communications to the top trolley are wireless, removing the need for any tether management.

Following these conceptual prototypes and tests, the final Tyrobot configuration was produced as shown in Figure 132. Tyrobot uses a friction-based system to traverse the cable, with both pulleys driven to maximize normal force used to generate friction. An idler pulley, shown in Figure 133, rises linearly on a screw system to increase the wrap angle around the drive pulleys and the normal force between the cable and the pulleys. An electromagnetic brake holds the idler fixed rather than relying on active motor torque. A spooling system with a level wind, similar to that of a fishing rod, releases a line to lower the platform and payload up to 50m below the Tyrolean line. Computing and most of the electronics are stored in the lowered platform, which controls the traverse and lowering actuators in the independently-powered top section. Inter-module communication utilizes wireless Ethernet to eliminate data tethers.




Figure 133: Cross section view of Tyrobot trolley illustrating internal motors, batteries and mechanism.




Figure 134: Latest Tyrobot being tested in the Astrobotic office

Tyrobot was designed to support a variety of modes, physical parameters, and payloads. For example, many internal components are fully customizable and can be shifted to different mount points for balancing. Drive pulley sizes, offsets, widths, and placement can be adjusted. The height of the central tensioner pulley can be raised or lowered to adjust the wrap angle on a taut tight rope line. This customizability supports adjustments after fabrication, since optimal system parameters will likely be determined only through meticulous field-testing. The system can be tuned during or after primary assembly without costly redevelopment.




Figure 135: Tyrobot during integration and preliminary testing in a high bay. A mapping sensor payload is mounted to the payload dock and lowered in an integration test in the image on the left.

11.1.2 Drive System Development & Characterization

Tyrobot's drive system and the impact of line tension and wrap angle on friction between drive pulleys and the rope were evaluated through a series of tests. Experimentation sought to determine the maximum angle that could be climbed, what tension had to be maintained in the tigtrope, and the impact of the geometry and material of the drive pulleys on frictional drive force and resilience against slipping.

A mock up of the drive system was developed in parallel with Tyrobot fabrication. It consisted of immobile drive pulleys, a tensioner pulley with adjustable height positions, and a load-bearing frame for payload weight (see Figure 136). This mock up was deployed on a test line tensioned with known weights and anchored to static end points in the lab. The assumption being tested is that the immobile drive pulleys and rope form a friction interface equivalent to driven pulleys under maximum motor torque. The torque of the drive motors used is sufficiently large that this assumption is valid in most circumstances. In operating conditions, Tyrobot will slip if friction is unable to counteract a strong lateral force on the line caused by events such as climbing a steep slope with heavy payload or instantaneous slack in the line due to stretching or breakage. By pulling laterally on the non-rolling drive system, slip

conditions were reproduced and measured (see Figure 137). A digital strain gauge was utilized to measure the pull force.




Figure 136: Experimental setup of friction pull tests. A student provides the lateral force needed to surmount the static friction of the system.




Figure 137: Friction pull tests are used to determine drive system effectiveness in traversing the tightrope. This consists of mocking up Tyrobot's drive system, but with immobile driven pulleys, and sliding it against the tightrope. Parameters that significantly change the system's coefficient of friction are used to inform further design iterations. Shown here are various drive pulley cross sections that were tested. The V-shaped pulley exhibited the best drive traction.

The aggregate effect of the aforementioned parameters such as pulley shape, wrap angle, line material, etc. can be approximated as a single static coefficient of friction (COF), which determines the holding force for a given payload weight and consequently the maximum slope that can be climbed. The methodology is simple: parameters of the drive mechanism are modulated, the coefficient of friction is determined by a pull test, and the experiment is repeated. The combination of parameters that maximizes the COF while remaining practical to other concerns is chosen.

The program conducted friction pull tests to optimize for pulley grooves, material, and placement. Tests were also conducted to select Tyrolean line material and characterize acceptable line tensions. A V-shaped pulley cross-section was chosen and fabricated from a hard polyurethane material. A flexible, multi-strand, braided cable made from synthetic fiber was selected as the Tyrolean line over steel cable.




Figure 138: Static friction forces are determined over the design parameter space - in this case, pulley shape and idler height. This graph informs that the V-shape design is more beneficial for drive traction (left). A strain gauge sensor continuously monitors the pull force. By increasing the lateral pull force over time, it is possible to determine the point at which the system slips as a break in the curve (right). A specific friction coefficient can be back-solved by knowing the payload weight and the maximum load before slip.

11.1.3 Tyrolean Line Tension Testing – optimal line tension

The tension of the Tyrolean cable affects the sag of the line and consequently the viewpoints for sensing. A perfectly horizontal, taut cable is ideal for capturing data from floor to rim and generating unbiased level slices. However, this scenario may require a complex mechanism and prohibitive amount of tension to hold. Conversely, a sagging line requires very little tension, but would negatively affect science modeling.

A test to correlate cable tension with expected sag and modeling viewpoints varied the tension by attaching one end of the Tyrolean line to a moveable anchor. The anchor was then moved to stretch the line and create a tension force. The resulting droop was measured by capturing rectilinear imagery from the across the pit. This camera features very low distortion optics for optimum measurement of angles.

Figure 139 shows several line tensions superimposed on a single image. While the line remains straight throughout the range of tensions, the payload droops significantly for low tensions.




Figure 139: Tension in the Tyrolean cable affects the maximum height attainable by the sensor payload at each point on the line. In low-tension cables, the maximum height at the midpoint is significantly lower than for high tension ones.

11.1.4 Cable Launching and Anchoring

A promising cable deployment strategy to deploy a light guide cable across a pit or void as shown in Figure 140 was demonstrated in the field. The launch mechanism consists of a pan-and-tilt platform that aims a rocket to a desired heading and (ideal) distance. The rocket is tethered to a cable, which unspools from a modified fishing reel. The field test used hobby-class rockets with slight modifications to propel the rocket across the pit. This system demonstrated unspooling and deploying of a cable over a 34m span. For comparison, medium sized pits on the Moon are about 50m in diameter, while large pits may be up to 100m in diameter. Greater distance is possible with larger hobby rockets; however, it is important to note that efficiency vastly improves in lunar gravity and with lack of atmospheric effects.



Figure 140: Rocket launch platform for tether deployment (left). This system might be mounted to a lander that is positioned near the edge of the skylight. Tests in earth gravity and atmospheric demonstrated deployed cable distance of 34m (center) with a hobby rocket; efficiency and distance only improve on the Moon. A fishing line tether is attached to the rocket that is used to pull the Tyrolean cable across (right).

In a lunar mission scenario, the lander-mounted launch mechanism deploys the guide cable across the pit. A multipurpose rover carrying one end of the Tyrolean cable travels to the other side of the pit by circumnavigation, locates the launched tether, and attaches it to the Tyrolean cable. The spooling system at the tether launch site then pulls in both the tether and cable to span the skylight, with the tightrape anchored using the deadweight of the lander on one side of the skylight and the rover on the opposite side. Scenarios where the rover might bury the anchor and be freed to explore are also envisioned.

Other methods of cable deployment were considered, including mechanical launching systems like catapults and fluid dynamic launching from a pneumatic cannon. Chemical propulsion by rocket was identified as the most promising.

11.1.5 Robotic Recharge for Floor Exploration

Tyrobot can also serve as a recharge link for subsurface exploration robots. Robust recharging is necessary for long-term robotic exploration of lunar pits and skylights with steep, nearly vertical walls where direct solar charging is not viable. This work posits that solar energy can be collected at the surface and carried to the pit floor via a power cable that docks with the rover to recharge its batteries. When integrated with the Tyrobot system, this power cable could be moved to multiple locations on the pit floor to deliver power to one or more floor explorers.

Robotic recharging on the Moon is complicated by the unpredictable terrain and dusty conditions, which could cause docking misalignment or foul electrical contacts. One possible solution is a mechanical device that mitigates both severe connector misalignment and dust interference to ensure reliable power transfer. The mechanism enables a rover at the bottom of a lunar pit to dock with and undock from the vertical, hanging tether carrying electrical power from the surface source. Power transfer between tether and rover can be facilitated by induction (primary mode) or conduction (redundant mode).

A prototype robotic recharging device was designed, built, and tested, and its performance was evaluated based on the mechanism's tolerance to imperfect approach angles and resistance to lunar regolith simulant accumulation. The device, shown in Figure 141, is mounted to a rover as it approaches the hanging power tether and after achieving a successful electrical connection.




Figure 141: The prototype recharging device is mounted to a rover as it approaches the hanging power tether during testing (left). Experimentation demonstrated successful mating under a variety of slope and terrain conditions that might cause misalignment. A successful electrical connection is indicated by a red LED for conduction and a green LED for induction (right).

The prototyped robotic recharging device was tested a total of 64 times, including 16 unique orientations, and established a successful connection (i.e., one capable of transferring electrical power to the rover) in 40 of those trials, for an overall success rate of 62.5%. The device was 69% reliable in clean conditions and 33% percent reliable when a significant amount of JSC-1A lunar simulant was applied to the device. The application of simulant was designed to test limits of accumulation that might result from material falls due to pit wall disturbance, an amount believed to be far beyond reasonable lunar accumulation. Observed failure modes were few in number, well-defined, and predictable. When avoiding 180 degree misalignment, steep pitch-downs, and unrealistic levels of regolith accumulation, all 31 of the remaining trials were successful (100%).

The results of this study indicate that the concept of using a passive alignment mechanism for conductive and inductive robot recharging works in principle but could be sensitive to regolith contamination. Further testing can better characterize what level of regolith accumulation is permissible. Interestingly, the conduction-induction redundancy did not significantly impact the overall reliability of the system. In all trials when induction succeeded, so did conduction, and only once did induction fail when conduction succeeded. A viable alternate strategy is to include only one mode of power transfer.

For example, using state-of-art inductive charging technology eliminates all of the observed failure modes and achieves reasonable power efficiency.

11.1.6 Sensing Payload Configuration

Figure 142 shows the sensor payload carried by Tyrobot during field tests. It features a fisheye camera (182 degree field-of-view), a planar scanning LIDAR (270 degree swept plane) and a tracking prism for validation. Internally, an inertial measurement unit tracks the pose of the sensor payload. The LIDAR slices a single plane of measurements, which is then rotated 360 degrees to double-sweep a hemisphere. The result is a bowl-shaped point cloud that covers up to 45 degrees above the horizon. The camera is coaxial to one side of the laser window and continuously records video.




Figure 142: Pit modeling sensor payload carried by a Tyrobot, adapted from a terrestrial borehole robot.

The dimensions of the mine pit analog used for this field test of Tyrobot closely resemble that of an actual lunar pit. Thus Tyrobot’s ability to image or scan desired details – at physically attainable positions within the pit – indicates reasonable expectations for science modeling quality. a significant finding of this work is that the 870nm infrared LIDAR utilized for these tests performed reasonably well only at distances much shorter than the reported 30m range. Much of the walls of the pit (and central rubble pile) at the test site consist of coal, which is known to have a visible albedo similar to lunar regolith. This finding informs the type and power of sensing to be carried by a pit modeling robot. It might mean that the robot must rely much more on vision and advantageous sunlight for modeling. Operationally, a Tyrolean robot could be expected to provide a superior range of up-close perspectives to compensate for sensor constraints.

11.2 KRAWLER

The Krawler robot was developed for both pit ingress (for pits with gentle ramps) and subsurface exploration.

Krawler can surmount extreme obstacles relative to its size using only four wheels. Its light, flexible chassis and elastic suspension keep all four wheels in constant contact with the ground, even over extreme terrain including large steps and gaps. Its high ground clearance enables it to surmount obstacles larger than its wheel diameter, a feature unique to rock crawlers. Krawler maintains a low center of mass due to its lightweight aluminum chassis construction and its use of pendulum-weighted wheels, each of which maintains a center of mass well below its axis of rotation, even as it rolls. The net effect dramatically lowers the vehicle's total center of mass, preventing tip-over on extreme slopes and steps and further extending its ability to traverse extreme terrain. In addition to weighted pendulums, each wheel utilizes a beadlock that anchors the tire to the wheel hub without relying on high tire pressure. This enables the use of lower tire pressure that results in a larger contact patch between the tire tread and the ground. The increase in ground contact area and ground compliance significantly reduces slip by improving traction.

Krawler's steering configuration further improves its mobility in confined space. Krawler incorporates steering mechanisms on both front and rear axle. This dual Ackermann steering system enables tighter turn radius than vehicles with only one set of steerable wheels, improving agility in confined spaces where minimal turning radius is essential.

When negotiating slopes in loose rubble terrain, choosing the path with the shallowest slope can greatly reduce the requirement on drive motor torque. However, a traditional steering configuration, where the front wheels steer to a certain direction while the rear wheel direction remains fixed, will increase the load on the drive motor because driving forward while changing heading direction requires far more torque than purely propelling forward. Crab steering mode effectively mitigates this issue by steering both front and rear wheels toward the same direction, enabling the rover to navigate to the desired gentle slope path without the need to overcome the extra resistance caused by steering. Since the heading directions of the four wheels are always parallel, the drive motors only need to deal with the resistance from the friction and slope.

Krawler also employs dual drive motors on both front and rear axle. This redundant configuration enables an electrical switch between 4WD/2WD drive modes and further increases traction in 4WD mode by introducing extra motor power. This configuration also increases safety, as the failure of one motor will not directly cause loss of traction.

The successful subsurface explorer must be able to traverse extreme terrain and make long forays into caves using only battery power. Section 11.2.1 examines the Krawler's ability to ascend and descend slopes with different rock sizes, even at reduced power. Section 11.2.2 classifies rock types from on-board sensor data to facilitate autonomous tailoring of the driving approach to the terrain type. Section 11.2.3 explores the energetics of long forays into caves. A field test demonstrating Krawler's descent down a scree slope into an analog pit is discussed in Section 9.4.1. Field test results for modeling from a Krawler in an analog lava tube cave are discussed in Sections 9.5.2 and 9.5.3.

11.2.1 Krawler Mobility Configuration: Evaluating Slope Traverse

Testing was conducted to analyze performance of the Krawler mobility system on slopes consisting of different materials. This testing occurred at an aggregate supply company that featured large rock piles with different sizes and shapes of material. Rocks ranged from pea-sized to softball-sized and included both smooth river rock and jagged crushed limestone. Slopes were naturally formed by each material’s angle of repose and ranged from 30-65 degrees. The variety of terrain types supported performance comparison across a spectrum of mobility challenges.

This experiment called for human-controlled driving of the crawler up and down the different slopes utilizing two patterns of driving: a linear trajectory (common in direct waypoint following in robotics) and swerving s-shaped curves. To record data, the rover was outfitted with a Beaglebone Black microprocessor, a Pololu Minimu-9 v2 inertial measurement unit, and two GoPro cameras in stereo configuration. A tracking survey instrument recorded the position of the vehicle. A high-resolution LIDAR scan of the slopes before testing was used to estimate average slope angle and provide ground truth geometry such as estimation of roughness and rock size (macroscopic geometry of the piles were modified as a result of testing). Rover power levels (10, 30, 50%) were also modulated during testing to expose the effect of energetics on climb performance.



Figure 143: Rock Crawler mobility platform ascending and descending piles of different material types (pea gravel, 1-inch limestone, pea river rock, 3-inch river rock).

11.2.1.1 Krawler Configuration: Conclusions from Field Testing

For small diameter pebbles, the crawler was able to ascend and descend the piles using either trajectory strategy even at 30% power, though ascending was slow. The high degree of material deformation around the tires did work to buttress the vehicle against predicted rolls. Dynamics came into play for larger rocks, as the crawler would “bounce” during descent, significantly altering the intended heading and trajectory. This has repercussion when factoring in the lower gravity of the Moon and Mars. For the largest smooth rocks, slip also had a substantial effect on rover performance, particularly for position holding and ascents.

11.2.2 Krawler Mobility Control: Learning Terrain Types for Active Control Compensation

This section covers the challenges inherent in traversing unknown scree slopes and bouldered terrain. It presents an approach to classifying terrain through machine learning based on vibration data. This terrain classification supports active adjustment of control strategies based on the material on which the Krawler is driving. This section addresses basis, experimentation, classification approaches, and results.

While the Krawler platform was selected for its terrainability, driving rock crawler platforms is an art. Hobbyists train for extended durations in order to drive these robots safely and efficiently.

Counterintuitive toe-in or zigzag maneuvers are often required to negotiate the most challenging terrain. This type of maneuvering is contrary to the state of practice in planetary robotics, which favors simple dynamics so that controls can be defined by direct actions like arc following and turn-to/go-to. These simple maneuvers would unacceptably limit Krawler terrainability.

Knowledge of the terrain would enable the rover to apply different driving modes or control compensation according to the actual environment. Modern cars and trucks use 2WD, 4WD, and even more detailed drive modes to negotiate asphalt paved roads, sand terrain, grass, snow, and other extreme terrain types. For mobile robots working in planetary environments, rocks that range from miniscule to almost wheel scale represent an effective generalization of the expected environment. Even materials such as loose regolith and sand can be viewed as small rocks with miniscule diameter. Applying appropriate drive modes on different rock sizes would help improve robot performance in mobility, agility, maneuverability, and/or energy efficiency. The first step in achieving control compensation is being able to recognize and classify the actual terrain that the rover is negotiating.

Inspired by human sensing modality, vision is the most direct way to recognize rock size. However, computer vision is too sensitive to trivial factors, such as environmental lighting, shade of the rover, etc. Vibration is a more robust, reliable and case invariant representative signature of rock sizes. More importantly, vibration is a direct estimation of mechanical property. Unlike vision, which measures visual features of the rock type and infers the correspondence to mechanical effect, vibration measures the mechanical effect directly. The amplitude, period, and other temporal and frequency features contain useful information about the rocks that display these characteristics in interaction with the rover. A field test was conducted to collect data for Krawler driving on different material types. Vibration data was recorded by an on-board IMU. Machine learning methods were then applied to learn terrain type classifications.

The field test was done on 4 different rock types with increasing rock size.

The 4 different rock types were:

- #3B River Gravel
- #4B River Gravel
- Pea Gravel
- #1B Limestone

The #3B and #4B river gravel are almost same size, but the #3B has sharp edges while the #4B's contour is smooth. 80,000 data points were recorded with an interval of ~2ms. Each data point is 6-dimensional, including 3 linear accelerations and 3 angular velocities. 40,000 points were classified as training data and the rest as test data. A moving window with length of 100 data points is used to assure simultaneously sufficient valuable information in one data point (600 dimensional) and sufficient data points (79,900).

11.2.2.1 Feature Extraction

Eight features were extracted for each of the training and test sets, four of which were in the time domain with the remaining four in the frequency domain. The time domain features are average value, zero crossing rate, short time energy, and energy entropy. The frequency domain features are standard deviation, skewness, kurtosis, and spectral roll off. Input to the classifier is 48-dimensional.

Zero crossing rate is the number of times per second that the signal crosses the zero axis. Short time energy is the sum of the squares of the amplitudes. Energy Entropy refers to a measure of the abrupt changes in energy present in the signal. A Fast Fourier Transform (FFT) is applied to obtain the features in the frequency domain. The spectral standard deviation, skewness, kurtosis, and spectral roll off are calculated based on the spectral amplitude.

11.2.2.2 *Classification*

Three classifiers were tested on this data: Artificial Neural Network (ANN) with different network parameters, Principle Component Analysis (PCA) combined with k-nearest-neighbor (kNN), and Support Vector Machine (SVM) with different kernel functions with different kernel parameters. For ANN, NetLab is used while SVM classification is based on LIBSVM.

For ANN, a hidden node number from 100 to 2000 is used, with an interval of 100, and train epochs from 10 to 100, with an interval of 10. The best result is 67.25% with 2000 hidden nodes and 50 epochs.

For PCA combined with kNN, the first 20 principal components were chosen and k was set to be equal to 1. The accuracy is 49%, which demonstrates that PCA plus kNN is not an effective classifier for the terrain recognition task.

For SVM, a linear, quadratic, polynomial, Gaussian Radial Basis Function (RBF), and Multilayer Perceptron (MLP) kernels were implemented and parameters for each kernel function are adjusted. The best accuracy, 70.25%, was achieved by polynomial kernel with the polynomial order to be 3. In order to achieve the best accuracy, not all the features described above are included in the classifier input. Different combinations of the features above were tried. The combination that yields the best accuracy is:

- All channels of Zero Crossing Rate
- All channels of Short Time Energy
- The 5th channel of Energy Entropy
- All channels of Average Value
- The 1st and 4th channel of Standard Deviation
- The 1st and 3rd channel of Skewness
- The 3rd channel of Kurtosis




Figure 144: Classification Result: green circles indicate the correctly classified data points while red crosses are the wrongly classified ones.

11.2.2.3 Discussion of Results from Field Testing

The best classification accuracy achieved was 70.25%, using Support Vector Machine classification based on LIBSVM with a polynomial kernel of order 3. The pea gravel and pea river rock data sets contained rocks of the same size with different contours, sharp versus smooth. By merging these two data sets into a single data set of rocks of the same size, the accuracy of classifying rocks of different sizes increased to 81%.

11.2.3 Krawler Power Considerations: Energetics for Long Forays

Maximum robot traverse range is a critical parameter for planetary exploration. Short range restricts the robot’s functionality and reduces the amount of data collected and mission coverage. In the absence of an isotope reactor or constant solar power, planetary exploration robots must rely on the limited energy carried in batteries. While recharging is possible for subsurface explorers (as shown in Section 11.1.5), range remains limited due to the need to return to a charging station. Battery-powered robots without replenishment provide the simplicity to focus on a constant energy amount while examining energy utilization. In building an energetic model for mobile robots for long forays, this work assumes that the robot is battery-operated.

Generally, in long forays, rover energy is expended for the following two purposes:

- Ancillary power for robotics functions like computing, sensing, communication, and payload that scales with duration of operations.
- Energy for motion, which predominantly scales with distance driven.

Since driving energy is primarily related to rover mass, gravity, distance traveled and terrain resistance, total energy for driving is mostly independent of speed and mission time. By comparison, the energy consumption for robotic functions such as sensing, computing, and communication is considerable, whether driving or sitting. While less concerning when recharge is possible from solar or radioisotope

source, this energy sink is paramount when operating from only single discharge from battery. In traditional exploration, the mission energy for sensing, computing, communication and payload far exceeds driving energy. Because most time is spent sitting or creeping, it traditionally requires days, weeks, months, or years to drive kilometers. Missions of this tempo require energies of kilowatt-hours because ancillary power is drained over such a long duration.

The total energy expended during whole mission time could be quantified as:

$$\begin{aligned} \text{Total Energy Consumption} &= \text{Robotic Consumption} + \text{Mobility Consumption} \\ &= \text{Robotic Power} \times \text{Mission Time} + \text{Constant} \times \text{Rover mass} \times \text{Distance} \end{aligned}$$

Mobility consumption includes all the energy needed to keep robot in motion, such as drive motor, steering motor, and their related energy losses. This comprises only a small fraction of total energetic consumption. Given a rover on a certain terrain, energy required for mobility is approximately constant for a certain range, with slight variation due to different steering activity. Robotic consumption is the part of energy used for the assigned tasks during traverse, such as computing, sensing, communication, etc. This energy is expended at all times whether moving or sitting, and energy consumption of robotic functions is mainly determined by total mission time (or rover speed, if given a certain range). Some rovers don't move all the time during traverse. They have to stop intermittently for navigation, planning, teleoperation, or data collection. Their "driving duty cycle" refers to the percentage of time that the rover is actually driving during payload operation in the total traverse time. In the time period when the rover stops and doesn't have mobility consumption, the robotic functions still continuously consume energy. This increases the ratio of robotic to mobility consumption.

Assuming a 10kg rover is to achieve a 2km traverse on Earth ($g=9.81\text{m/s}^2$), the terrain resistance coefficient (ratio of terrain resistance to rover weight) is 0.15. The direct propulsion energy to overcome terrain resistance and drive forward consists only 30% of total mobility energy, considering motor loss, internal friction, etc. The robotic power is assumed to be 30 watts. The relationship of required energy, driving duty cycle, rover velocity and achievable traverse range in accomplishing mission objectives can be seen in Figure 145 and Figure 146.

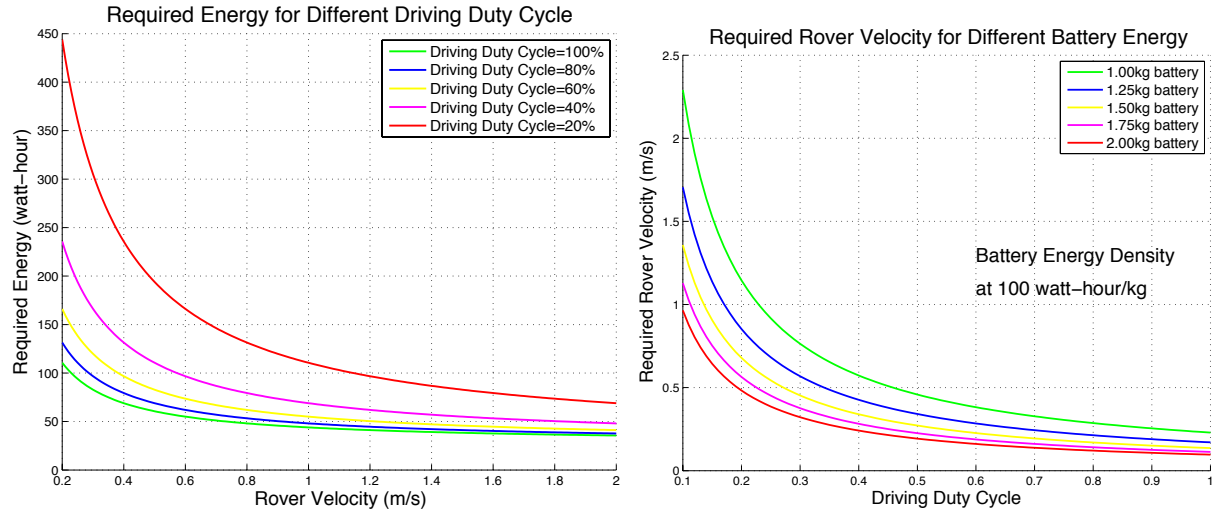


Figure 145: Left: Required total battery energy for a 10kg rover to achieve a 2km traverse for different driving duty cycles with respect to different rover velocities. Right: Required rover velocity to achieve a 2km traverse for different battery energy with respect to different driving duty cycles. 30-watt robotic power assumed




Figure 146: Achievable range of a 10kg rover with 100 watt-hour battery energy under different driving duty cycles with respect to different rover velocities (30-watt robotics power)

In Figure 145 left), total required energy increases with decreasing rover velocity, for a specified driving duty cycle value. This relationship is especially important in the low rover velocity range where the required energy increases dramatically for any duty cycle. Figure 145 (right) shows that for a specified total battery energy, required rover velocity decreases with increasing driving duty cycle. The greater range achievable with a greater rover velocity for a specified driving duty cycle is shown in Figure 146.

12 TECHNOLOGY ROADMAP

This research identified the critical enabling technologies and formulated those into a technology roadmap for further development toward pit and cave missions. Technologies identified are classified as Autonomy, Sensing & Modeling, Power & Communication, Mechanisms, and Ancillary Beneficial Technologies. Specific technology items in each category are defined in Section 12.1. These technologies are mapped to mission concepts (see Section 7) and identified as baseline or extended capabilities in Section 12.2. Enabling technologies are described in detail in Section 12.3, including comments on state of art and needs for targeted development.

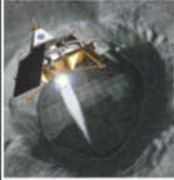



12.1 ENABLING TECHNOLOGIES FOR PIT & CAVE EXPLORATION

T1: Autonomy	T2: Sensing & Modeling	T3: Power & Comm	T4: Mechanisms	T5: Beneficial Technologies
T1.1: Precision landing & hazard avoidance	T2.1: Detecting & modeling pits and caves from orbit	T3.1: Lightweight, high capacity batteries	T4.1 Crawler suspension & wheels	T5.1: Cryogenic survivability
T1.2: Autonomy for scree slope driving	T2.2: Surface modeling capability	T3.2: Remote recharging	T4.2 Tether mechanisms	T5.2: Computing power
T1.3: Autonomy for bouldering operations	T2.3: Flyover modeling capability	(beamed power, and/or charging from Tyrobot)	T4.3 Emplacing infrastructure in caves	T5.3: Long distance traverse
T1.4: Autonomy for tunnel operations	T2.4: Fusion of orbital, flyover, surface, and subsurface data	T3.3: Cave radio (low frequency link for supervision)		
T1.5: Autonomous survey and stamped plans for surveying the surface	T2.5: Long-range sensing in the dark	T3.4: Communication relaying		
T1.6: Supervised autonomy	T2.6: Miniaturization of science instruments	T3.5: Lightweight Tethers		
	T2.7: User interfaces			

12.2 MISSIONS ROADMAP

Table 3 associates technologies with mission concepts developed in this research. A notable finding is that the Scout and Wayfarer missions are possible now and can provide significant value, both scientific and in terms of planning for future robotic and human subsurface exploration. Technologies for the Spelunker mission have been developed and initial feasibility demonstrated, but they require explicit research and development to advance TRL for missions in 10-20 years. Infrastructure emplacement missions for human exploration require many of the same technologies as subsurface robotic missions, and other technologies for long-duration human presence on planetary bodies are already in development for orbit, asteroid, lunar, and Mars missions. With targeted technology development, human missions to pits and caves are possible in 20-30 years. With appropriate planning, these missions may enable lower infrastructure masses and costs through utilization of natural shelter; extend human stay durations through minimizing radiation exposure; and enable unique subsurface science.

Table 3: Technologies mapped to pit and cave mission types with indication of baseline or extended.

	2015	2025	2035	2050
	Scout & Wayfarer Missions		Spelunker Missions	
				
	Baseline	Extended	Baseline	Extended
T1	T1.1, T1.5		T1.2, T1.3, T1.4, T1.6	
T2	T2.1, T2.2, T2.3	T2.4, T2.7	T2.4, T2.6	T2.5, T2.7
T3			T3.4 T3.2 &/or T3.5	T3.1, T3.3
T4			T4.1, T4.2	T4.3
T5		T5.1, T5.2, T5.3	T5.1, T5.2	T5.3

12.3 TECHNOLOGY AREA DESCRIPTIONS

12.3.1 T1. Autonomy

T1.1 Precision landing & hazard avoidance

Autonomous terrain relative navigation and hazard avoidance technologies are essential for flying over and landing near a skylight rim. Registering camera images and LIDAR data to previously generated maps enables localization relative to the mapped terrain. Navigation could be advanced for skylight missions through technologies such as Simultaneous Localization And Mapping (SLAM), which can improve navigation relative to local features (e.g., a pit) on the fly for optimal positioning for sensor views. For hazard detection, the lander scans landing zone with LIDAR, builds a model on the fly, and detects and avoids hazards for safe touchdown. This capability enables landing close to sites of interest located in rougher terrain (e.g., skylights).

Precise, safe landing technologies are currently under development by NASA's [ALHAT](#) group and private enterprises, including [Astrobotic Technology & Carnegie Mellon University](#). The technology has had extensive terrestrial testing, including closed loop propulsive vehicle campaigns by ALHAT and Astrobotic, and is ready for integration into a flight mission.

T1.2 Autonomy for scree slope driving

Driving on loose soil is an art. Detecting dangerous situations and points of no return before they occur is critical to successfully navigation. Traversal is often dynamic and may venture outside of communication coverage, making direct control from Earth infeasible. Counterintuitive toe-in or zigzag maneuvers are often required to negotiate the most challenging terrain. This type of maneuvering is contrary to the state of practice in planetary robotics, which favors simple dynamics so that controls can be defined by direct actions like arc following and turn-to/go-to. The need is for autonomy that can infer terrain properties and automatically adjust during scree slope descent.

This research performed preliminary investigations into terrain sensing for sloped descent. Related work at Carnegie Mellon and Astrobotic is developing terrain property sensing through thermal diffusion (Cunningham, et al. 2015). Prior to incorporation in a mission, significant additional development is required to put these concepts into practice and to develop appropriate control strategies and rules for automated recovery in the event of anomalies. Targeted technology development is recommended.

T1.3 Autonomy for bouldering operations

Similar to scree slope driving, bouldering is a complex and often dynamic activity for planetary rovers, unsuitable for state-of-art autonomous operations and degree of human-in-the-loop. Rock crawlers on Earth can surmount boulders many times wheel diameter and climb dynamically up near vertical rock faces, but this has not been reliably demonstrated robotically.

Prior to incorporation in a mission, significant additional development is required to put these concepts into practice and to develop appropriate control strategies and rules for automated recovery in the event of anomalies. Targeted technology development is recommended.

T1.4 Autonomy for tunnel operations

As robots venture beyond the skylight floor and into caves, autonomous driving and mapping becomes exceedingly important. While operators may be able to communicate at reasonable data rates while in line of site to a comm relay for control, data rates may be extremely limited to zero in periods without line of sight.

Terrestrial robotic systems have demonstrated remarkable capability for exploring and mapping subterranean voids to gather information on where problems such as encroachment, collapse, flooding, and subsidence can occur. Autonomous operation has demonstrated coverage, quality, and economy of robotic approaches relative to traditional human approaches. Prior work at Carnegie Mellon University demonstrated the mechanisms, sensing, and software of subterranean operations in active, abandoned, and submerged subterranean spaces (Morris, et al. 2006). Additional development for planetary application is required before incorporation in missions and targeted technology development is recommended.

T1.5 Autonomous Survey

Most rover exploration during planetary missions has involved traveling from one interest point to the next and doing scientific investigations at these points. This approach can build detailed models along the rover's route, but it does not attempt to achieve complete coverage of an area. A full coverage survey of terrain around a pit, one of the mission objectives discussed for the Scout and Wayfarer missions, could provide evidence of subsurface voids and their extent. The simplest approach is to follow a lawnmower or spiral pattern where the distance between passes is determined by the field of view of the robot's sensors. This works for aerial or underwater vehicles, but surface rovers must adapt survey plans to terrain obstacles.

Some work has been done on this problem for terrestrial systems (Hodo 2007). More development would be needed to ensure that the technology is ready for flight, including development of a planetary-relevant localization method with sufficient precision for this application.

T1.6 Supervised Autonomy

Unique mission operations software is essential for supervised autonomy for descent and subsurface operations. Mission operators must clearly understand what a robot is doing when and why and be able to command actions, at least at a high level to adjust mission objectives in operation. Appropriate feedback and command levels under differing conditions of data rate, terrain hazard, power level, and type of operation (e.g., driving versus sampling) must be developed.

Supervised autonomy is well studied in terrestrial robotics, and is applied in current planetary missions like MER and MSL. The planetary rovers operate in considerably more moderate terrain conditions and the terrestrial rovers typically operate without the data rate restrictions or lack of physical human intervention of planetary missions. Additional development for planetary pit application is required before incorporation in missions and targeted technology development is recommended.

12.3.2 T2. Sensing & Modeling

T2.1 Detecting & modeling pits and caves from orbit

Prior work has demonstrated automated detection of pits in imagery, though this work has limited applicability to pits at high latitudes. Stereo reconstruction of pits from orbital imagery has been demonstrated, including the example presented in this report. Models could be enhanced using techniques that take advantage of known illumination or assumptions based on geological models. Geometric models of terrain combined with geological models could predict cave locations. GPR or gravimetric sensing (using existing or new gravimetry data) could also be used to detect subsurface voids. Additional research to model pits from orbital data and targeted collection of new orbital data for improved modeling of specific pits could be performed now.

T2.2 Surface modeling capability

This work demonstrated modeling a pit from a surface perspective with an analog field test. Several additional developments are needed to get the technology ready for missions. Camera parameters, such as FOV, must be carefully selected. Depending on the onboard computing capabilities of the robots intended for flight, it may be possible to build pit models on the planetary surface and then send them back to Earth. If this is judged to be feasible (or, for a very low communications-bandwidth mission, necessary), then flight software for model building would need to be developed. Adaptations to surface modeling are also needed to consider time-varying illumination.

View trajectory planning enables rovers to efficiently create high-quality models of pits from the surface, even as illumination conditions change. Work during this project laid out the view trajectory planning problem and began algorithm development. Further work is needed to prepare this technology for flight.

T2.3 Flyover modeling capability

Flyover modeling happens quickly; there is only one chance to get it right for a given mission. Sensor parameters such as FOV, exposure time, and range must be carefully considered. Better resolution can be achieved with smaller FOV sensors, but then those sensors need to be swept over the pit during flyover. This sweep must be carefully planned such that it is robust to variations in trajectory. The lander's position during flyover must also be carefully tracked to facilitate stitching multiple sensor measurements into a coherent model. Inspired by this research and funded by NASA's Flight Opportunities, a related project at Carnegie Mellon is investigating these issues and working toward a field test of flyover pit modeling, slated to occur early in 2015. This test will significantly advance the technology readiness level of flyover modeling capabilities.

T2.4 Fusion of orbital, flyover, surface, and subsurface data

Combining 3D models from different sources that share a similar perspective is relatively straightforward to do with existing techniques. Combining models from very different perspectives is a much harder problem. In combining flyover and surface models, for example, a lander and a rover may have seen different sides of the same rock. If there were not enough overlap, it would be quite difficult to determine that the two views corresponded to the same feature. Development of technology for combining models captured from different perspectives is needed.

T2.5 Long-range sensing in the dark

State-of-art surface sensing techniques like stereovision require innovations in active illumination for relevance in this domain. Beyond hazard detection, dark modeling technologies that provide the same high-resolution immersion expected of surface robots to inspire and engage the public do not exist. Radically new modalities of sensing require investigation for application to subsurface missions, and could dramatically enhance the quality of models generated by these missions. Targeted development of sensing modalities for planetary caves is required to enable this.



Figure 147. Waitomo Cave is lit by thousands of glow worm insects. Radical new modalities of sensing like thrown light beacons could be ideal for subsurface sensing. These could provide corrective landmarks for mapping as well as illumination for long-range photography, hazard detection, and model building. Prior work has shown that even throwing a light source can provide 3D model estimates.

T2.5 Miniaturization of science instruments

For subsurface exploration, power, mass, and volume are significantly limited. Small, low-power instruments are needed. Significant work has been done terrestrially to miniaturize cameras, and sensors like flash LIDAR and thermal cameras are also getting smaller. Work will be needed to adapt existing miniaturized sensors for space, and instruments for scientific investigations in caves will also have to be miniaturized and readied for space.

T2.6 User interfaces

Most planetary data collected to date have been from environments where geometry is easily represented as 2.5-D (i.e., by digital elevation maps). For pits and caves this is not the case. Research is needed into representation of data collected in a substantially 3D environment to humans on the ground. 2D maps on a computer screen will not be sufficient to enable scientists to effectively conduct experiments in planetary caves. One option is to build immersive models of pits and caves that scientists can fly through in virtual reality, but how can the user's experience then be enhanced to draw attention to important features that may not be noticed otherwise? This research produced several models and visualization methods to enhance user understanding, but more could be done to tailor these methods to specific mission objectives and to evaluate user responses.

Also, given limited communications, an immersive model may not be possible, or it may not be feasible in the short timeframe that mission controllers have to make decisions. What kinds of data are most important to send back to assist in decision-making, and how should they be represented to users?

12.3.3 T3. Power & Communications

T3.1 Lightweight, high capacity batteries

High energy density batteries enable longer cave excursions with low battery masses. Batteries provide significant cost and programmatic advantages over radioisotope options for power. High energy density batteries for planetary missions are in development for numerous missions under NASA, in accordance with NASA roadmap element 3.2.1 *Batteries*. Therefore, targeted development for pit and cave application is not required.

T3.2 Remote recharging

Subsurface robots have limited views of the Sun, and dragging tethers limits range and adds mass and risk of entanglement. Remote recharging could enhance robot range, either through capacitive/conductive recharge by docking as demonstrated in this research (Section 11.1.5) or by power beaming. Power beaming enables recharging of subsurface robots from a solar-powered lander. Power beaming is under development by several groups including LaserMotive, who won the NASA Power Beaming Challenge and are presently working for NASA to design the architecture to use lasers to launch rockets and power satellites, and, eventually, power lunar bases (LaserMotive 2012). Charging from Tyrobot would require targeted technology development. Power beaming is under development for other applications and may not require specific targeted development for pit and cave missions.

T3.3 Cave radio

Low-frequency “cave” radio enables communication from a base inside the tube to mobile subsurface explorers to pass through some rock obstructions. This reduces risk associated with communication loss. Limited data link through rock can be achieved with very low-frequency radio or magneto-inductive comm. Cave radios are in use terrestrially for mine disaster rescue. They have recently undergone significant reductions in mass and power, presenting promise that they could be modified for planetary missions. Lack of moisture in planetary environments, which is major source of signal degradation, has the potential to improve performance over terrestrial radios.




Figure 148: A limited data-rate link through cave ceilings can be achieved using very low-frequency radio or magneto-inductive comm

T3.4 Communication relaying

Surface-to-surface radios could enable communication from a Tyrobot or a base inside the tube to mobile subsurface explorers. Surface radios are under development at NASA JSC for future planetary missions on Mars and it is not expected that targeted technology development is needed.

T3.5 Lightweight Tethers

Lightweight power and data cabling enables deploying tethered robots into subsurface voids to establish power and communication nodes. Low mass reliable cabling will reduce mission cost and risk. This is a key NASA technology development associated with planetary base infrastructure and is associated with *NASA roadmap element 3.3.3, Power Distribution and Transmission*.

Low mass reliable cabling will reduce mission cost and risk. Tether development for pit and cave exploration is unique and challenging. The tether may need to function as a rappel rope. The tether is subject to bending and abrasion at the lip of the skylight, since the mechanical, power and data connectivity must span back to the lander. Exposure of the surface segment leads to huge thermal swings with day/night cycle. Length approaching 300 meters requires attention to compatible mechanical stiffness and thermal expansion in the coaxial layering of data, power, insulation, and strength and abrasion layers. Miniaturization and light weighting are paramount, since the tether must be carried, then reeled out from the rover in order to avoid dragging during deployment or extensive sliding at the skylight rim. Targeted development for pit and cave exploration is recommended.




Figure 149: Cross section of possible tether design, showing key elements




Figure 150: Thermal gradient on a tether used for skylight descent will be large

12.3.4 T4. Mechanisms

T4.1 Crawler suspension & wheels

Terrestrial rock crawlers perform with flexible suspension and large contact patch rubber wheels for bouldering and driving on loose material. The thermal ranges and vacuum of planetary environments challenge terrestrial suspension elements (e.g., springs and dampers). Rubber outgasses in vacuum, making it impractical for rover wheels. Not-rubber, non-pressurized flexible wheels are in development by various companies, including Michelin's Tweel and the metallic wheels on Apollo's Lunar Rover Vehicles. Dynamic suspensions require development for crawlers and may require targeted technology development. This type of suspension has broader applicability to high speed and rough terrain surface operations.

T4.2 Tether mechanisms

Specialized mechanisms for managing and utilizing tethers are essential to subsurface missions. Mechanisms might include: tether unspooling and spooling; methods to connect and disconnect from

Tyrolean lines; and deploying, emplacing, and anchoring Tyrolean lines. Preliminary versions of these mechanisms were demonstrated in this research as well as in prior work including JPL's Axel rover (Nesnas, et al. 2008) and [Tethers Unlimited](#). This work lays the foundation for development and targeted technology development is required to advance for mission readiness.

T4.3 Emplacing infrastructure in caves

Emplacing infrastructure enables habitat building and more complex robotic operations. Infrastructure may include: surface communication and power stations with connections to subsurface; habitat modules inside a cave; descent "elevators" for entry and exit by robots or astronauts; or resource processing facilities. Technologies must be developed to emplace these infrastructure elements. This involves modification of technologies suited for robot exploration such as tethered descent for large payloads and safe, reliable conveyance of humans. Preliminary investigations of mission concepts for human habitation and feasibility studies of technology are recommended as a next step.

12.3.5 T5. Ancillary Beneficial Technologies

Ancillary beneficial technologies are broadly applicable to planetary exploration rather than being explicitly driven by or required by pit and cave exploration.

T5.1 Cryogenic survivability

Extended robotic missions on planetary surfaces confront extreme conditions, where violent thermal cycles, hard vacuum, and relentless radiation exposure conspire to degrade electronics. Innovative cryo-tolerant avionics technologies for planetary landers and rovers are crux enablers of reliable, multi-year operations on planetary scale. In prior lunar missions that attempted overnight survival, such as Lunar Surveyor, the battery and power system was the weak link. Preliminary investigations indicate that off-the-shelf components have promise for cryogenic survivability (Minogue 2012). In particular, non-aqueous battery chemistries and electronics without sensitive capacitors or thermal strain weak points have been shown to survive multiple cycles of cryogenic survival. Targeted technology development would be required to bring this to fruition, as no ongoing funded programs exist. This capability would be highly beneficial to planetary missions by reducing the cost and program associated with radioisotope for overnight heating.

T5.2 Computing power

A key challenge for autonomous control of surface and subsurface explorers and model generation off-Earth is operating the high computation algorithms on space-relevant hardware. Derivatives of intensive algorithms like keyframe SLAM have been ported to parallel computing environments like graphics cards (Hendeby, Karlsson and Gustafsson 2010). Astrobotic has performed initial testing of GPU acceleration for terrain relative navigation and has demonstrated sufficient speed of visual feature detection on flight-relevant computing. This is a key element for pose determination and map building, relative to conventional computing.

A promising approach that is gaining traction is to incorporate redundant, rugged, military grade COTS computers to perform GPU-accelerated visual computations and high computational load SLAM updates. The COTS computers perform modeling and navigation algorithms, while a traditional space-qualified processor performs filtering operations. This has the potential to broadly impact NASA missions from Earth orbit to planetary rovers to deep space explorers.

T5.3 Long distance traverse

Wayfaring missions and missions with human explorers are significantly enhanced by long distance and high speed traversal relative to state-of-art planetary rovers. Wayfarers may spend significant time investigating a pit, operating at low speed and low driving duty cycle while taking imagery and sampling. When investigation of a pit is complete, it is highly advantageous to shift into a high-speed, high-duty cycle driving mode to rapidly traverse between one pit of interest and another. This demands autonomy for high speed and mechanisms for high cycle life. This has the potential to broadly impact NASA planetary missions and is aligned with goals of *NASA Technology Roadmap Element TA4.0*.




Figure 151: Approximate locations of a number of known pits (pink dots, images) on the lunar surface. With long distance traverse technologies, a Wayfarer mission visit widespread locations, such as these for global surveys on the Moon, Mars or Mercury.

13 IMPACT

This research has had broad impact, both here in the US and in the international community. NASA's Science Mission Directorate recognizes this research as having

[...] far-reaching connections to the kinds of new science that the AGENCY needs to be doing. It is responsive to new locales and vantage points for high-value science and ultimately human exploration of accessible solar system destinations, and it is broadly applicable to several NRC Decadal Survey (DS) thematic goals as described in the 2011 Planetary DS.

The progress to date is "beyond impressive" with a tapestry of practical field demonstrations tied to real questions with direct relevancy to science at the frontier (and well connected to several Mars priorities). The connection to near-term flight opportunities (Lunar X-prize) [...] but also as enabling capabilities suitable for openly-competed SMD missions offered via AO's such as DISCOVERY and NEW FRONTIERS is value-added for NASA and the community.

The demonstration of new approaches for data fusion to support advanced surface (and sub-surface) robotic exploration is particularly attractive as it presents a pragmatic demonstration of what could be possible (and affordable) for NASA in the 2020's and beyond. [...] Red's NIAC activities showcase many of the elements of what would be needed to lower risk, improve science yield, and ultimately to reduce cost. If we are to pursue new environments with science potential that was unrecognized only a few years ago (such as in Pits, Caves, lava tubes), then this is the right step at the right time.

- Dr. James B. Garvin
Chief Scientist
NASA Goddard Space Flight Center

Follow-on funding has been provided by the Undergraduate Student Instrument Project for a team at Carnegie Mellon to perform flyover pit modeling with a Masten Space Systems' lander in the Mojave Desert (see Section 13.3). The Japanese aerospace agency, JAXA, has begun funding pit exploration research. A vignette about cave exploration, pulling on the language and ideas of this research, was included for the first time in the national robotics roadmap (A Roadmap for U.S. Robotics: From Internet to Robotics 2013).

This research resulted in publications in both technical and scientific conferences (Section 13.1) and inspired related thesis research (Section 13.2). Funding from NASA and other sources for related projects has grown the pit and cave exploration research community at Carnegie Mellon and Astrobotic beyond the NIAC work (Section 13.4). Connections forged with the science community have helped scientists see the possibilities of robotic exploration (Section 13.5). Ongoing mission development work at Astrobotic and Carnegie Mellon, including Astrobotic's selection for the CATALYST program, will help make the first pit exploration mission a reality within the next few years (Section 13.4).

13.1 PUBLICATION

During phase I of this project, the team submitted an abstract to the First International Planetary Caves Workshop entitled “Mission Design for Combined Lander-Rover Modeling of a Skylight” (Peterson, Jones and Whittaker 2011). Two project team members attended the workshop to give an oral presentation of the work. A paper was also published in the Field and Service Robotics conference entitled “Complementary Flyover and Rover Sensing for Superior Modeling of Planetary Features” (Jones, Wong, et al. 2012) This paper compared modeling a skylight pit from lander flyover, modeling the pit from surface rover reconnaissance, and modeling the pit from a combination of lander flyover and surface rover reconnaissance planned from the lander model.




Figure 152: Figure showing the mission concept evaluated in “Complementary Flyover and Rover Sensing for Superior Modeling of Planetary Features” (Jones, Wong, et al. 2012)

In phase II, the project team submitted an abstract to the 2013 Lunar and Planetary Science Conference entitled “Skylight: Mission to Investigate and Model a Lunar Pit” (Jones, Peterson, et al. 2013). PI Red Whittaker attended the conference to present a poster on the work.

The project team also submitted an abstract that was accepted to the Lunar Exploration Analysis Group, an interdisciplinary group that provides NASA analysis on scientific, technical, and commercial developments for the Moon. That abstract was titled, “Astrobotic Technology: Planetary Pits and Caves for Science and Exploration” (Huber, et al. 2014). Authors Kevin Peterson and Dan Hendrickson attended the conference and presented a poster on the work.

A related project that focused on the issues encountered by multi-robot teams exploring planetary caves, a paper entitled “Mapping Planetary Caves with an Autonomous, Heterogeneous Robot Team” was published in the 2013 IEEE Aerospace Conference (Husain, et al. 2013).

Abstracts and oral presentations were submitted to the Workshop on Planetary Volcanism at USRA – “Robotic Technologies for Exploration of Planetary Pits and Caves” covering results-to-date of primarily NIAC- funded work and “Pit Modeling from Lander and Rover Reconnaissance,” discussing various mission scenarios with near-term applicability. Unfortunately, the workshop was canceled.

Future publications are intended for this research. A paper on surface rover reconnaissance and modeling of planetary pits entitled “Planning Views to Model Planetary Pits and Other Features under Transient Illumination” was submitted to the 2015 IEEE Aerospace Conference and accepted for publication (Jones, Tabib and Whittaker n.d.). A journal paper on the field testing done for this project is also intended.

This NIAC project was also featured in a Nature article (Powell 2012).


13.2 RELATED THESIS WORK

Uland Wong, an author of this report whose work in Lumenhancement produced some of the strongest results in this research, completed his PhD in robotics at Carnegie Mellon during phase I (U. Wong 2012).

Heather Jones, another author of this report, is currently pursuing a PhD in robotics at Carnegie Mellon. Her research has looked at multiple aspects of planetary cave exploration with robots as a part of this project. Her current research focus is on view trajectory planning for modeling pits from a surface rover with consideration for transient illumination on planetary bodies.

13.3 MISSION DEVELOPMENT

In 2011, the project team submitted to NASA a Technology Demonstration Mission proposal entitled “Skylight: Pinpointing Planetary Destinations,” for a mission that would fly over a lunar skylight during landing and return to model the pit with a rover. Although this proposal was not funded, it served as a forcing function for the team to align on a mission-oriented view for robotic exploration of planetary caves. It was also an important step in a longer-term effort to build a partnership between Astrobotic, Carnegie Mellon, and NASA toward a commercial lunar landing capability. The Astrobotic/NASA side of that partnership was formalized in April 2014 in NASA’s Lunar Cargo Transportation and Landing by Soft Touchdown (Lunar CATALYST) program. Astrobotic was selected as an industry partner for development of robotic lunar landers that can be integrated with U.S. commercial launch capabilities to deliver payloads to the lunar surface.



Lunar pits and caves are exciting new destinations for future science and exploration missions and crewed lunar outposts.

Pit and Cave Science:
Pits and caves on bodies throughout the solar system are opportunistic targets to study planetary origins, geology, climate, and potentially even biology. Pit crater chains have been identified on Venus, Earth, the Moon, Mars, Eros, and Europa. Caves, pits, and lava tubes present unique opportunities to investigate planetary volatiles, mineralogy, origins, morphology, geologic history, and exposed surfaces without regolith cover.

Pit and Cave Exploration:
Planetary caves could provide a safe haven from a number of hazardous conditions that are inherent to planetary environments. For instance, human habitats in lunar caves would protect crews from the extreme temperature variations, radiation bombardment, and micrometeorite hazards inherent in living and working on the lunar surface. The potential for risk to human explorers demands robotic precursors as the first step to cave exploration.

Robotic explorers can approach, enter, and model these features at thousand-fold increased resolution relative to orbital means, and perform physical sampling and analysis to achieve unique science not otherwise possible.

Missions:
Planetary pit and cave exploration demands bold new technology for flyover, apert viewing, pit descent, rock crawling, and robotic caving. Modeling, visualizing, and compressing the multi-physical data associated with scientific investigation of this kind will require new techniques that must be prototyped and integrated with science, and evaluated in analog experimentation.

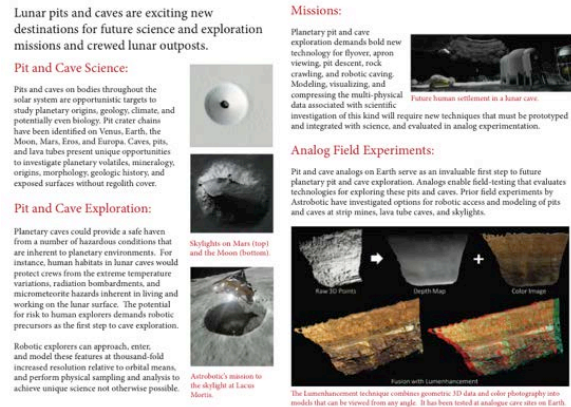
Analog Field Experiments:
Pit and cave analogs on Earth serve as an invaluable first step to future planetary pit and cave exploration. Analog enable field-testing that evaluates technologies for exploring these pits and caves. Prior field experiments by Astrobotic have investigated options for robotic access and modeling of pits and caves at strip mines, lava tube caves, and skylights.

Future human settlement in a lunar cave.

Skylights on Mars (top) and the Moon (bottom).

Astrobotic's mission to the skylight at Lavinia Mine.

The Lumenhancement technique combines geometric, 3D data and color photography into models that can be viewed from any angle. It has been tested at analogue cave sites on Earth.





Astrobotic's first target mission – to a skylight located at the Lacus Mortis region on the Moon – has been strongly guided by this research. The robotic configurations and modeling technologies developed for exploration of planetary pits and caves have directly informed and enabled planning for Astrobotic's mission. Significant planning and hardware development for the mission has taken place. For instance, the lander has been designed and constructed, and its communications system tested in concert with a rover built by CMU at a lunar analog test site. An autonomous landing system has been tested in the Mojave Desert using three suborbital launch vehicle test flights that successfully demonstrated the system's ability to scan terrain, compare the landing area to a map, avoid hazards, and make a safe powered descent to the ground. Mission flight trajectories have been calculated, and mass budgets for the spacecraft and its carrier payloads have been developed. Moreover, the landing site at the Lacus Mortis skylight was carefully selected using observational data from Lunar Reconnaissance Orbiter. Analysis of the Lacus Mortis skylight to determine entry points for potential caves has also been conducted.

The work completed by the project team has also enabled submission of a NASA SBIR proposal that proposes the development of detailed mission concepts for future human habitation in lunar caves. This study would build on the work of the project team and investigate emplacing habitats in lunar caves. The study would develop a roadmap of critical enabling technologies for these missions that address robots for emplacing habitats, safe entry and exit for humans, architecture for power generation, architecture for communication to Earth, and sufficiency of existing habitation modules.

Skylight: Pinpointing Planetary Destinations

BAA NNM11ZDA001K / June 24, 2011
Astrobotic Technology Inc.




Figure 153: Cover page for Skylight Technology Demonstration Mission proposal

13.4 RELATED RESEARCH AND FUNDING

An undergraduate group at Carnegie Mellon affiliated with this project team submitted and won a proposal to NASA's Undergraduate Flight Opportunities program entitled "Flyover Mapping and Modeling of Terrain Features" to do flyover modeling of a pit from a propulsive lander. This team is currently working toward a flight on a Masten Space Systems' lander in early 2015. Students in this group have also won funding from Carnegie Mellon for pit and cave related research. Rick Shanor received a Summer Undergraduate Research Fellowship entitled "Subsurface Exploration of Lunar and Martian Caves." Rick Shanor and Ashrith Balakumar won a Small Undergraduate Research Grant (SURG) for "Mechanical Design of Skylight Survey Instrument," and Eddie Nolan, Brent Strysko, and Neal Bhasin won a SURG for "Skylight Survey Electronics and Avionics."

Heather Jones secured a NASA Space Technology Research Fellowship in 2012 on a topic closely aligned with this project. Renewing this fellowship for multiple years, she has built relationships between the robotic pit/cave exploration project team at Carnegie Mellon and Astrobotic and scientists and engineers at NASA Ames. Through her involvement at Ames, she was able to leverage a trip to Craters of the Moon in Idaho by the Ames-led FINESSE SSERVI team and facilitate participation by Carnegie Mellon and Astrobotic NIAC team members. This not only enabled some field testing and fostered interaction with a larger community of geologists and planetary scientists.

The project has involved graduate and undergraduate students through Carnegie Mellon’s Robotics Institute Summer Scholars program and through Astrobotic’s Internship program. Graduate and undergraduate students in the Carnegie Mellon courses “Mobile Robot Design” and “Advanced Mobile Robot Development” have also been involved in research for this project.

A NASA STTR Phase I project was funded on the related topic of Multi-Robot Systems for Subsurface Planetary Exploration. This project explored issues specific to multi-robot teams used in planetary cave exploration.

Several other related proposals were submitted, including phase 2 for the multi-robot cave exploration STTR, a NASA Moon And Mars Analog Missions Activities proposal, and a NASA SBIR, entitled “MapSense: Innovative 3D Mapping System,” which would have taken the LIDAR and camera fusion methods in this work and made a more portable and space relevant sensing and modeling package. Although these proposals were not funded, they helped craft a roadmap for future research in robotic exploration of planetary pits and caves.



Figure 154: Multi-robot team in front of a building used to construct an analog tunnel and conduct experiments in multi-robot cave exploration for the NASA STTR Multi-Robot Systems for Subsurface Planetary Exploration



13.5 SCIENCE EFFORTS

The Carnegie Mellon and Astrobotic researchers on this project built a multi-institutional team of scientists and planetary exploration experts to submit a proposal to NASA's Solar System Exploration Research Virtual Institute (SSERVI) Cooperative Agreement Notice. This proposed SSERVI, titled "CAVES: Cave and Void Exploration Science," laid out an agenda to research pits and caves on the Moon, the moons of Mars, and asteroids. While CAVES was not selected for funding, the proposal effort grew the community of scientists and engineers pursuing exploration of pits and caves on other planetary bodies and fostered a greater understanding of pit/cave science issues among the Carnegie Mellon and Astrobotic pit/cave team.

14 CONCLUSIONS

Pits and caves are opportunistic study targets for unique origins, geology, climate, and astrobiology that will broadly impact planetary science. Although missions for deep-cave access require substantial development, minimalist near-term missions could yield great gains. In these missions, pits can be modeled from bird's eye view during lander flyover at much greater resolution and accuracy than achievable from orbit and at close enough range for active illumination techniques such as LIDAR. Pits can be modeled by robot circumnavigation at even greater resolution and accuracy, with access to lower view angles and varied lighting conditions.

Robotic missions can access the full range of pit morphologies. Pits can be accessed through flyover, perimeter traverse, Tyrolean lines, and free or tethered descent. Some pits offer ramps for possible rover descent without specialized descent mechanisms. Small, battery-powered, wheeled rovers can explore caves and lava tubes. Innovative techniques based on mathematical models of the physics of light, such as Lambertian algorithms, improve pit and cave models by an order of magnitude. Good pit models can be constructed from images only. Time history of illumination must be considered in planning robotic pit exploration routes. Experimental methods in this research compare results of innovative techniques with ground truth survey.

The study of pits and voids lays a foundation for human exploration of planetary bodies. Voids such as caves and tubes that are accessible through pits may provide shelter from harsh planetary environments, reducing the complexity of human survival beyond Earth. That is authentic exploration. Robots must precede humans, since the promise of caves is tempered by their great unknown. Safety and viability must be insured before humans explore and utilize these enchanting, invaluable destinations. Much discovery awaits below the surface. Subsurface explorers may discover signs of life, origins and resources, and possibility of havens unseeable from orbit. Robots and technologies envisioned here will explore, pioneer, and prepare our way to the subsurface worlds that await beyond Earth.

15 ACKNOWLEDGEMENTS

The authors would like to thank Kevin Peterson of Astrobotic Technology for early leadership in the flyover and surface modeling mission concept. Joshua Yee, Zachary Medeiros, and Carson Coulter contributed to the tension testing for Tyrobot. Preston Ohta, formerly of Astrobotic Technology, contributed to the Krawler work. Carnegie Mellon University students in 16-861 Mobile Robot Design and 16-865 Advanced Mobile Robot Development also made contributions to the research.

Doug Owen provided invaluable guidance on field sites and operations at Craters of the Moon. Scott Hughes of Idaho State University, Darlene Lim of NASA Ames Research Center, and Linda Kobayashi of NASA Ames Research Center facilitated collaboration between this NIAC project and the FINESSE SSERVI project. W. Brent Garry of NASA Goddard Space Flight Center and the FINESSE SSERVI project provided DGPS readings to enable better alignment of models with geographic coordinate systems. FARO provided their x330 laser scanner for ground truth measurements. Operators at FIEG Brothers Mine sites graciously allowed this NIAC team to conduct analog field tests at several mine locations. PBS Coals Inc. provided the site and training for the Sheep Ridge mine pit experiments. Penny Boston of New Mexico Tech provided invaluable insights into the science of planetary caves. James Ashley of Arizona

State University helped the team come to an early understanding of lunar pits. Mark Robinson and Robert Wagner of Arizona State University have provided ongoing insights into lunar pits and how to interpret orbital images of them.

16 REFERENCES

"A Roadmap for U.S. Robotics: From Internet to Robotics." Georgia Institute of Technology, Carnegie Mellon University, Robotics Technology Consortium, University of Pennsylvania, University of Southern California, Stanford University, University of California-Berkeley, University of Washington, Massachusetts Institute of Technology, 2013, 90-91.

Ashley, J. W., et al. "Lunar Pits: Sublunarean Voids and the Nature of Mare Emplacement." *Lunar and Planetary Science Conference*. The Woodlands, TX, 2011.

Ashley, James W., et al. "Lunar Caves in Mare Deposits Imaged by the LROC Narrow Angle Cameras." *First International Planetary Caves Workshop*. Carlsbad, NM, 2011.

Bares, John E, and David S Wettergreen. "Dante II: Technical Description, Results, and Lessons Learned." *International Journal of Robotics Research* 18, no. 7 (1999): 621-649.

Bleacher, J E, R Greeley, D A Williams, S C Werner, E Hauber, and G Neukum. "Olympus Mons, Mars: Inferred changes in late Amazonian aged effusive activity from lava flow mapping of Mars Express High Resolution Stereo Camera data." *Journal of Geophysical Research*, 2007.

Bleacher, J E, R Greeley, D A Williams, S R Cave, and G Neukum. "Trends in effusive style at the Tharsis Montes, Mars, and implications for the development of the Tharsis province." *Journal of Geophysical Research*, 2007.

Broxton, M J, and L J Edwards. "The Ames Stereo Pipeline: Automated 3D Surface Reconstruction from Orbital Imagery." *Lunar and Planetary Science Conference*. 2008.

Cignoni, P, M Corsini, and G Ranzuglia. "MeshLab: an open-source 3d mesh processing system." *ERCIM News*, April 2008.

"CloudCompare (version 2.5.5) [GPL software]." EDF R&D, Telecom ParisTech, 2014.

Coombs, Cassandra R, and B Ray Hawke. "A Search for Intact Lava Tubes on the Moon: Possible Lunar Base Habitats." In *NASA. Johnson Space Center, The Second Conference on Lunar Bases and Space Activities of the 21st Century*. 1992.

Cunningham, Chris, Uland Wong, Kevin Peterson, and William Whittaker. "Predicting Terrain Traversability from Thermal Diffusivity." *Field and Service Robotics* (Springer) 105 (2015): 61-74.

Cushing, G E. "Candidate cave entrances on Mars." *Journal of Cave and Karst Studies*, 2012: 33-47.

Cushing, G. E., T. N. Titus, and E. Maclennan. "Orbital Observations of Martian Cave-entrance Candidates." *First International Planetary Caves Workshop*. Carlsbad, NM, 2011.

- Diebel, James, and Sebastian Thrun. "An Application of Markov Random Fields to Range Sensing." *Conference on Neural Information Processing Systems (NIPS)*. 2005.
- Dubowsky, S, K Iagnemma, and P Boston. "Microbots for Large-Scale Planetary Surface and Subsurface Exploration." NIAC Phase I study, 2005.
- Dubowsky, S., J.S. Plante, and P. Boston. "Low Cost Micro Exploration Robots for Search and Rescue in Rough Terrain." *IEEE International Workshop on Safety, Security and Rescue Robotics*. Gaithersburg, MD, 2006.
- Fairfield, Nathaniel, George A Kantor, and David Wettergreen. "Real-Time SLAM with Octree Evidence Grids for Exploration in Underwater Tunnels." *Journal of Field Robotics*, 2007.
- Fairfield, Nathaniel, George A Kantor, and David Wettergreen. "Segmented SLAM in Three-Dimensional Environments." *Journal of Field Robotics*, 2010.
- . "Three Dimensional Evidence Grids for SLAM in Complex Underwater Environments." *International Symposium of Unmanned Untethered Submersible Technology*. 2005.
- . "Towards Particle Filter SLAM with Three Dimensional Evidence Grids in a Flooded Subterranean Environment." *International Conference on Robotics and Automation*. 2006.
- Frankot, R., and R. Chellappa. "A Method for Enforcing Integrability in Shape from Shading Algorithms." *IEEE Pattern Analysis and Machine Intelligence* 10, no. 4 (1988).
- Furukawa, Y, and J Ponce. "Accurate, Dense, and Robust Multi-View Stereopsis." *IEEE Transactions on Pattern Analysis and Machine Intelligence* 32, no. 8 (August 2010): 1362-1376.
- . "Accurate, Dense, and Robust Multi-View Stereopsis." *IEEE Conference on Computer Vision and Pattern Recognition*. 2007.
- Gaddis, L R, et al. "An Overview of the Integrated Software For Modeling Imaging Spectrometers (ISIS)." *Lunar and Planetary Science Conference*.
- . "An Overview of the Integrated Software For Modeling Imaging Spectrometers (ISIS)." *Lunar and Planetary Science Conference*. 1997.
- Gould, S., P. Baumstarck, and M. Quigley. "Integrating Visual and Range Data for Robotic Object Detection." *Proceedings of European Conference on Computer Vision (ECCV)*. 2008.
- Grossman, J., and W. Dally. "Point Sample Rendering." *Rendering Techniques*, 1998.
- Haruyama, J, et al. "Possible lunar lava tube skylight observed by SELENE cameras." *Geophysical Research Letters*, 2009.
- Heiken, Grant H, David T Vainan, and Bevan M French. *Lunar Sourcebook: A User's Guide to the Moon*. New York, NY: Cambridge University Press.
- Hendeby, Gustaf, Rickard Karlsson, and Fredrik Gustafsson. "Particle Filtering: The Need for Speed." *EURASIP Journal on Advances in Signal Processing*, 2010.

- Hodo, David. "Development of an Autonomous Mobile Robot-Trailer System for UXO Detection." MS Thesis, Auburn University, Auburn, AL, 2007.
- Huber, S A, D B Hendrickson, H L Jones, J P Thornton, W L Whittaker, and U Y Wong. "Astrobotic Technology: Planetary Pits and Caves for Science and Exploration." *Annual Meeting of the Lunar Exploration Analysis Group*. 2014.
- Huntsberger, Terry, Vivek A Sujan, Steven Dubowsky, and Paul S Schenker. "Integrated System for Sensing and Traverse of Cliff Faces." *In Proceedings of Aerosense'03: Unmanned Ground Vehicle Technology V*. 2003.
- Husain, A, et al. "Mapping Planetary Caves with an Autonomous, Heterogeneous Robot Team." *IEEE Aerospace Conference*. 2013.
- Intelligent Robotics Group. "The Ames Stereo Pipeline: NASA's Open Source Automated Stereogrammetry Software." NASA Ames Research Center, 2013.
- Jones, H L, K M Peterson, W L Whittaker, and U Y Wong. "Skylight: Mission to Investigate and Model a Lunar Pit." *Lunar and Planetary Science Conference*. 2013.
- Jones, H L, U Wong, K Peterson, J Koenig, A Sheshadri, and W L Whittaker. "Complementary Flyover and Rover Sensing for Superior Modeling of Planetary Features." *Proceedings of the 8th International Conference on Field and Service Robotics*. 2012.
- Jones, Heather L, Wennie Tabib, and William L Whittaker. "Planning Views to Model Planetary Pits under Transient Illumination." *IEEE Aerospace Conference*.
- Kazhdan, M, and H Hoppe. "Screened Poisson Surface Reconstruction." *ACM Transactions on Graphics* 32, no. 3 (2013).
- Kazhdan, M, M Bolitho, and H Hoppe. "Poisson Surface Reconstruction." *Eurographics*. 2006.
- Kerl, C., J. Sturm, and D. Cremers. "Dense Visual SLAM for RGB-D Cameras." *International Conference on Intelligent Robotics and Systems (IROS)*. 2013.
- Koudelka, M. L., P. Belhumeur, S. Magda, and D. Kriegman. "Image-based Modeling and Rendering of Surfaces with Arbitrary BRDFs." *International Conference on Computer Vision and Pattern Recognition*. 2001.
- Krishnan, A K, P McGarey, S Saripalli, and J F Bell. "NIR-CAM : Development of a Near Infrared camera." *IEEE International Symposium on Robotic and Sensors Environments (ROSE)*. 2013.
- LaserMotive. "Laser Power Beaming Fact Sheet." 2012. www.lasermotive.com.
- Lee, Kuang-Chih, J. Ho, and D. Kriegman. "Nine points of light: acquiring subspaces for face recognition under variable lighting." *IEEE Conference on Computer Vision and Pattern Recognition*. 2001.
- Léveillé, Richard J., and Saugata Dattab. "Lava tubes and basaltic caves as astrobiological targets on Earth and Mars: A review." *Planetary and Space Science* (Elsevier Ltd.) 58, no. 4 (March 2010): 592-598.
- Li, Stan Z. *Markov Random Field Modeling in Image Analysis*. 3rd. Springer, 2009.

- Litwiller, Dave. "CCD vs. CMOS: Facts and Fiction." *Photonics Spectra* (Laurin Publishing Co. Inc.), January 2001.
- Lowe, D. G. "Object recognition from local scale-invariant features." *International Conference on Computer Vision*. 1999. 1150-1157.
- Mallick, S., T. Zickler, D. Kriegman, and P. Belhumeur. "Beyond Lambert: Reconstructing Specular Surfaces Using Color." *International Conference on Computer Vision and Pattern Recognition*. 2005.
- McEwen, Alfred, and Sarah Mattson. "DTEEC_023531_1840_023953_1840_A01." *MRO MARS HIGH RESOLUTION IMAGING SCIENCE EXPERIMENT DTM V1.0*. 2013.
- Mei, C, and P Rives. "Single View Point Omnidirectional Camera Calibration from Planar Grids." *International Conference on Robotics and Automation*. 2007.
- Minogue, Ethan. "Avionics for Hibernation and Recovery on Planetary Surfaces." *NASA NSTRF Final Report*, 2012.
- Mitchell, K L, and M Malaska. "Karst on Titan." *First International Planetary Caves Workshop*. Carlsbad, NM, 2011.
- Moratto, Z M, M J Broxton, R A Beyer, M Lundy, and K Husmann. "Ames Stereo Pipeline, NASA's Open Source Automated Stereogrammetry Software." *Lunar and Planetary Science Conference*. 2010.
- Morris, Aaron, David Ferguson, David Silver, and Scott Thayer. "Topological Exploration of Subterranean Environments." *Journal of Field Robotics*, 2006: 395-415.
- Nenas, Issa, Pablo Abad-Manterola, Jeffrey A Edlund, and Joel W Burdick. "Axel Mobility Platform for Steep Terrain Excursions and Sampling on Planetary Surfaces." *IEEE Aerospace Conference*. 2008.
- Nutecher, A, S Augustin, H Surmann, K Lingemann, and J Hertzberg. "6D SLAM with an application in Autonomous Mine Mapping." *International Conference on Robotics and Automation*. 2004.
- Oberbeck, V R, W L Quaide, and R Greeley. "On the origin of Lunar sinuous rilles." *Modern Geology*, 1969: 75-80.
- Peterson, K M, H L Jones, and W L Whittaker. "Mission Design for Combined Lander-Rover Modeling of a Skylight." *First International Planetary Caves Workshop*. Carlsbad, NM, 2011.
- Pfister, H., M. Zwicker, J. Baar, and M. Gross. "Surfels: Surface Elements as Rendering Primitives." *Proceedings of IEEE SIGGRAPH*. 2000.
- Pirjanian, Paolo, et al. "Distributed control for a modular, reconfigurable cliff robot." *International Conference on Robotics and Automation*. 2002.
- Powell, Devin. "Roaming robot may explore mysterious Moon caverns." *Nature*, November 2012.
- Prinz, Martin. "The Great Rift and King's Bowl Lava Field, Snake River Plain, Idaho." *Annals of the New York Academy of Sciences* (Blackwell Publishing Ltd) 163, no. 1 (1969): 90-93.
- Robinson, M S, et al. "Lunar Reconnaissance Orbiter Camera (LROC) Instrument Overview." *Space Science Review* 150 (2010): 81-124.

Robinson, M.S., et al. "Confirmation of sublunarean voids and thin layering in mare deposits." *Planetary and Space Science* 69, no. 1 (2012).

Robinson, Mark. "WAC_GLD100_E300S0450_100M." *LRO MOON LROC 5 RDR V1.0*. 2011.

Schenker, P S, et al. "Lightweight rovers for Mars science exploration and sample return." *Proc. SPIE 3208, Intelligent Robots and Computer Vision XVI: Algorithms, Techniques, Active Vision, and Materials Handling*. 1997.

Scholten, F, et al. "GLD100: The near-global lunar 100 m raster DTM from LROC WAC stereo image da." *Journal of Geophysical Research* 117, no. E00H17 (2012).

Skonieczny, Krzysztof. "Lightweight Robotic Excavation." PhD Thesis, Robotics Institute, Carnegie Mellon University, 2013.

Snaveley, N, S M Seitz, and R Szeliski. "Modeling the World from Internet Photo Collections." *International Journal of Computer Vision*, 2007.

Snaveley, N, S M Seitz, and R Szeliski. "Photo Tourism: Exploring image collections in 3D." *ACM Transactions on Graphics (Proceedings of SIGGRAPH 2006)*, 2006.

Spenko, M. J., et al. "Biologically inspired climbing with a hexapedal robot." *Journal of Field Robotics*, April 2008.

Thrun, Sebastian. "Robotic mapping: a survey." In *Exploring artificial intelligence in the new millennium*, edited by Gerhard Lakemeyer and Bernhard Nebel, 1-35. 2003.

Thrun, Sebastian, et al. "Autonomous Exploration and Mapping of Abandoned Mines." *IEEE Robotics and Automation Magazine*, December 2004: 79-91.

Torres-Mendez, L., and G. Dudek. "Inter-Image Statistics for 3D Environment Modeling." *International Journal of Computer Vision (IJCV)*, 2008.

Triggs, Bill, Philip F McLauchlan, Richard I Hartley, and Andrew W Fitzgibbon. "Bundle Adjustment - A Modern Synthesis." In *Vision Algorithms: Theory and Practice*, edited by Bill Triggs, Andrew Zisserman and Richard Szeliski, 298-372. Springer Berlin Heidelberg, 2000.

Wagner, Michael D., Dimitrios Apostolopoulos, Kimberley Shillcutt, Benjamin Shamah, Reid Simmons, and William "Red" Whittaker. "The Science Autonomy System of the Nomad Robot." *IEEE International Conference on Robotics and Automation*. Seoul, Korea, 2001.

Werker, J, S M Welch, S L Thompson, B Sprungman, V Hildreth-Werker, and R D Frederick. "Extraterrestrial Caves: Science, Habitat, and Resources." NIAC Phase I study, 2001.

Wettergreen, David, et al. "Second Experiments in the Robotic Investigation of Life in the Atacama Desert in Chile." *International Symposium on Artificial Intelligence, Robotics and Automation in Space*. Munich, Germany, 2005.

Wong, U, B Garney, C Whittaker, and W Whittaker. "Image-Directed Sampling for Geometric Modeling of Lunar Terrain." *International Conference on Field and Service Robotics*. 2012.



Wong, U, et al. "Comparative Evaluation of Range Sensing Technologies for Underground Void Modeling." *Proc. Intelligent Robotics and Systems (IROS)*. 2011.

Wong, Uland. "Camera and LIDAR Fusion for Mapping of Actively Illuminated Subterranean Voids." *Conference on Field and Service Robotics (FSR)*. 2009.

—. *Lumenhancement: Exploiting Appearance for Planetary Modeling*. PhD Thesis, Robotics Institute, Carnegie Mellon University, 2012.

Worthington, P. "Re-illuminating single images using Albedo estimation." *Pattern Recognition*, 2005.

Wu, Changchang. "Towards Linear-time Incremental Structure From Motion." *3DV*. 2013.

Zhang, Zhengyou. "Iterative Point Matching for Registration of Free-Form Curves and Surfaces." *International Journal of Computer Vision* 13, no. 2 (October 1994): 119-152.

Zlot, Robert, and Michael Bosse. "Efficient Large-Scale 3D Mobile Mapping and Surface Reconstruction of an Underground Mine." *Field and Service Robotics*. Matsushima, Japan, 2012.

APPENDIX A: ANALOG FIELD TESTING SITES

This section describes the field test sites used during this research and references the experiments performed in each analog.

MINE PITS

Mine pits were used for experiments related to the pit descent feasibility study (Section 11.1.1 & 11.1.3) and demonstration of pit descent and modeling methods (Sections 9.4.1 & 9.4.2). Several pits in active open-pit coalmines were used as analogs for planetary pits in this study. The mine pits were within a few hours drive of the project team’s primary work location in Pittsburgh, so tests could be conducted in single day trips enabling more tests to be conducted than would otherwise have been possible.

The processing of open-pit coal mining is to dig to create a high wall that exposes seam(s) of coal; use explosive charges to break up the high wall; dig out the boulders created by blasting (creating a new high wall); and back-fill the blasted dirt (with coal removed). This creates a pit that moves continuously over the life of the mine. The high wall, with its layers of coal and other rock material, provides a good analog for layered walls in planetary pits. Additionally, coal has similar albedo to lunar regolith, producing accurate reflectivity. The fact that these pits were in active mines ensured that vegetation would not have time to grow, and meant that a pit would not remain the same for multiple visits.

These mine pits served as analogs for testing pit access and modeling technologies. There were no associated tunnels in which to test subsurface exploration and modeling.



Figure 155: Robotic exploration and modeling tests in a mine pit. Yellow streamers mark the edges of the portion of the mine’s high wall that was the focus of modeling for this experiment.

PLUTO'S CAVE

Pluto's Cave was used for an experiment in subsurface modeling described in Section 9.5.1. Pluto's Cave is a lava tube in Siskiyou, CA that served as an analog for lava tubes on the Moon and Mars. The site has several features observed on the Moon: a steep walled pit, a pit with a ramp, and a natural bridge between two pits. It also has sections of lava tube tunnel. While not yet observed on the Moon or Mars, lava tube tunnels are believed to exist on both planetary bodies. The accessible extent of lava tube cave is approximately 366 meters long. This lava tube is from an older lava flow (approximately 190,000 years old), so sediment has collected on the surface around the tube and some has been blown into the tube, giving it a sandy floor in areas not covered by rubble. This makes the tunnel floors more relevant to Mars than the Moon. The tunnel walls and ceilings are relevant to both bodies. Another result of this feature's age is the growth of vegetation on the surface and on pit floors, which is not planetary-relevant.

Pluto's cave served as an analog for testing subsurface exploration. The site is open to the public with very little restriction, and there is graffiti on the walls. This did not pose a problem for the testing conducted at the site for this project, but it would make the site less suited for high-resolution wall modeling.

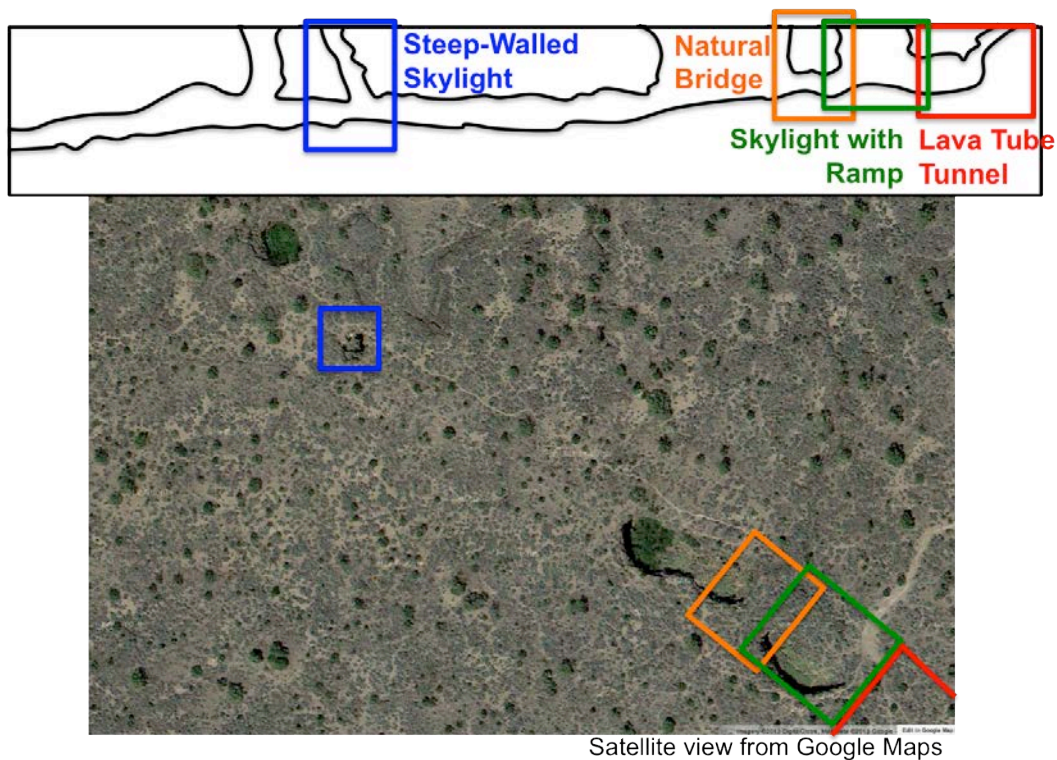


Figure 156: Top: Map of Pluto's Cave with features of interest marked. Bottom: Features marked in satellite view.



Figure 157: View from inside a collapse pit in Pluto's Cave. A natural bridge between two pits is visible in the center of the image. To the left, a rocky, vertical wall can be observed. To the right a ramp leads down from the surface. Considerable vegetation has grown on the surface and in the pits of Pluto's cave.



Figure 158: View from inside a tunnel segment in Pluto's Cave. Where not buried in rubble, the floor is covered with sediment.

KING'S BOWL PIT

King's Bowl pit was used for an experiment in surface reconnaissance and modeling (Section 9.3.2) and data from that same test was used to demonstrate point cloud rendering for visualization (10.1). It is part of the Craters of the Moon National Monument and Preserve is an approximately 76m long by 30m wide by 30m deep pit along Idaho's great rift. It sits in a 2.6 km² lava field. A steam explosion along the volcanic rift formed the pit (Prinz 1969). This is more relevant for Mars than the Moon. It is unknown whether sufficient water would have been present during Martian volcanic activity to cause steam explosions, but pits associated with rift structures have been identified on Mars (G. E. Cushing 2012). Whether or not the pit is an analog in terms of morphology and formation mechanisms, its layered, blocky walls made of volcanic rock and its rubble-covered floor made it a good functional analog for modeling pits on both the Moon and Mars. The walls and floor are primarily barren, though some vegetation made the site slightly less planetary-relevant.

King's Bowl served as an analog for surface reconnaissance of a pit. There is a cave entrance at one end of the pit that can be accessed by walking trail from the surface. While the trail is not planetary-relevant, future testing for a Tyrobot-like access system could be conducted at this site.



Figure 159: FARO instrument used for ground truth data collection at the edge of King's Bowl pit.

INDIAN TUNNEL

Indian Tunnel was used for an experiment in surface reconnaissance and modeling (Section 9.3.1.1) as well as for experiments in subsurface modeling (Sections 9.5.2, 9.5.3, and 9.5.4). Indian Tunnel is part of the Craters of the Moon National Monument and Preserve in Idaho and is a lava tube cave with multiple skylight pits. The cave is approximately 244 meters long. This lava tube tunnel is in a relatively recent lava flow (approximately 2000 years). The terrain is still quite barren and the lava tube floor is not covered in sediment.

Indian Tunnel served as an analog for both subsurface exploration and modeling (tests inside the cave) and surface reconnaissance and modeling (tests from outside the cave at collapse pit locations).

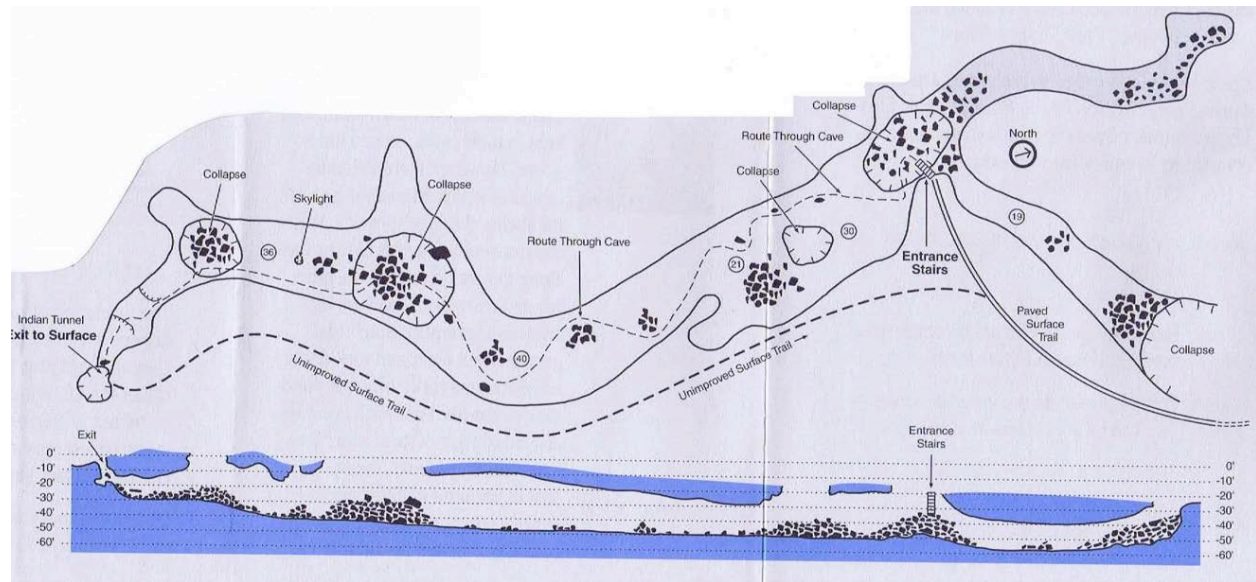


Figure 160: Map of Indian Tunnel cave.



Figure 161: View from inside the cave at Indian Tunnel, looking back at the entrance.



Figure 162: Rough lava rock is still apparent on the floor of Indian Tunnel in areas not covered in rubble. While this surface is traversable by a robot with sufficient terrainability, it is not covered in smooth layers of sediment, as seen in older lava tubes on Earth.



Figure 163: Conducting experiments from the rim of a skylight pit at Indian Tunnel

APPENDIX B: CRATERS OF THE MOON FIELD DEMONSTRATION AND SURVEYING

This section presents a description of fieldwork performed at Craters of the Moon, National Monument & Preserve, Idaho from July 30 through August 5. The two areas modeled are the Kings Bowl and Indian Tunnel. King's Bowl is a phreatic explosion pit 280 feet (90 m) long, 100 feet (30 m) wide, and 100 feet (30 m) deep, caused by lava meeting groundwater and producing a steam explosion 2,200 years ago. Indian Tunnel is the largest of the lava tubes in the caves in the park. It is over 800 feet long, up to 50 feet wide and 30 feet high, and it has several openings to the surface. Indian Tunnel and the other caves were formed by the Blue Dragon Flow 2,100 years ago.

KING'S BOWL



Figure 164: Kings Bowl

Kings Bowl pit modeling was conducted entirely from the surface rim, since it was intended as an analog for surface reconnaissance and modeling during a Scout or Wayfarer mission (see Sections 7.1 and 7.2). Data collected at King's Bowl included camera images, 3D LIDAR scans, and survey measurements.



LIDAR scans in 3 dimensions, with corresponding photographic data, were captured using a FARO Focus^{3D} X330 at several locations along the rim in different lighting conditions as shown in Table 4.

Table 4: Scan locations and lighting conditions

Date	Time of Day	Location	Scans
Wednesday July 30, 2014	Afternoon	West Rim (1W-14W)	14 (#40 – #53)
Thursday July 31, 2014	Morning	East Rim (1E-16E)	16 (#54 - #69)
Friday Aug 1, 2014	Morning	West Rim (1W-14W)	14 (#71 -#84)

The 44 scans were aligned and matched into 1 contiguous model using Autodesk ReCap software. By conducting scanning at different times of the day and from opposite sides of the pit the imagery gathered from the FARO Focus provides well-illuminated images for colorizing the scan points. The final model contains colorized points in UTM world coordinates.

Each matched scan location in an arbitrary coordinate frame was extracted from the ReCap software. During the field test, 30 scan locations were surveyed using a reflective 360° survey prism and a 3arc-sec survey instrument. The survey points are used as a check for the quality of the scan matching from the model.

Survey points were transformed into the model coordinate frame using a least square transformation with the following statistics:

- common points 1W and 16E used for transformation
- 1.1mm base error
- 90.748m baseline
- 0.9999987 scale factor

Comparison of the 30 transformed survey points to model points statistics:

- 8mm average horizontal delta
- 40mm maximum horizontal delta
- 55mm average elevation delta
- 128 mm maximum elevation delta

As the statistics show the horizontal data was evaluated standalone from the vertical data. The separation of the horizontal data from the vertical for analysis is due to much better quality of horizontal data from the survey data and output of the scan matching software. Note the survey and model points are well within expected tolerances for an outdoor environment.

The FARO Focus also gathers GPS data however this data is only good to within about 1 meter and is much coarser than the survey or model in quality. Idaho State University Department of Geosciences conducted research at Kings Bowl simultaneously with Carnegie Mellon. Their research team was able to provide 4 highly accurate Differential GPS points in world coordinates that tied to our control survey.



With these DGPS points the horizontal survey and model data is transformed into world Zone 12 UTM coordinates with the following matching statistics.

- 2 common points were used for transformation
- 0.7mm base error
- 17.897m baseline
- 1.000039 scale factor

One elevation point referenced from the DGPS data provided was used to correct the model and survey elevation data.

While 3D LIDAR data similar to that collected in this test could feasibly be collected by a surface rover during a pit exploration mission, GPS is not available beyond Earth. Survey measurements could be collected between multiple surface robots, but may or may not be used during pit exploration. The primary reason for building a LIDAR model using survey and GPS data was because this is the best model we can construct of the pit, and it can thus be used as “ground truth” for camera modeling.

Additional Survey Data Collected at Kings Bowl

Imagery experiments were also conducted as part of the fieldwork. Large flat washers with surveyor flagging were evenly distributed on the rim at 14 locations on the West and 16 locations on the East of the Kings Bowl pit. These physical markers provided reference locations for the imagery experiments and scan locations over the 3 days. As with the scan locations many of the imagery experiment locations were surveyed, providing information on the time of day and detailed coordinates for the camera position relative to the model.

The survey instrument is robotic and has automated tracking capability where range and angles to a survey prism are collected as the prism moves. On Thursday, July 31, perimeter walks were conducted tracing out the edge of the rim on both the East and West sides of the pit. This perimeter data served as an additional visual check for the model and as a reference while reviewing data from the field.



INDIAN TUNNEL

Indian Tunnel is the largest of the lava tubes in the Caves area of Craters of the Moon. The 800 ft tunnel runs North / South. Modeling data was collected underground in the tunnel and on the surface using a FARO Focus^{3D} X330 laser scanner.

Date	Time of Day	Location / Purpose	Scans
Saturday Aug 2	Afternoon	Ranger Orientation to Caves	
Sunday Aug 3	Night - Dawn	Underground / Long Tube Run	8 (#86 - #94)
Sunday Aug 3	Day	Surface / Collapse Areas North of Tunnel	3 (#96 -#98)
Sunday Aug 3	Day	Surface / Collapse Entries North	8 (#99 -#106)
Sunday Aug 3	Day	Surface / Collapse Entries South	8 (#107-#114)
Monday Aug 4	Night - Dawn	Underground / Hi Res Wall	3 (#115-#118)
Monday Aug 4	Night - Dawn	Underground / Long Tube Run	3 (#119-#121)
Monday Aug 4	Night - Dawn	Underground / Long Tube Run Fill	3 (#122-#126)
Monday Aug 4	Day	Underground/ Field Team Scan	1 (#128)
Monday Aug 4	Day	Underground/ Fill to Stair Entry	3 (#129-#131)
Tuesday Aug 5	Day	Presentation Data Transfer	

Scans collected were matched to create models using both ReCap and FARO Scene software. 4 groups of models were created for Indian Tunnel by scan matching adjacent scans:

- Surface North
- Underground from stair entrance to the large collapse zone
- Surface South
- High Resolution Wall

Additionally a fifth model was created by combining Surface North, Underground, and Surface South. The FARO GPS data from the surface scan points is used to align the large 5th composite model to UTM coordinates.

For night operations underground, FARO scanner imaging was lighted from LED rings appended to the scanner. This provided color data to overlay and colorize the laser point data. Long tube run data was a thread of 12 scans through the widest part of the tunnel starting at the large collapse in the South and extending North approximately 100 ft. Long tube run fill collected occluded areas of the long tube run to supplement the model. Results for the High Resolution wall experiment can be found in Section 9.5.4.