

Robotic Perception for Exploring Lunar Poles

Varsha Kumar

CMU-CS-22-112

May 2022

Computer Science Department
School of Computer Science
Carnegie Mellon University
Pittsburgh, PA 15213

Thesis Committee:

William “Red” Whittaker, Chair

David Wettergreen

*Submitted in partial fulfillment of the requirements
for the degree of Master of Science in Computer Science.*

Copyright © 2022 **Varsha Kumar**

Keywords: Rover, Micro-Rover, Lunar Pole, Stereo Vision, Autonomy, Illumination, Cameras, HDR

ABSTRACT

No rover to date has traveled to or imaged the lunar poles. The lack of atmosphere on the Moon and low elevation of the sun at the lunar poles creates regions of brilliant illumination, absolutely black shadow, and scenes containing both extremes. Rover perception at the lunar poles must create accurate terrain models in these unprecedented circumstances. The research investigates stereo geometry, algorithms, lensing, exposures, material appearance, and terrain illumination unique to lunar poles. Additional considerations include processing on space-relevant computing and camera hardware and achieving the short cycle times essential for achieving continuous rover motion. This research formulates, implements, and evaluates an illuminated stereo system and software specialized for lunar polar perception and deployment to the Moon. The resulting perception system will guide Carnegie Mellon's autonomous MoonRanger micro-rover that will map ice on the lunar south pole in 2023.





Figure A.1. The Extremes of Illumination. *Top Left: Lunar analog terrain illuminated at lunar polar sunlight intensity. Top Right: The same terrain illuminated by rover-powered laser dot projectors. Bottom: Associated depth models. Lunar polar sunlight is three orders of magnitude brighter than rover-powered illumination.*

The perception system generated in this work is being incorporated into MoonRanger, compelling distinct simplicity in design. Key features include eight infrared laser dot projectors to illuminate terrain, a fixed lens aperture size of F/4.0, a 16 millisecond exposure, and the semi-global matching stereo algorithm that enables continuous imaging and stereo depth perception on lunar polar terrain without switching camera hardware or software when imaging different extreme illumination conditions. The research additionally suggests an extension to the flight perception system that leverages High Dynamic Range techniques within the 16 millisecond exposure to minimize image whiteout in edge-case roving conditions.

Figure A.2. Comparison of Depth Model in Darkness to Depth Model in Sunlight. *The depth model of the region of terrain illuminated by rover-powered laser dot projectors is nearly identical to the corresponding section of the "sunlight" depth model (Avg. difference < 3 mm).*



“We choose to go to the Moon...” -JFK

TABLE OF CONTENTS

1. Introduction	5
1.1 AUTONOMOUS ROVERS AND MICRO-ROVERS	5
1.2 CAMERA APERTURE AND EXPOSURE TIME	6
1.3 SUNLIGHT AT THE LUNAR POLES	6
1.4 NON-VIABLE ALTERNATE SOLUTIONS	9
1.5 SCOPE OF RESEARCH	9
2. Thesis Statement	10
3. Priors	11
3.1 PRIOR LUNAR ROVER PERCEPTION TECHNOLOGY	11
3.2 PRIOR MARS ROVER PERCEPTION TECHNOLOGY	12
3.3 PRIOR EARTH PERCEPTION TECHNOLOGY	13
4. Rover Lens and Exposure Selection	14
4.1 EXPERIMENT TO DETERMINE APERTURE SIZE AND EXPOSURE TIME	14
4.2 LENS APERTURE ANALYSIS AND SELECTION	16
4.3 EXPOSURE TIME ANALYSIS AND SELECTION	17
5. Illuminating Terrain	20
5.1 IDEAL ILLUMINATION	20
5.2 LASER DOTS VS. LED	20
5.3 DETERMINING LASER DOT POSITIONS	21
5.4 VERGENCE FOR MULTIPLE PROJECTORS	23
5.5 INCREASING THE NUMBER OF PROJECTORS	24
6. Stereo Algorithm Selection and Implementation	26
6.1 STEREO ALGORITHM BACKGROUND	26
6.2 STEREO ALGORITHM SELECTION	27
6.3 ARCHITECTURE OF FLIGHT CODE STEREO IMPLEMENTATION FOR MOONRANGER LUNAR POLAR MICRO-ROVER	30
6.4 UNIT TESTING FLIGHT CODE STEREO IMPLEMENTATION FOR MOONRANGER LUNAR POLAR MICRO-ROVER	33
6.5 ACCURACY EVALUATION AND RUNNING LIVE	36
7. Space-Readiness and Verification	38
7.1 QUALITY AFTER LAUNCH VIBRATION	38
7.2 FUNCTIONALITY IN VACUUM	41
7.3 FUNCTIONALITY IN SPACE THERMAL CONDITIONS	43
7.4 FUNCTIONALITY ON DARK REGOLITH	45
8. High Dynamic Range for Continuous Polar Roving	48
8.1 SIMULATING SUNLIGHT AT DIFFERENT HEADINGS	48

8.2 HEADING LIMITS WITHOUT SHORT-EXPOSURE HDR	49
8.3 SHORT-EXPOSURE HDR	49
8.4 SHORT-EXPOSURE HDR MAINTAINS DEPTH MODEL CORRECTNESS IN NON-WHITEOUT IMAGES	52
9. Conclusion	54
9.1 SUMMARY	54
9.2 REGARDING THE RESEARCH'S PERCEPTION SYSTEM	54
9.3 REGARDING ROBOTIC PERCEPTION FOR EXPLORING LUNAR POLES	55
10. Acknowledgements	57
11. References	58

1. Introduction

Rover perception of lunar polar terrain must image regions of intensely brilliant illumination, absolutely black darkness, and both extremes simultaneously. Traditional perception methodologies for such high-contrast environments, such as terrestrial high dynamic range (HDR) [1], require stationary imaging of terrain at various exposures, precluding continuous rover motion. Instead, a combination of many laser dot projectors, a lens aperture size large enough to accommodate a low exposure time, and the semi-global matching stereo algorithm enables continuous autonomous roving.




Figure 1.1. The Extremes of Illumination. Left: Terrain illuminated at lunar polar sunlight intensity. Right: The same terrain illuminated by rover laser dot projectors. Sunlight at the lunar pole is three orders of magnitude brighter than the active illumination of the rover.

1.1 AUTONOMOUS ROVERS AND MICRO-ROVERS

Roving autonomously and utilizing micro-rovers are two strategies to maximize the planetary surface explored over the duration and cost of a mission. Autonomous roving removes humans from control cycles, eliminating hours spent in communication delays where a rover sits in place waiting for humans to receive data, process it, and send the next command. The most value from autonomy comes when a rover can continuously navigate terrain without stopping to do its own processing. Micro-rovers are orders of magnitude less mass than typical rovers (Figure 1.2), which in turn makes them orders of magnitude less expensive to launch. It takes less fuel to launch a light micro-rover than a larger rover, and a micro-rover can be added onto a mission rather than designed as the primary purpose. Though this research's results and conclusions are fundamental in their own right and many applicable to lunar polar roving in general, the investigation heavily considers autonomous micro-roving so that the results can be directly applied to Carnegie Mellon's MoonRanger, an autonomous micro-rover that will map ice at the South Pole of the Moon in 2023.




Figure 1.2: Micro-Rover vs. Rover Size [2]. Left to Right: Sojourner, MoonRanger, Spirit/Opportunity, Curiosity/Perseverance. Micro-rovers (MoonRanger, Sojourner) are orders of magnitude lighter and smaller than car-sized rovers such as Spirit and Perseverance. As such, simplicity in micro-rover design is both valuable and necessary because they cannot carry heavy hardware.

1.2 CAMERA APERTURE AND EXPOSURE TIME

To enable a human to see in both bright light and in a dark room, the pupil contracts or dilates to control the amount of light entering the eye. To enable an amateur photographer to image the outdoors both at noon and at night, a digital camera includes an auto-exposure feature to control the duration for which light is captured in an image. Both a smaller pupil and shorter exposure time regulate the brightness of a scene whereas a larger pupil or longer exposure time enhance vision in the dark, making a scene appear brighter. In a DSLR camera, the camera “pupil,” known as an aperture, can be adjusted by changing the lens [3]. The resulting brightness of a pixel is an integration of the magnitude of light let in by the lens aperture over the duration governed by the exposure time. To perceive lunar terrain, rover camera aperture size and image exposure time must be such that pixels are neither too dark nor too bright to interpret.

Setting aperture size and exposure time for rover cameras is more complex than for the human eye, digital cameras, or DSLR cameras. To physically change “pupil” size on a camera lens requires either sensitive hardware that can be damaged by the violent shaking that is experienced during launch [5] or complex hardware too heavy for a micro-rover to carry [4]. Adjusting exposure like a digital camera would be a reasonable solution if a rover were to stop to take an image of long enough exposure to perceive in darkness. However, the ambition of a rover in a short mission is to maintain continuous motion in order to maximize its average driving velocity and total exploration distance.

1.3 SUNLIGHT AT THE LUNAR POLES

Lunar polar rover cameras must image terrain in brilliant sunlight that is unlike anything on Earth despite the immense range of scene lighting and the inability to utilize commonplace terrestrial solutions for controlling the brightness of an image. The absence of an atmosphere on the Moon means that light is neither absorbed nor diffused before it reaches the lunar surface, so sunlight on the Moon is approximately 30% brighter than sunlight on the earth [6]. The brightness perceived by a camera is the result not only of illumination by the sun or rover, but also the albedo of the terrain and relative incidence angles of sun, terrain slope and camera [7]. These albedo and incidence matters are quantified in the Bidirectional Reflectance Distribution Function (BRDF) [8]. The albedo of lunar surface material (regolith) is low due to its high absorptivity, making rover-illuminated imaging in darkness even more challenging.

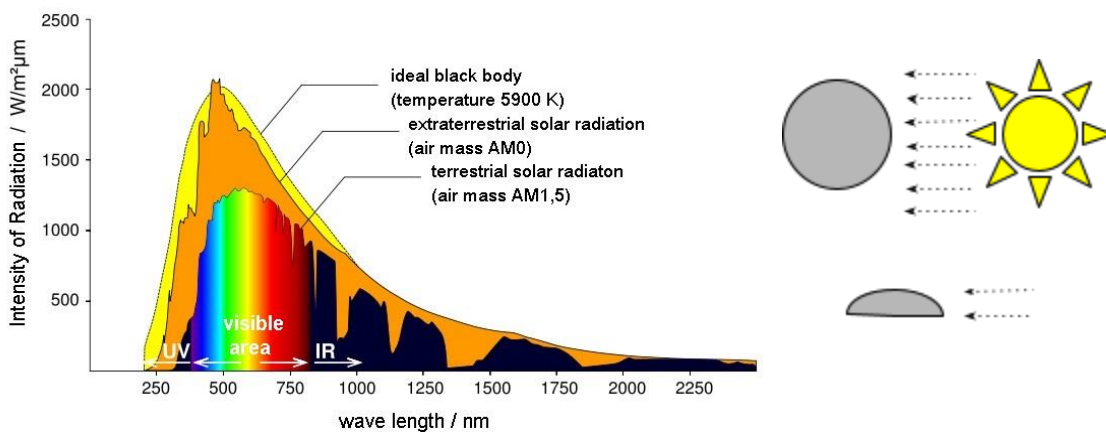


Figure 1.3. Brilliant, Grazing Sunlight. Left [6]: Earth’s atmosphere absorbs approximately 30% [6] of the sunlight that reaches the Earth, reducing the portion that reaches the surface (colored and black envelope). Note that reductions in the infrared spectrum (Black) are substantial at some wavelengths. The absence of an atmosphere on the Moon means that no radiation is absorbed prior to impinging the lunar surface (orange envelope). The yellow envelope is the idealized black body spectrum at 5900K that most matches the Sun’s actual spectrum (orange). Right: At the lunar pole, sunlight reaches the surface at a low azimuth, clocking around the pole rather than rising overhead. Incoming lunar polar sunlight is 162,720 lux in magnitude whereas rover-powered illumination is a few hundred lux in magnitude [9]. For perspective, the sometimes-blinding change in lux due to headlights of an oncoming car at night on Earth is only a few hundred lux [10].

The sun when viewed from the lunar pole is at a low angle above the horizon. This low-angle lighting combined with the non-diffuse nature of sunlight in a vacuum creates long, sharp, absolutely black shadows from hills or rocks on the terrain [11] that can only be illuminated by rover-powered light sources. Even small pebbles on terrain create pockets of shadow, so even in sunlight, images exhibit a high-contrast mix of brightly lit regions and pockets of black shadow. A lunar polar rover must create depth models of terrain in these mixed-lighting conditions and must drive perpendicular to the sun so as to not be blinded by the sun or discern no features due

to the opposition effect. The opposition effect is the appearance of whiteout due to camera orientation directly opposite to the bearing of the sun. When this occurs, all shadows are occluded in the camera's view by the rocks that create them, so the image contains no perceivable shadowing, is entirely lit, and is devoid of the texture and visual featuring that otherwise makes it useful for rover perception.




Figure 2.4. Low-Angle Sunlight [12]. Low-angle sunlight (even lower at a lunar pole) creates long, dark, absolutely black shadows from obstacles on terrain. The bi-directional distribution function (BRDF) of a surface, which determines how much light is seen on a surface, is a function of material, angle of incidence of light, and viewing angle. The albedo of most lunar regolith and rocks is reasonably similar, so large differences in the appearance of dark or light objects and regions are mostly a function of the illumination or relative incidences within the BRDF equation. Bright objects like the camera-pointing face of the large rock are oriented directly at the camera, so the incidence is primarily at play in the high BRDF. The intermediate darkness of the horizontal terrain is primarily due to its more acute orientation to the camera in the BRDF equation. The black appearance of the shadowed regions is due to the near-zero magnitude of the illumination magnitude that reaches the shadowed regions in the BRDF.



Figure 1.4: Opposition Effect. The opposition effect creates a region of no shadow on terrain directly opposite a light source [13][14]. This occurs when the bearing of the camera is 180 degrees opposite the light source (in this case the sun). The appearance of whiteout is due to a maximum of BRDF combined with lack of visible shadowing that otherwise introduces discernable feature and texture. This region cannot be interpreted and appears as a featureless

locus of bright terrain. Left: Simulated lunar polar terrain with apparent opposition effect in center of image [15]. Right: Image cropped to region directly surrounding opposition effect.

1.4 NON-VIABLE ALTERNATE SOLUTIONS

Imaging terrain in these unprecedented circumstances of overwhelmingly brilliant, low-angle sunlight and in near-absolute darkness without adjusting aperture defies traditional approaches like changing filters, HDR, or adjusting exposure. Although Spirit and Opportunity changed filters via a motorized “filter wheel” for spectral imaging studies [16], the approach was non-viable for autonomous micro-rover relevance for two reasons. (1) The mechatronic volume, mass and complexity of a filter wheel is a non-starter for a micro-rover, and multiple stereo cameras would require multiple of these. (2) Some autonomous policy would be required to make decisions of which filter to utilize under what lighting/terrain conditions. Development of such a policy is non-viable as it requires high confidence in the lighting characterization algorithm applied to an environment that nobody in the world has imaged to date as well as sufficient rover development timeline.

HDR or adjustable exposure is only effective if the long exposure time needed to image terrain in darkness does not introduce motion blur to images acquired while the rover is moving. Moreover, HDR is non-viable since it requires that all images for an HDR calculation be acquired from the same vantage point with respect to a scene. Since the images are acquired sequentially at different times while a rover is moving, the image series is necessarily from different locations which violates the fundamental assumption of HDR. HDR would be the solution of choice for historical rovers that stopped between incremental steps, but is a non-starter for a rover that is in continuous motion.

1.5 SCOPE OF RESEARCH

The research determines that a single lens aperture, exposure time, and source of illumination in darkness is suitable for perceiving lunar terrain in scenes of brilliant sunlight and absolutely black shadows characteristic of lunar poles. Considerations, design, implementation, and testing of illumination source and geometry are discussed. Perception research agenda additionally selects one stereo algorithm to generate viable depth models in all illumination circumstances and implements rigorous stereo flight code. The resulting perception system was verified as space-ready.

The research determined which headings relative to the sun cause image whiteout, and hence are nonviable headings for robot bearings. High Dynamic Range (HDR) using only short exposures mitigated image whiteout, enabling calculation of depth models at headings otherwise precluded by whiteout.

2. Thesis Statement

This research investigates whether a perception system using a single lens aperture, exposure time, source of illumination in darkness, and stereo algorithm could be suitable for perceiving lunar terrain in scenes of brilliant sunlight and absolutely black shadows characteristic of lunar poles and implements the results to create space-ready perception system.

The research addresses stereo algorithms, lensing, exposures, material appearance, and terrain illumination unique to lunar poles. Additional considerations include processing on space-relevant computing, achieving the short cycle times essential for achieving continuous rover motion, space-readiness verification of camera hardware, and algorithm implementation as rigorous stereo flight code. Whereas the research is fundamental in its own right, the results will additionally guide Carnegie Mellon's autonomous MoonRanger micro-rover as it maps ice near the lunar south pole.

3. Priors

Perception technology has been used on the Moon, Earth, and Mars. Most planetary rovers have carried stereo cameras to the Moon and Mars, and perception systems that illuminate scenes are ubiquitous in Earth robotics. However, rover perception technology has yet to be deployed at the lunar poles, and no prior perception technology exists that is both suitable for perceiving lunar polar terrain and simple enough to use on micro-rovers.




Figure 3.1. Semi-Autonomous Rovers. Left [17]: VIPER, a large rover with stereo cameras and flood lighting mounted at the top of a mast. Images are processed on Earth, then a plan for traversing the visible terrain is sent to VIPER to follow. Right [18]: Perseverance, with stereo cameras for navigation (NavCams) mounted on a mast. This large rover can navigate terrain autonomously for short distances and is also monitored by Earth.

3.1 PRIOR LUNAR ROVER PERCEPTION TECHNOLOGY

Lunar rovers that carried/will carry perception systems include NASA’s Volatiles Investigating Polar Exploration Rover (VIPER), ISRO’s Pragyan, the Yutu and Yutu-2 rovers, and the Lunokhods. Each of these rovers relied/will rely heavily on human guidance, and VIPER is the only rover to explore the lunar poles.

VIPER, a large lunar rover launching in 2023, uses stereo cameras and powerful flood lights mounted high above the rover on a mast to perceive lunar polar terrain (NavCams, Figure 3.1) [17]. The stereo cameras enable depth perception, and the flood lights enable visibility in shadow. Images are sent back to Earth, processing and planning occur on Earth, and the rover follows the plan for about 25 feet. The rover follows the plan blindly except for monitoring the wheels via HazCams (Figure 3.1) [19]. Whereas the mast-mounted camera system and powerful flood lights are a beautifully engineered system for a large rover at the lunar poles, such a system is too heavy and draws too much power to be carried by a micro-rover. Further, VIPER had the

advantage of building a rover system around illumination, whereas the MoonRanger illumination system, discussed in this work, was added onto main rover design. As such, MoonRanger illumination had to be as simple as possible and significantly simpler and lighter than VIPER's.

Pragyan, a micro-rover slightly larger than MoonRanger launched in 2019 and launching again in 2022, has a pair of stereo cameras mounted to the top of its rover body. Its path planning occurs on Earth as well. While Pragyan suggests a micro-rover could use stereo cameras to perceive terrain, its perception system does not carry illumination to perceive terrain in darkness and does not generate depth models for live, on-rover planning. [20][21][22]

Yutu and Yutu-2 (2010s), as well as the Lunokhods (1970s), also carried stereo camera systems, but were teleoperated from Earth. The Lunokhods especially were directly joysticked by human drivers. [23][24][25]

3.2 PRIOR MARS ROVER PERCEPTION TECHNOLOGY

Mars rovers include NASA's Sojourner (1997) [26], NASA's Mars Exploration Rovers (Spirit and Opportunity, 2004-2018)[16], the Zhurong rover (2021-Present)[27], and NASA's Curiosity (2012-Present) and Perseverance (2021-Present)[28][18][29] rovers. These rovers used stereo cameras to perceive terrain, though processing and command generation primarily occurred on Earth. Two exceptions were Sojourner's step-stop-drive hazard avoidance and Perseverance's short-distance autonomous navigation. The Mars rovers have never addressed perceiving terrain in both brilliant illumination and darkness, as is characteristic of the lunar poles.




Figure 3.2. Mars Perception Systems. Top: Pancam [16], a system of changing a filter in front of a lens, much like wearing sunglasses, flew as part of NASA's Mars Exploration Rover cameras. Such a system could not be used on MoonRanger due to development timeline. Bottom: A set of laser line stripes and stereo cameras that flew on Sojourner [26] for hazard avoidance. A three-dimensional contour was triangulated for each stripe, and the combination of stripes was analyzed as a "bumper" for hazard avoidance. This research investigates illuminating an entire region of terrain in front of a rover rather than a few strips.

The Mars Exploration Rovers' (MER) Pancam filter wheel and Sojourner's laser lines could have inspired alternate approaches to perceiving in brilliant sunlight and in darkness (Figure 3.2), but were dismissed due to mass constraints and an insufficiently large region of illumination.

The MER Pancam included a filter wheel, a series of colored filters in front of a camera lens much like sunglasses. Operators commanded different filters depending on the portion of the light spectrum they needed to image [16]. Whereas a “sunglasses” system could address perception in brilliant lunar polar sunlight, the additional mass and development time for an algorithm that would correctly determine when to “wear sunglasses” precluded using such a system on the short-development-timeline MoonRanger rover.

Sojourner projected five laser stripes on terrain in front of the rover. Using triangulation, it calculated the three-dimensional contours along these lines and used the set of contours as a “bumper” for autonomous hazard avoidance [26]. Though groundbreakingly successful in space, laser line stripes only illuminate a small portion of the terrain in front of a rover. To enable calculating a more complete model of terrain in front of a rover, this research investigates illuminating an entire region of terrain in front of a rover rather than a few strips.

3.3 PRIOR EARTH PERCEPTION TECHNOLOGY

Non-space technology such as High Dynamic Range and depth cameras are used to perceive in various illumination scenarios or to perceive depth.

High Dynamic Range (HDR), which merges images taken at many exposures [1], would be a perfect solution to perceive terrain in brilliant sunlight and in darkness weakly illuminated by a rover. However, it would require a rover to be stationary as images ranging from less than a millisecond of exposure to hundreds of milliseconds of exposure are taken. This would force a rover to stop every control cycle, drastically reducing exploration speeds. For a week-long lunar polar micro-rover mission, a reduction in rover exploration speed is a critical detriment to mission scope.

Depth cameras on Earth, such as the Intel Realsense [30], are commonly used in warehouse robotics and on drones. These cameras project an array of dots and use an image and the structure of dots to calculate depth. Such cameras are not yet verified to survive the harsh environment of space. Other perception systems such as driver-assisting cameras on a car or LIDAR on an autonomous vehicle might be too heavy to carry on a micro-rover or are also not yet verified to survive space.

4. Rover Lens and Exposure Selection

Lunar polar rovers must image and create accurate depth models in brilliant illumination, absolutely black shadow, and scenes containing both extremes. Short exposures and smaller apertures are used in bright lighting to eliminate whiteout, whereas long exposures and large apertures are necessary for any visibility in the dark. For terrestrial environments that exhibit a mix of these lighting conditions, adaptive exposure and apertures are the primary technologies. However, applying adaptive exposure and dynamic aperture size technology to moon roving is not straightforward for even large, heavy and high-budget rovers [16], let alone micro-rovers. For a micro-rover like MoonRanger, constrained to orders of magnitude smaller mass and budget, adaptive exposure and aperture size is entirely infeasible. The need is to use a single lens aperture, exposure time, and source of illumination in darkness that, though not optimal for any circumstance, is “good enough” for most and simple enough for use on a micro-rover.

4.1 EXPERIMENT TO DETERMINE APERTURE SIZE AND EXPOSURE TIME



Figure 4.1. Simulating Lunar Lighting on Lunar Terrain. Arrimax 18/12 cinematic light [31] projecting 18 kilowatts of “sunlight” onto MoonRanger’s perception testbed.

Images of lunar analog terrain were analyzed to determine which lens aperture size and camera exposure time enable best overall visibility across scenes ranging from brilliant illumination to darkness. Images were taken using each of two apertures, $f/5.6$ and $f/4.0$, and at exposures ranging from 1 to 200 milliseconds, both in light and in darkness. To simulate sunlight at the

lunar poles, an ArriMax 18/12 [31] cinematic day light cast 16720 lux [9] on the lunar analog terrain. A Hydrargyrum Medium-Arc Iodide (HMI) bulb was used to approximate the spectral distribution of sunlight [9]. In darkness, infrared laser dot projectors cast a few hundred lux onto the terrain.

The terrain for this investigation included two craters, a hill, and rocks covered in gray Portland Cement powder. Portland Cement mimics the optical properties of lunar regolith at lunar polar latitudes [32]. Rocks and wooden blocks placed at half-meter increments served both as positive obstacles and key points for rapid analysis of regions of visibility. The objective in aperture and exposure selection was to maximize area of interest for rover navigation discernable by cameras in all lighting circumstances. A direct measurement of this is the distance at which terrain at least as wide as a rover is visible in a generated depth model. The blocks and rocks were placed along the line created if a rover were to drive forward in a straight line, marking the bounds of the critical region of interest for rover navigation.




Figure 4.2: Lunar Analog Terrain Incorporating Craters, Hill, and Rocks. Blocks (boxed in cyan) placed on the lunar analog terrain at 0.5 meter intervals facilitated rapid analysis of depth models to measure the area of visibility. The blocks were placed at rover width on the left, and because the camera and illumination system is symmetric, the line of blocks marks the area of interest for rover navigation. The depth model above exhibits exceptional visibility in sunlight for a small rover, a viewing horizon of over 3 meters. Additional features of the terrain include a ruler to mark the centerline of the stereo pair, a crater in the front right and one far away directly ahead, and a hill and rocks on the right.

4.2 LENS APERTURE ANALYSIS AND SELECTION

A larger lens aperture size admits more light, enabling image capture at shorter exposure times. This decreases the chance of motion blur in images. With a smaller aperture, shorter exposures can render images too dim to analyze. By increasing lens aperture size, the camera equivalent of pupil dilation, the instantaneous illumination captured by a camera sensor is increased by allowing additional light to pass through the lens. Because pixel brightness is an integration of instantaneous illumination over a duration of exposure, increasing instantaneous illumination by increasing aperture size negates the dimming effects of shorter exposure times.

Two potential problems created by increasing aperture size are shrinking depth-of-field and oversaturation of images. A larger aperture size typically decreases depth-of-field, the distances at which terrain begins to appear blurry rather than sharp. Too large an aperture size can additionally let in too much light, rendering images captured, at any exposure, oversaturated or too bright to interpret. The oversaturated regions appear completely white in an image, exhibiting “whiteout”. For the perception system developed in this work, depth-of-field effects and whiteout are shown to be non-issue below.

Lens Aperture Size	f/5.6 (smaller)	f/4.0 (larger)
Depth-of-Field Start	24.3 cm	31.2 cm
Depth-of-Field End	Infinity	Infinity

Figure 4.4: Lenses and Depth-of-Field Calculations. [33] Both lenses have infinite depth-of-field in the far field, so objects far from the rover are still in sharp focus. For the larger aperture lens, the region of sharpness begins at 31 cm away from the cameras rather than the 24 cm of the smaller aperture lens. However, because the depth model begins at 28 cm in front of the rover (Figure 4.5), at most the first three centimeters of the depth model are generated from blurry images. This small section would have been modeled in the exact previous set of images [34], so the slightly smaller depth-of-field of the f/4.0 lenses is a non-issue.




Figure 4.5: Closest Point in Depth Model. The closest point of the depth model is at 27.9 cm. For an f/4.0 lens, this means that only the first three centimeters of the depth model might be

generated from blurry images. These first three centimeters, less than 2% of the model, are modeled during previous time steps [34] as a rover navigates, and exhibit no visible degradation in quality. As such, the three centimeter smaller depth-of-field is a non-issue.

Depth-of-field was calculated for both the f/5.6, smaller aperture, lens and the f/4.0, larger aperture, lens. The f/4.0 has a smaller depth-of-field by only the first three centimeters of the depth model (Figure 4.4). Because the first three centimeters are less than 2% of the depth model, are modeled during previous time steps [22] as a rover navigates, and exhibit no visible degradation in quality, the three centimeter smaller depth-of-field is a non-issue.

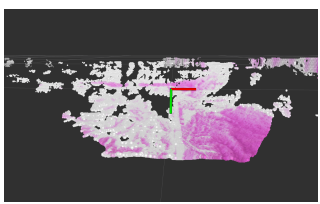
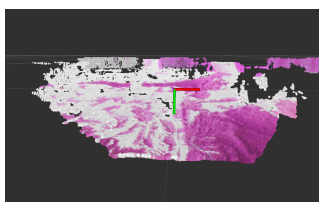
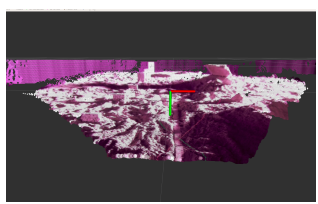
Exposure	100 ms	50 ms	10 ms
Depth Model			

Figure 4.6: Exposure Controls Whiteout in f/4.0 Lenses. Depth models of the terrain in Figure 4.2 generated from images taken at three different exposures are shown above. Regions of color are points in the depth model, while gray marks “no solution”. Whiteout causes regions of the 100 ms depth model to be empty because no reasonable interpretation was possible. Lower exposure times mitigate and eliminate whiteout.

The second concern with the larger aperture was that all images would exhibit whiteout due to too much light being allowed through the lens. However, this is also not an issue because exposures of 10 milliseconds and below using the f/4.0 lens exhibit no whiteout (Figure 4.6).

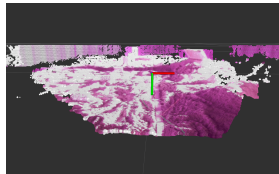
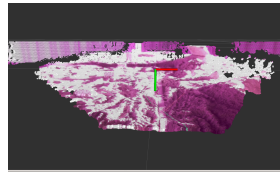
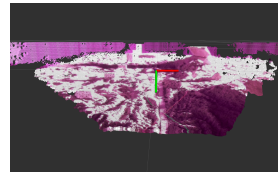
Result 4.1: After calculations and testing, neither depth-of-field shrinkage nor constant whiteout are issues with the larger f/4.0 aperture. Therefore, the f/4.0 aperture can be used to take bright enough images using short exposures that minimize motion blur.

4.3 EXPOSURE TIME ANALYSIS AND SELECTION

For the f/4.0 aperture lens, an exposure time that optimally minimizes whiteout in sunlight and maximizes visibility in darkness was selected. However, shorter exposure times minimize whiteout while longer exposure times enable further visibility of terrain in darkness illuminated by the weak rover-powered laser dot projectors. The exposure selected could therefore not be optimal for either case, but was instead “good enough” for both. The selection was conducted by

first eliminating exposures that exhibit excessive whiteout, then selecting the exposure that maximizes visibility in darkness.

In simulated lunar polar sunlight, exposures above 25 milliseconds exhibited significant whiteout (gaps in top right of depth model, Figure 4.7), while exposures below 5 milliseconds were not sufficiently illuminated. Within the 5 to 25 millisecond range, a 5-10 millisecond exposure exhibited the least whiteout and a 16 millisecond exposure exhibited very minor whiteout.

Exposure	30 ms	25 ms	16 ms
Depth Model			

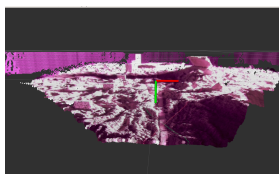
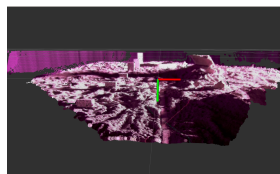
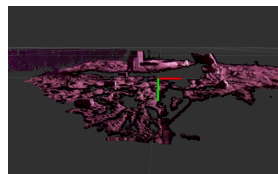
Exposure	10 ms	5 ms	1 ms
Depth Model			

Figure 4.7: Depth Models of Lunar Analog Terrain in Simulated Lunar Polar Sunlight.

Whiteout was minimal in depth models created from images at 5 to 25 milliseconds of exposure and least in the 5 to 10 millisecond exposure range. Below 5 milliseconds of exposure, regions of the depth model were missing due to insufficient brightness in the associated images. Note that 16 milliseconds of exposure time had very minor whiteout.

In darkness, laser dots enabled 0.5 meter of visibility at 5 milliseconds of exposure and 1 meter of visibility at 10 milliseconds of exposure (Figure 4.8). For context, the MoonRanger micro-rover is 0.64 meters long, so exposures in the 5 to 10 millisecond range have a viewing horizon slightly beyond a rover length ahead. The limited visibility is not conducive to pre-planning obstacle avoidance. The next longest exposure tested was 16 milliseconds, which has 1.5 meters of visibility in darkness over two rover lengths. This is a much further horizon for obstacle avoidance planning, for minor losses (Figure 4.7) from whiteout. Higher exposures up to 25 milliseconds have additional gains for similar losses from whiteout. However, because there is a higher confidence that exposures 16 milliseconds and below will not manifest motion blur in images [35], 16 milliseconds was selected as the longest reasonable exposure time.

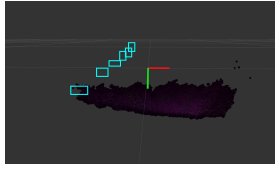
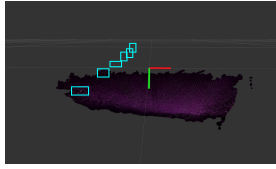
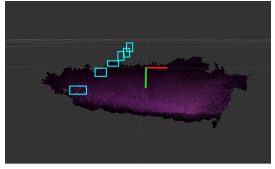
Exposure	5 ms	10 ms	16 ms
Depth Model			

Figure 4.8: Depth Models of Dark, Rover-Illuminated Terrain. The same terrain as Figure 4.7 was imaged in darkness illuminated by rover-powered laser dot projectors. Each teal box outlines a wooden block placed at 0.5 meter increments. At 5 milliseconds of exposure, the viewing horizon in darkness is 0.5 meters. At 16 milliseconds of exposure, the viewing horizon jumps to 1.5 meters.

Result 4.2: After depth model analysis of images in brilliant light and illuminated only by laser dots, a 16-millisecond exposure for the f/4.0 aperture lens was selected.

5. Illuminating Terrain

Lunar polar rovers encounter regions of absolutely black shadow. To perceive in this absolute darkness, a lunar polar rover must illuminate its terrain. Considerations, design, implementation, and testing of illumination for a small rover are discussed here.

5.1 IDEAL ILLUMINATION

The ideal illumination in darkness is super-powerful sunlight-intensity flood lighting projected on as wide of a region as far out in front of a rover as possible. Sun-intensity illumination in darkness could keep illumination constant across sunlit terrain and absolute darkness, precluding the need for careful camera hardware and exposure selection (Section 4). However, the only lights that could accomplish this would burn tens of kilowatts of power and mass over 60 kilograms [31]. Such a light would draw four orders of magnitude more power than a micro-rover can generate and be thrice the mass of an entire micro-rover, making it infeasible for micro-roving.

5.2 LASER DOTS VS. LED

Two sources of illumination were investigated, infrared laser dots [36] and infrared LEDs [37]. A dot projector and LED with the same electrical properties and same measured lux on terrain over a four square inch area were compared. Illumination from the dot projector generated a more complete disparity map of more area than that generated from illumination by the LED (Figure 5.2). Therefore, infrared laser dots were chosen as the source of illumination.



Figure 5.1. Scenes for Testing Illumination. A cement-powder scene with a small crater in front and a rock far away was used for illumination testing. Cement powder served as an optical

simulant similar to lunar regolith. A meter stick marked the center line between the cameras. Right: A lux meter taped on a vertical wooden board was placed on the scene during most tests.

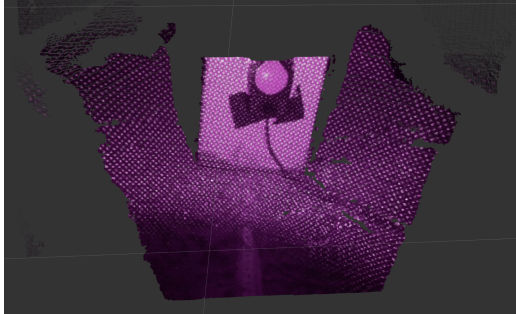
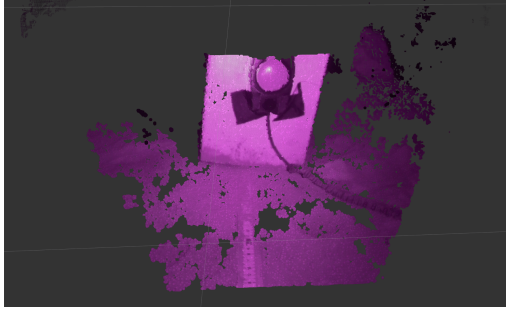
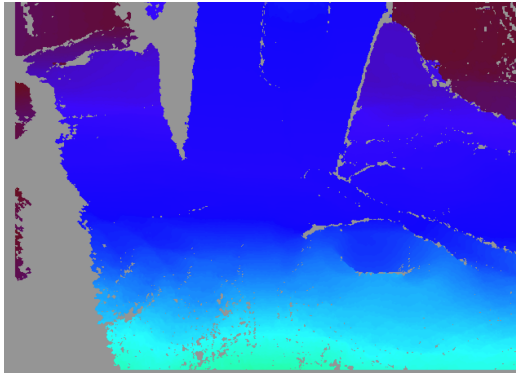
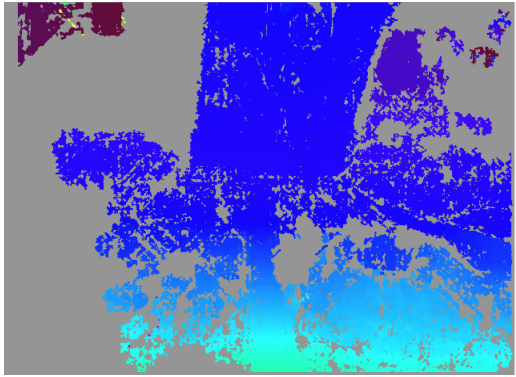
	INFRARED LASER DOTS	INFRARED LED
Depth Model		
Disparity Map		

Figure 5.2: Choosing Laser Dots (left) Instead of LED (right). *Depth models of a cement-powder-covered surface and a vertical board with a lux meter attached are pictured above. The depth model created using the laser dots (top left) is more complete than that created using the LED as illumination (top right). The disparity maps provide additional clarity that the depth model from laser dots is more complete, where gray is the absence of data in a depth model (bottom row).*

Result 5.1: Infrared laser dot projectors are superior sources of illumination for stereo reconstruction compared to wattage-equivalent LEDs.

5.3 DETERMINING LASER DOT POSITIONS

A laser dot projector could be placed above, in line with, or below cameras at any given tilt down from the horizon. Based on the noise in depth models created from placing a projector above, in line with, and below the cameras, placing the laser dot projector above the cameras is ideal. Further, experimenting with different angles relative to horizontal suggested a 10 degree tilt

down from horizontal for cameras that are angled 26 degrees downward. This 10 degree tilt minimized noise in the depth model.

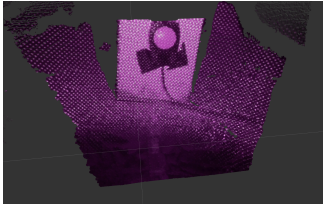
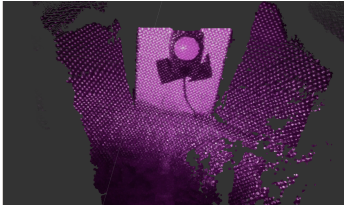
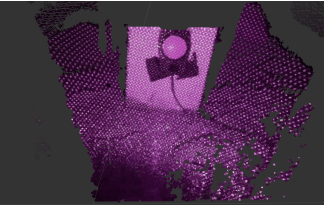
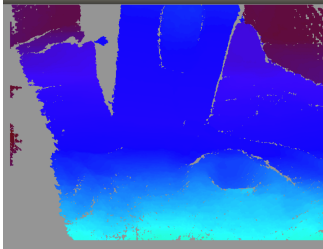
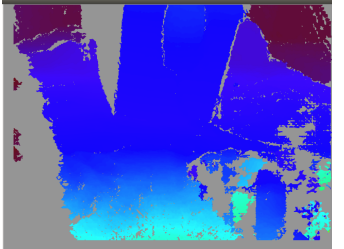
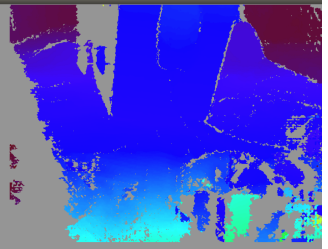
	Above	In Line With Cameras	Below
Depth Model			
Disparity Map			

Figure 5.3: Choosing Dot Projector Height. *The depth model created by placing the laser dots above the cameras lacks the noise observed in the bottom right of depth models and disparity maps generated when the projector is placed in line with or below the cameras. As such, placing the laser dot projector above the cameras is preferred.*

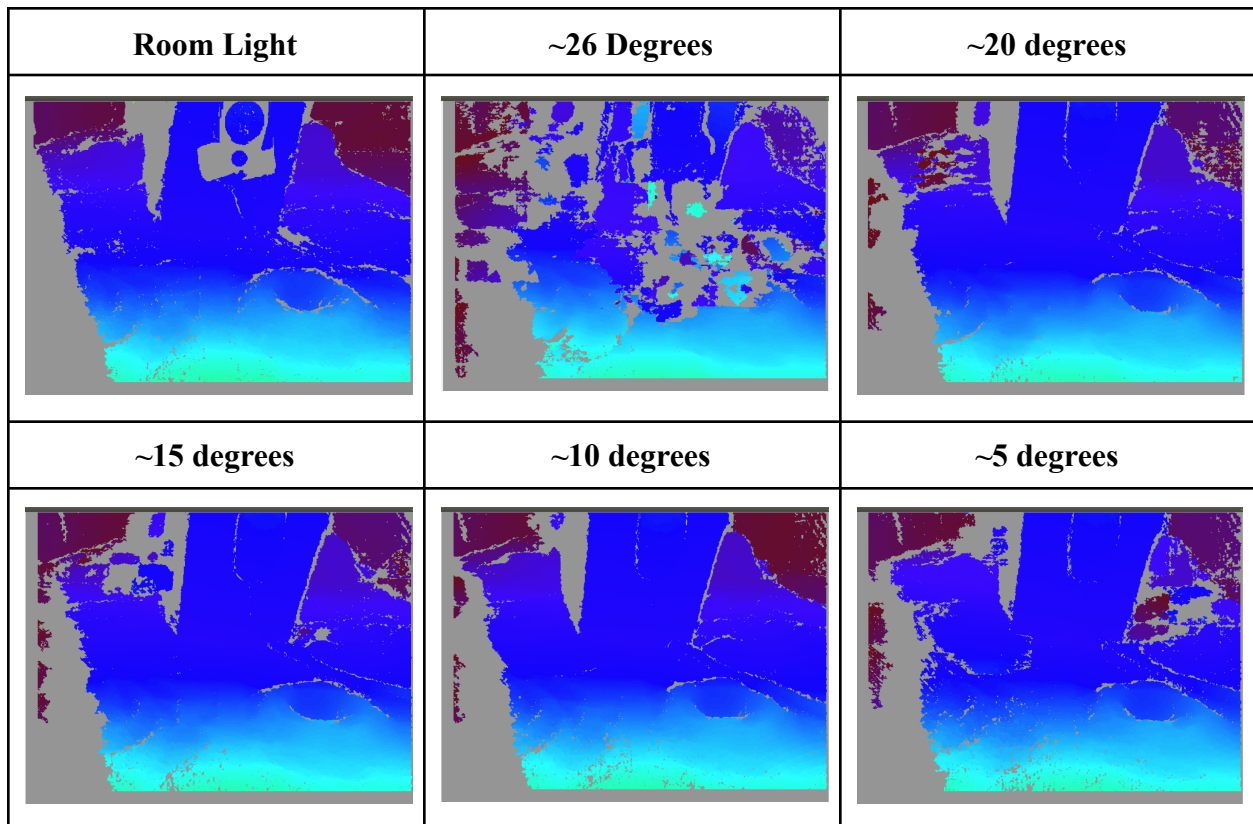


Figure 5.4: Choosing Laser Dot Tilt Downward. Disparity maps created by tilting the laser dot projectors down by various amounts were compared. The “Room Light” disparity map (top left), created from images illuminated by room lighting, serves as a “ground truth,” where the colors encode distance away from the cameras (darker is farther). A 10 degree tilt down from horizontal exhibits the least noise, mismatches between the depth encoded (color of a region) compared to the “Room Light” disparity map.

Result 5.2: Placing a laser dot projector above cameras at a 10 degree tilt downward generates the most complete and least noisy depth model.

5.4 VERGENCE FOR MULTIPLE PROJECTORS

To increase the illumination projected onto terrain, multiple projectors could be used such that the areas where they cast light overlap. This would additively increase the incident lux on terrain over that area. For symmetry, assume two sets of projectors. The lateral separation between projectors should be as small as possible to increase overlap. The divergence between the sets of projectors could be varied but should always be outward to maintain overlapping areas of illumination at distance. A 10 degree vergence was chosen to both maximize the completeness of

the generated depth model and compensate for a gap in illumination in the lower right of the image because of the internal roll angle of the dot projector.

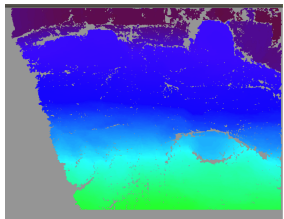
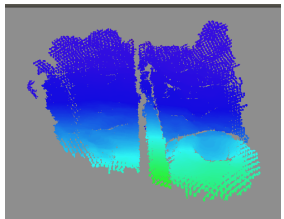
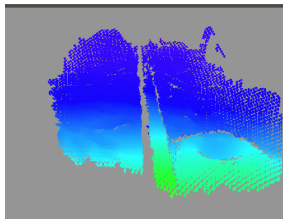
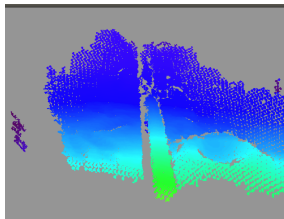
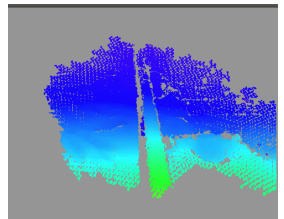
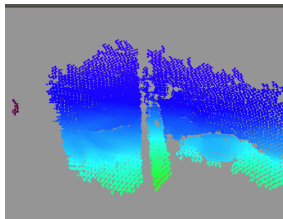
Outward Divergence	Room Lighting (200 ms Exposure)	0 Degrees (25 ms Exposure)	20 Degrees Outward (25 ms Exposure)
Disparity Map			
Vergence Per Projector Set <td>30 Degrees Outward (25 ms Exposure) <td>40 Degrees Outward (25 ms Exposure) <td>50 Degrees Outward (25 ms Exposure) </td></td></td>	30 Degrees Outward (25 ms Exposure) <td>40 Degrees Outward (25 ms Exposure) <td>50 Degrees Outward (25 ms Exposure) </td></td>	40 Degrees Outward (25 ms Exposure) <td>50 Degrees Outward (25 ms Exposure) </td>	50 Degrees Outward (25 ms Exposure)
Disparity Map			

Figure 5.5: Testing Laser Dot Divergence. Disparity map representations of depth models mark color where there is data in the depth model and gray where there is none. The “Room Light” disparity map serves as a “ground truth” of what a fully illuminated disparity map should look like. The remaining disparity maps cover smaller regions because they are generated from images taken at 25 milliseconds exposure time rather than 200 milliseconds. The most complete disparity map is the one without vergence. However, a 10 degree divergence was chosen to compensate for a gap in illumination (Figure 5.5).




Figure 5.6: Selecting a 10 Degree Divergence. Images without projector vergence (above) exhibit a dark spot in the bottom left corner of the image. This is because the projector is internally rolled by 15 degrees, so the square region of dots projected on terrain is actually at an angle. To “fill in” some of this dark area, the left set of projectors must be verged out relative to

the right, and to maintain symmetry, the right set of projectors are also verged out by the same amount. A 10 degree outward divergence was chosen as a small vergence that might fill this dark spot and not have the losses in depth model completeness seen in a 20 degree divergence (Figure 5.4).

Result 5.3: A 10 degree outward divergence between two sets of laser dot projector best illuminates terrain to maximize completeness of a depth model while also compensating for a projector-inherent gap in illumination.

5.5 INCREASING THE NUMBER OF PROJECTORS

To increase the illumination cast on terrain, two sets of laser dot projectors were considered above in 5.3. Maximizing the number of projectors, keeping the same number per set to maintain symmetry, maximizes total illumination cast. The limiting constraint on the number of projectors is total allowed power draw. Rover solar panels capture energy at a specific rate, generating a specific wattage of power available at any time to power the entire rover. As such, the maximum number of dot projectors is determined by power affordable in a “rover power budget” allocated to illumination. For MoonRanger, the maximum number of projectors is two sets of four projectors.


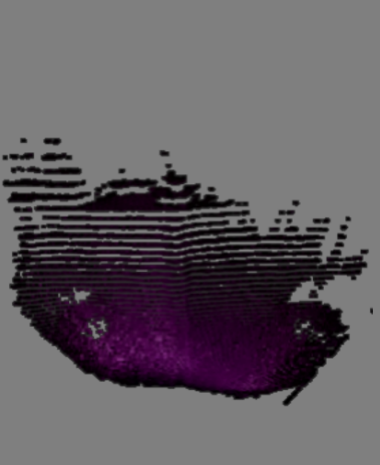

Scene	Depth Model From 6 Projectors	Depth Model From 8 Projectors
		

Figure 5.7. Increasing The Number of Dot Projectors. *A terrain (left) with two optical simulants including two craters and three sets of ridges was used to measure the increase in viewing horizon when increasing the number of projectors. The set of craters and each set of ridges were sculpted at approximately half-meter increments measured away from the rover. The*

brighter optical simulant represented slightly darker than expected optical properties of lunar polar regolith while the darker simulant represents the absolute darkest possible on the moon.

The two simulants are further discussed in Section 7. On both types of regolith, the viewing horizon using 8 projectors (right) is approximately 1.3 times longer than the viewing horizon using 6 projectors (center), improving the viewing horizons discussed in Section 4.

Result 5.4: Increasing the number of laser dot projectors to the maximum powerable under rover power budget maximizes illumination cast. For MoonRanger, this maximum is two sets of four projectors (though Section 4 used an earlier board with sets of three).

6. Stereo Algorithm Selection and Implementation

Three dimensional awareness of the surrounding environment is necessary to plan routes around obstacles and create maps. Lunar polar rovers are challenged to generate these depth models from images taken in brilliant sunlight and in darkness when faintly self-illuminating the terrain. In the case of the MoonRanger rover, the terrain is illuminated by laser dot projectors.

The canonical method for generating depth models in daylight is stereo vision. Traditionally, in structured lighting, such as illumination from laser dot projectors, depth model generation algorithms leverage the pattern of dots projected. However, the short development time of MoonRanger precluded incorporating the favorably patterned dots that would have afforded computational advantages. Also precluded was any intelligent determination of whether an image was sunlit or rover-illuminated to adjust the stereo algorithm. Perception research agenda related to selecting and implementing one stereo algorithm from the myriad of possibilities that would generate viable depth models in all illumination circumstances. The experimentation, selection and flight-code implementation of the stereo algorithm that will guide the MoonRanger lunar polar micro-rover is discussed here.

6.1 STEREO ALGORITHM BACKGROUND

A stereo algorithm is the computational equivalent to humans perceiving depth using two eyes. A stereo algorithm processes a pair of images and creates a three-dimensional depth model of the environment seen by the cameras.




Figure 6.1: Canonical Stereo Algorithm Pipeline. A stereo image is first rectified via defined matrix operations to computationally remove distortion and any imperfections in the horizontal alignment of the two cameras. Then, pixels in the left image are “matched” to the corresponding pixels in the right image by searching the right image for a correct match. For example, a “corner” of a rock in the left image would match to the same “corner” in the right image. However, because the left and right images are taken by cameras horizontally spaced apart by some baseline, the x-coordinates of the rock “corner” will be different in each image. The difference in x-coordinates of “matches” is known as disparity and is inversely proportional to

depth. Depth calculations “invert” the disparities to get depth from cameras, generating a set of points in three-dimensional space, known as a point cloud.

6.2 STEREO ALGORITHM SELECTION

Because rectification and depth calculations (Figure 6.1) are direct formulaic mathematical calculations, stereo algorithm selection focused instead on the disparities generated by various stereo matching algorithms. Five stereo matching algorithms were considered and libSGM, a GPU-accelerated implementation of the Semi-Global Matching (SGM) algorithm [38] was chosen.

The five stereo algorithms compared were: libSGM [38], OpenCV Semi-Global Block Matching (SGBM) [39], UT AMRL’s stereo_dense_reconstruction [40], a dynamic programming stereo algorithm [41], and NVIDIA Visual Programming Interface (VPI) Stereo [42]. The Robot Operating System (ROS) wrappers of libSGM and OpenCV SGBM were used and were included in the selection process because they were already implemented within MoonRanger’s prototype software. The UT AMRL stereo_dense_reconstruction and dynamic programming stereo algorithms were included based on the results of a trade study of open-source stereo algorithms in low-fidelity lunar polar sunlight [43]. The NVIDIA VPI algorithm was included because it is the native stereo algorithm on an NVIDIA TX2i GPU, which is MoonRanger’s autonomy computer [44]. The disparities generated by the five algorithms in sunlight (Figure 6.2) and in darkness (Figure 6.4) were used to determine which stereo algorithm would be optimal for depth perception in all lighting conditions.

(Continued on next page)

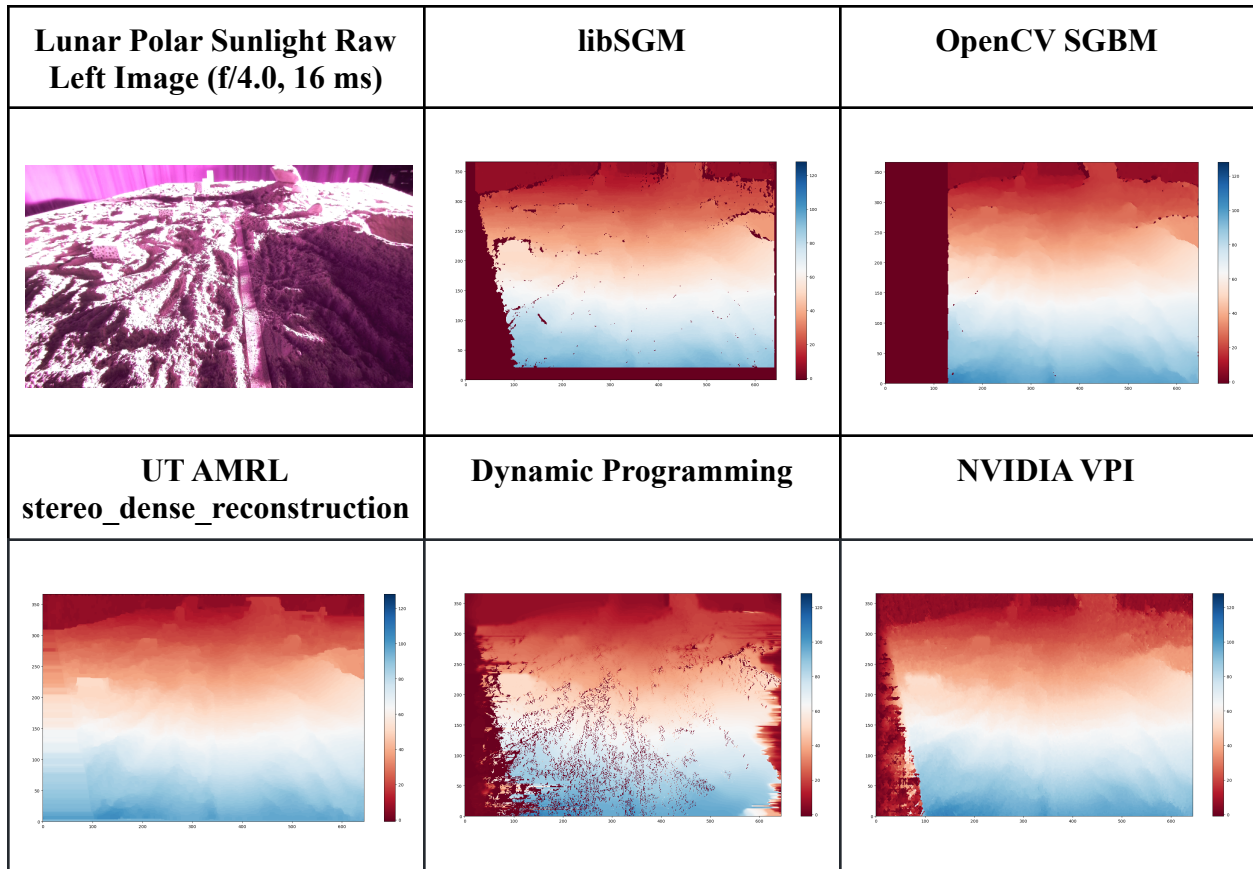


Figure 6.2: Comparison of Disparity Maps in Lunar Lighting. Images of lunar analog terrain in lunar polar sunlight (Section 4.0) were analyzed. For each stereo algorithm, a disparity map, or array of the disparities calculated per pixel of the left image, was generated. Pictured above, larger disparities, representing depths closer to cameras, are bluer. The darkest red represents the absence of a solution, such as the left edge of the disparity map, where the right camera's field of view does not overlap the left camera's field of view (Figure 6.3). Disparity maps generated were reasonable in that they had bluer (larger) disparities in the near field and redder (smaller) disparities in the far field. However, UT AMRL's stereo_dense_reconstruction (bottom left) and NVIDIA VPI's stereo algorithm (bottom right) calculated disparities along the left edge of the disparity map where there should be none. Additionally, the dynamic programming stereo algorithm (bottom center) marked "no solution" along the right edge of the disparity map where there should be disparities because the right camera's field of view covers that portion of the left camera's field of view (Figure 6.3).

Disparities generated by each stereo matching algorithm from images of lunar analog terrain in illumination approximating lunar polar sunlight (Section 4.0) were analyzed. Whereas disparities calculated were reasonable overall, three algorithms either calculated disparities when there should have been none or marked "no solution" in areas where there should be disparities (Figure 6.2). UT AMRL's stereo_dense_reconstruction and NVIDIA VPI's stereo algorithm calculated disparities along the left edge of disparity maps where there should be none because a

right camera's field of view cannot cover the left edge of a left camera's field of view (Figure 6.3). However, because the left edge could theoretically be masked out using known geometry, these two algorithms could still be considered. In contrast, the dynamic programming stereo algorithm marked "no solution" along the right edge of the disparity map where there should indeed be disparities because a right camera's field of view covers the right portion of a left camera's field of view (Figure 6.3). Because there is no simple and accurate solution for re-generating depths along this right edge, the Dynamic Programming algorithm should no longer be considered.




Figure 6.3. Overlapping Camera Field of View. Stereo cameras that are collinear and separated horizontally will have overlapping fields of view. The left camera's field of view is depicted in blue and the right camera's field of view is depicted in yellow (left). The overlapping field of view is green. Disparities, which are calculated from a "match" between pixels in the left and right images, can only be calculated in the green region. When seen in the left camera image (right figure), the green region does not reach the left edge of the image but does cover the right edge. As such, when displaying disparities as an array that corresponds to their pixel in the left image, disparities should have "no solution" along the left edge and should have solutions along the right edge.

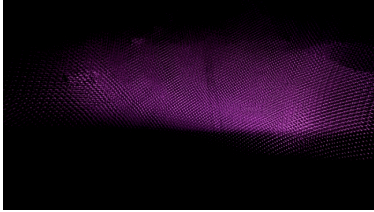
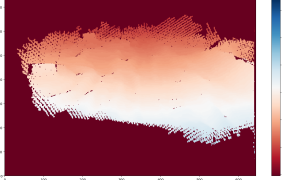
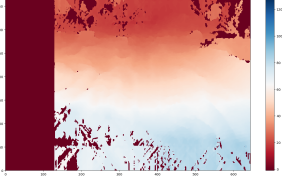
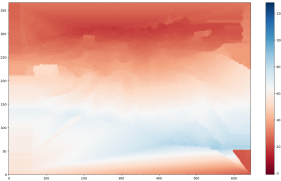
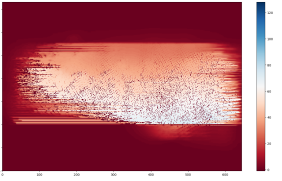
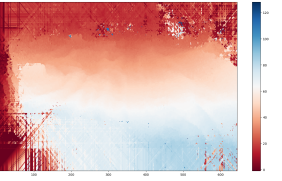
Laser Dots in Darkness Raw Left Image (f/4.0, 16 ms)	libSGM	OpenCV SGBM
		
UT AMRL stereo_dense_reconstruction	Dynamic Programming (For Completeness)	NVIDIA VPI
		

Figure 6.4. Comparison of Disparity Maps in Laser-Dot-Illumination. The terrain and position of cameras is identical to Figure 6.1, but the terrain was illuminated only by six laser dot projectors. Note that libSGM (top, center) correctly marked the lower region (black, not illuminated in the raw image) as having “no solution”, whereas the OpenCV, UT-AMRL stereo_dense_reconstruction, and NVIDIA VPI implementations (top right, bottom left, bottom right) solved for the lower region. These solutions were incorrect disparities, not evidence of a superior algorithm, because instead of large near-field disparities, which would be depicted by darker blue and represent depths closer to the cameras, the near-field was full of small disparities, depicted by the near-white of depths in mid-field or the red colors of the far-field.

Of the four algorithms left in consideration, OpenCV’s SGBM, UT-AMRL’s stereo_dense_reconstruction, and the NVIDIA VPI algorithms incorrectly marked near-field black pixels as being at mid-field or far-field depth rather than marking the region as having no solution due to insufficient illumination. In contrast, libSGM correctly marked that the black regions of the image could have no reasonable disparities calculated. As such, libSGM was chosen as the stereo matching algorithm for lunar polar micro-rovers.

Result 6.1: The libSGM stereo matching algorithm is best suited to depth model generation for lunar polar micro-roving.

6.3 ARCHITECTURE OF FLIGHT CODE STEREO IMPLEMENTATION FOR MOONRANGER LUNAR POLAR MICRO-ROVER

Although the algorithms profiled herein pertain to extreme lighting of all types, the implementation is distinct from a purely academic inquiry by manifesting it as rigorous flight code enabling navigation and autonomy on the MoonRanger lunar polar rover. For this reason, the implementation must meet a gold standard of fault tolerance, memory management, and readability. To accomplish this, the campaign to implement stereo for the flight rover included designing a modular the C++ implementation for use in the NASA Core Flight System (cFS) environment, careful consideration about memory use and segmentation fault prevention, and rigorous testing of both success and failure conditions. After multiple code reviews during development, the implementation was merged into the flight codebase and verified on a preflight rover.

The MoonRanger flight rover uses the NASA Core Flight Systems (cFS) as the environment for running all autonomy software. NASA cFS is a spaceflight-grade fault-tolerant message passing software and system executive, somewhat similar to the more common but less rigorous Robot Operating System (ROS). Major software components are written as cFS applications, classes that include specific initialization and message processing methods, as well as application-global data storage. The application created before this work contained a message-passing skeleton and rectification [45]. For modularity, this work architected the remainder of the stereo pipeline as a separate algorithm class (Figure 6.4, golden color) of which one instance was stored in the stereo cFS application's global data. The stereo cFS application (Figure 6.4, green) was then improved from a pipeline publishing placeholder data to a pipeline that called the algorithm class' functions to process a pair of stereo images.

(Continued on next page)

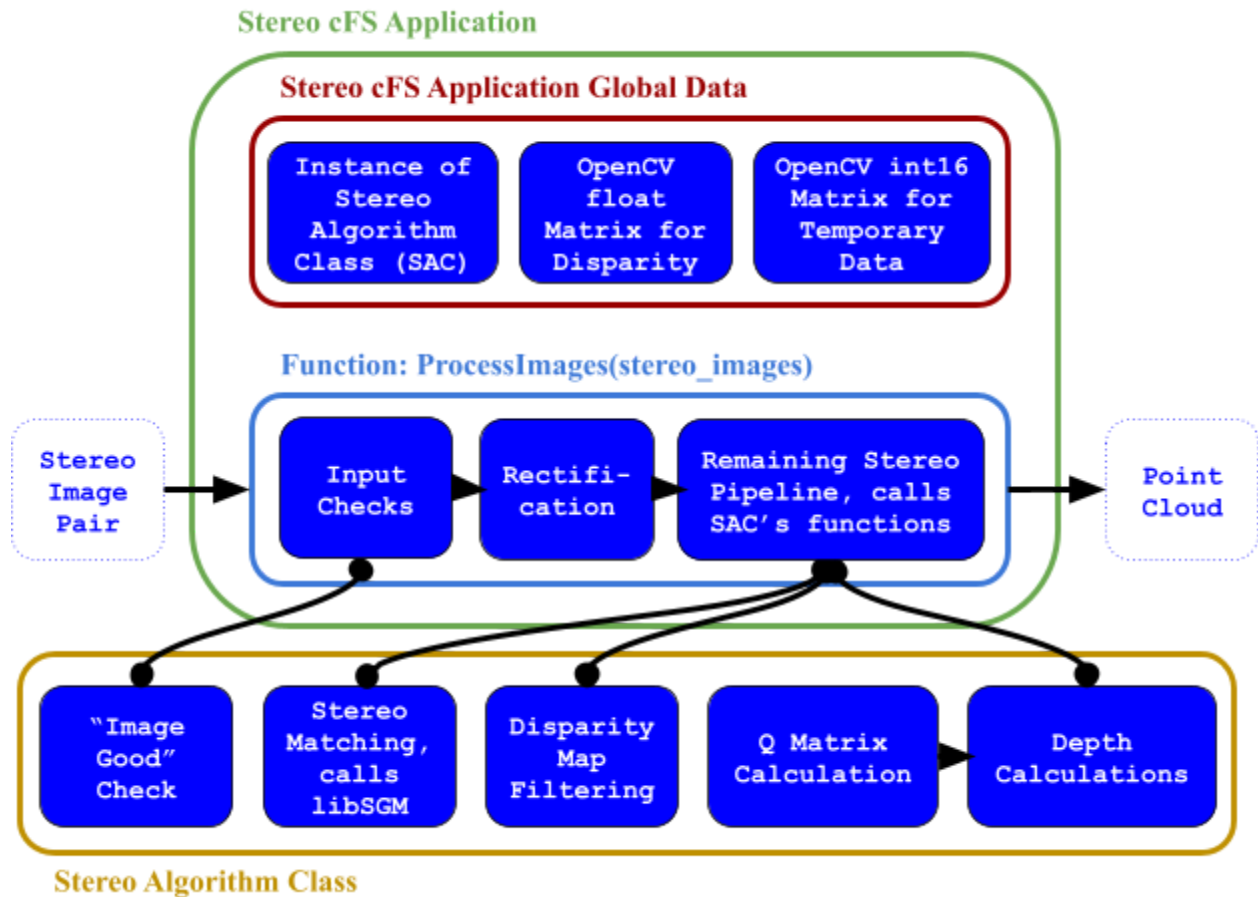


Figure 6.5. Flight Implementation Architecture. For modularity, stereo algorithm implementation was partitioned into a separate class from the main stereo cFS application. The application stored an instance of the stereo algorithm class, whose methods it called to run the complete stereo pipeline (Figure 6.1), as well as two OpenCV matrices to store intermediate results. By globally storing the disparity matrix in the cFS application, memory was reused when filtering a disparity map or processing later images. Similarly, any intermediate image processing occurred reusing the temporary data matrix. The memory reuse prevented memory fragmentation from allocating and deallocating memory. Within the Stereo Algorithm Class, additions to the canonical pipeline included filtering out noisy disparities, implemented by William Lee-Moore [47] then modified by this work. Additionally, calculating the Q matrix, a matrix of calibration parameters, from camera calibration matrices was separated out from depth calculations implemented originally by Amy Lin [46]. This work also augmented each method of the Stereo Algorithm Class to check input data and return an error flag instead of segfaulting on malformed inputs.

To prevent memory fragmentation from allocating and deallocating memory, the stereo cFS application stored globally two OpenCV matrices for storing intermediate results. One stored the float disparity values after stereo matching and disparity map filtering, and the other was used as

a location for temporary data storage within the stereo matching and disparity map filtering functions. By reusing the OpenCV matrices during these function calls and for every new pair of images, memory used for disparity calculation was only allocated once.

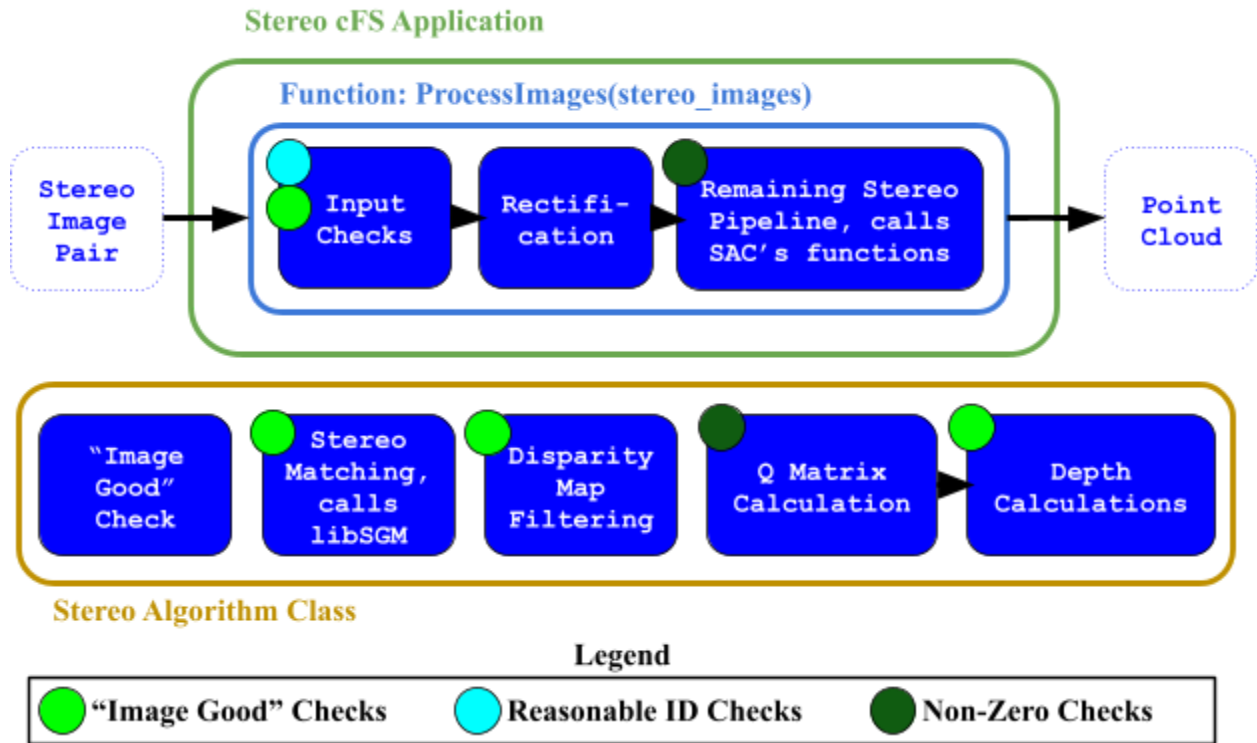


Figure 6.6. Input Checks to Prevent Segmentation Faults. To prevent segmentation faults that can occur by processing malformed inputs, checks are run on inputs at the beginning of every method implemented by this work. This work also added checks to the beginning of depth calculations and disparity filtering [46][47]. “Image Good” checks confirmed that image or disparity map OpenCV Matrices were non-empty, of the expected size, and single-channel or three-channel images. Reasonable ID checks confirmed that the two images in a stereo pair were taken at the same time. Non-Zero checks confirmed that certain indices of calibration matrices, which were used as divisors in division calculations, were non-zero to avoid faults caused by dividing by zero. When an input did not pass the checks, meaning it was malformed, the functions in the Stereo Algorithm Class returned failure instead of success, and the ProcessImage function in the stereo cFS application raised a failure to the cFS executive.

Segmentation faults were prevented by checking input before running the stereo pipeline in the stereo cFS application and at the beginning of each function in the stereo algorithm class (Figure 6.5). Upon receiving a new pair of images, the images were checked to have been captured at the same time and to be well-formed. “Image Good” checks performed upon the receipt of a new image pair and at the beginning of stereo matching checked that camera images were of the correct size and either 3-channel (color) or 1-channel (grayscale) images. “Image Good” checks

performed at the beginning of disparity map filtering and depth calculations confirmed that the input disparity maps were of the correct size. Non-zero checks were performed to ensure that calculations involving specific parameters from calibration matrices would not raise a “Divide by Zero” fault.

If a check failed, the methods in the stereo algorithm class returned failure and the image processing function in the stereo cFS application raised an error to the cFS executive. In this case, the stereo algorithm class’s functions did not calculate anything and did not alter their inputs, except for setting a Q matrix to an identity matrix.

Result 6.2: For modularity, the flight stereo system was partitioned into a main cFS stereo application that called helper functions and a stereo algorithm class storing these functions. Memory management and explicit error handling elevated the code to flight standard robustness and fault tolerance.

6.4 UNIT TESTING FLIGHT CODE STEREO IMPLEMENTATION FOR MOONRANGER LUNAR POLAR MICRO-ROVER

Unit and integration tests confirmed the robustness and accuracy of the implementation before it was merged, along with the test cases, into the flight code. Unit tests were created for each of the stereo algorithm class’s functions, and integration tests were created for testing the whole pipeline in the stereo cFS application. The test cases included testing that failure conditions raised errors instead of crashing and testing that non-failure runs had accurate output.

Both the class and the cFS Application were unit tested using images taken by the stereo cameras as input. Each method above, including error cases, was tested for accuracy and correct error case behavior. Disparity maps were compared pixel-by-pixel to those created by the ROS wrapper of libSGM. Point clouds were compared to what was declared as an acceptable ground truth.

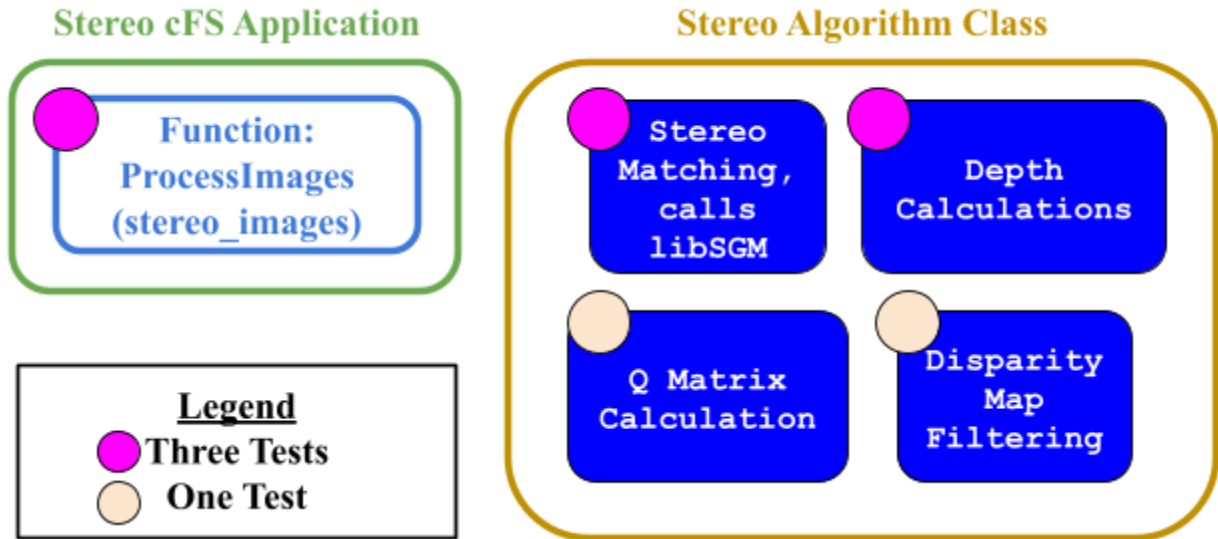


Figure 6.7. Non-Trivial Test Cases. The functions marked in pink tested on “Blocks World”, “Lunar Sunlight”, and “Dot-Illuminated” datasets (Figure 6.8), while the ones marked in beige were simpler functions tested only on “Blocks World” data. Tests on stereo matching compared calculated disparity maps with a ground truth from the ROS libSGM wrapper. Testing depth calculations and the ProcessImages function compared the calculated point clouds to “ground truth” (Figure 6.8). Because Q matrix calculation was a simple formula, its output was compared to a known Q matrix for the “Blocks World” camera calibration. Because disparity map filtering was also a simple, known set of operations [48], the test that a filtered disparity map was still similar to the original was only run on a “Blocks World” disparity map.



Figure 6.8. Test Case Datasets. Three datasets were used to evaluate the accuracy of the stereo pipeline, disparity maps generated by stereo matching, and point clouds generated by depth calculations. The three datasets each tested a unique environment. Left: The “Blocks World” dataset tests an environment of fixed geometry. Center: The “Lunar Sunlight” dataset comprised lunar analog terrain illuminated at lunar polar sunlight intensity approximating solar spectral distribution. Right: The “Dot-Illuminated” dataset comprised lunar analog terrain illuminated by laser dot projectors, what a rover would perceive in the shadows characteristic of the poles.

For disparity map and point cloud accuracy, the flight implementation was checked relative to the Robot Operating System (ROS) wrapper of the selected libSGM stereo algorithm [38]. Disparity maps were compared pixel-by-pixel to the ROS wrapper output and point clouds were

compared coordinate-wise to a “ground truth.” A complication in testing disparity map accuracy was transforming “ground truth” ROS disparity output, which was a ROS message [49], to a format readable by the test case. A CSV file representing the image was used as the fastest implementable method of converting from a ROS disparity message [50] to an “image file.” A complication in testing point cloud accuracy was that the order of points in the ROS point cloud differed from that of the flight implementation. To remedy this, the output from a run of the flight implementation was stored into a CSV file and compared to the “ground truth” point cloud in CloudCompare, a graphical point cloud analysis software [51]. If the flight code output was sufficiently similar to the “ground truth” ROS wrapper’s point cloud, that flight code output was used as a “ground truth” of the correct order to check coordinate error [52].




Figure 6.9. “Ground Truth” Point Clouds. *The ROS wrapper’s point cloud was ordered differently from that of the flight implementation, making a direct comparison to determine accuracy infeasible. The fastest-implemented solution, rather than using the point cloud generated by the ROS wrapper (top) to represent itself, used an output of the flight implementation point cloud (bottom). Because the two clouds were geometrically nearly identical, as shown by the side-by-side comparison above and as could be seen by overlaying the two clouds, the bottom point cloud was a valid representation of the ROS wrapper’s point cloud. This bottom point cloud also had the same order as any point cloud output from the flight implementation, and therefore could be used as “ground truth” in test cases. The same comparative approach was used to generate “ground truth” for each of the three test case datasets (Figure 6.8).*

The functions that did not require testing on all three datasets were still checked for accuracy. For Q matrix accuracy, a calibration on the “blocks world” dataset was used as input, and the

output was compared to the hard-coded expected output of the function. The expected output had been derived mathematically [59] and verified via a calibration tool [60] different from the tool used by MoonRanger [60], which does not calculate the Q matrix. To test disparity filtering, the output of the method was checked relative to the input to be similar within ten percent, determining that over-filtering did not occur.

Result 6.3: The flight code implementation contains extensive fault detection, fault mitigation, and unit testing code. The unit tests ran on three datasets which stressed calculation on defined geometry, on images under an approximation lunar polar sunlight, and in darkness.

6.5 ACCURACY EVALUATION AND RUNNING LIVE

The implementation had high accuracy relative to a dense accurate LIDAR scan [53][54] of ground truth and successfully ran live on rover hardware. The root-mean-square error of the calculated point cloud relative to ground truth blocks on “Blocks World” (Figure 6.8) was less than 1 centimeter (Figure 6.10). This is one-fourth the size of the smallest (four centimeter) height required by the system. The final verification was successfully running the implementation on rover computers using rover cameras (Figure 6.11).




Figure 6.10. Implementation Point Cloud vs. Ground Truth. *Left: A front view of “Blocks World”. Right: Angled view of generated and ground truth points. A dense accurate LIDAR [53] scan of “Blocks World” generated the points in pink [54]. The flight code implementation generated the points in white. The LIDAR scan contains blocks, not the black drape in the background, plywood base, or the wall, all of which are seen in the white point cloud. The RMS (root-mean-square) error between these two clouds, assuming the blocks comprise 20% of the scene, was sub-centimeter.*




Figure 6.11. Live On-Rover Stereo Calculation. Top: Live on-rover calculation of depth model. Bottom: Cropped image of the scene from right camera. The large rock in the center was 1 meter away from the rover.

Result 6.4: The flight code stereo, a fault tolerant implementation that carefully manages memory and successfully creates depth models when run on rover hardware, is ready to fly to the Moon and enable MoonRanger’s autonomous navigation in 2023.

7. Space-Readiness and Verification

A rover’s perception system must be proven to function in the conditions it will encounter in space before launching. The hardware described here will be violently shaken during launch. Temperature changes could adversely alter the stereo vergence and baseline. On the Moon, the hardware will be exposed to vacuum and specific thermal conditions. Further, the system will image terrain regolith of various shades and illumination of brilliant and low intensity. The system must continue to function after launch vibrations and throughout lunar vacuum, thermal, and regolith scenarios. Proof that the perception system is space-ready is presented here (illumination-readiness in Section 4).

7.1 QUALITY AFTER LAUNCH VIBRATION

Analysis done as part of MoonRanger Avionics Vibration Test conducted by Daniel Scher, Tim Angert, and Dream Vo. Camera mount designed by Calvin Boyle.

The camera system and the projectors were mounted to a vibration table and vibrated at rocket launch loads. The cameras maintained focus, depth models created before and after were identical, and the laser dots maintained brightness and position.




Figure 7.1. Camera Focus Before and After Launch Vibrations. High-resolution images of a focus target 1 meter in front of the rover before and after vibration testing show no observable differences in the quality of camera focus. The images above are cropped to the focus target for ease of understanding.




Figure 7.2. Depth Model Before and After Launch Vibrations. Depth model generated from images of “Blocks World” (Figure 6.8) before (white, top left) and after (pink, top right) cameras were vibrated at launch loads. The bottom image depicts both depth models in the same frame. The RMS error in matching the denoised point clouds is sub-centimeter (0.0088 meters). With further denoising, the RMS error becomes 0.0044 meters.




Figure 7.3. Point Cloud Distances Before and After Launch Vibrations. The clouds are nearly identical (blue), except for occasional noise exhibited when calculating depth for black drapes behind “Blocks World”, which were expected to have no solution at all (green/yellow/red).




Figure 7.4: Laser Dots Before and After Launch Vibrations. Infrared laser dots projected onto Blocks World before (left) and after (right) vibration testing. Neither the region nor intensity of illumination appreciably change. The maximum intensity of a pixel on the image was maintained at 255 on each color channel of each camera, further suggesting no degradation of laser dots.

Result 7.1: The perception system will survive the violent shaking of a rocket launch.

7.2 FUNCTIONALITY IN VACUUM

Tests led by the author and run in conjunction with teammates Serene Feng and Samyank Jain.

The primary concern in vacuum is that electronic components and lenses will fail. The camera and lens subsystem was tested by imaging a focus target at atmospheric air pressure and approximately 10^{-5} kPa of vacuum and showed that the camera maintains focus. Imaging the dots in atmospheric pressure and 10^{-5} kPa of vacuum showed that the dots maintain functionality.

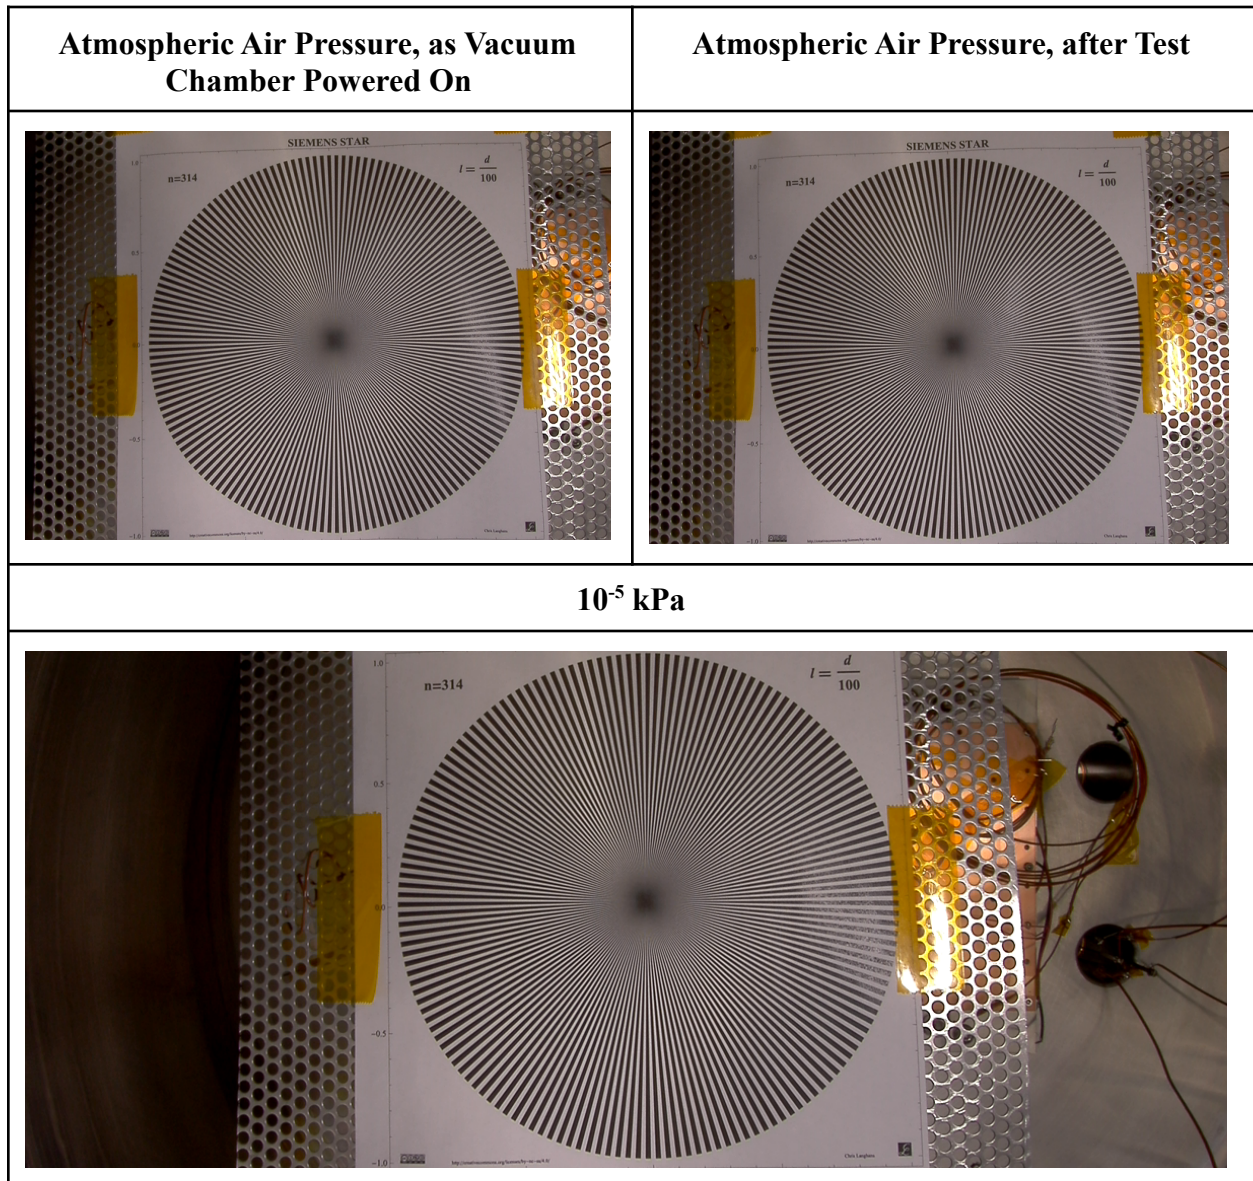



Figure 7.5: Camera Sensor and Lenses in Vacuum. A camera and lens subsystem was placed at the bottom of the vacuum chamber facing up, and a focus target [62] was hung directly above it.

There is no noticeable change in focus between images taken with the camera and lens subsystem in atmospheric air pressure (top) and vacuum (bottom).

ATMOSPHERIC AIR PRESSURE



VACUUM




Figure 7.6: Laser Dots in Vacuum. The pattern and intensity of laser dots with a single projector in air (top, atmospheric pressure) and in vacuum (bottom) are nearly identical. The scene is a dented aluminum plate in a cylindrical vacuum chamber. The images above are cropped to the main section of dots. A purple haze appeared in some “vacuum” images; however, dot intensity appeared unaffected.

Result 7.2: The cameras, lenses, and laser dot projectors function in vacuum.

7.3 FUNCTIONALITY IN SPACE THERMAL CONDITIONS

Tests led by the author and Tim Angert, run in conjunction with Serene Feng and Samyank Jain.

Thermal simulations determined that the perception system components must function from -25°C to 50°C [55] when on the Moon. The components were placed in a thermal chamber, heated to slightly above 50°C and cooled to slightly below -25°C . Images of a focus target taken at the extremes of that range suggested that the camera sensor and lens subsystem will function across the required thermal range. Images of the laser dots suggested that the dot projectors will function across this range as well.

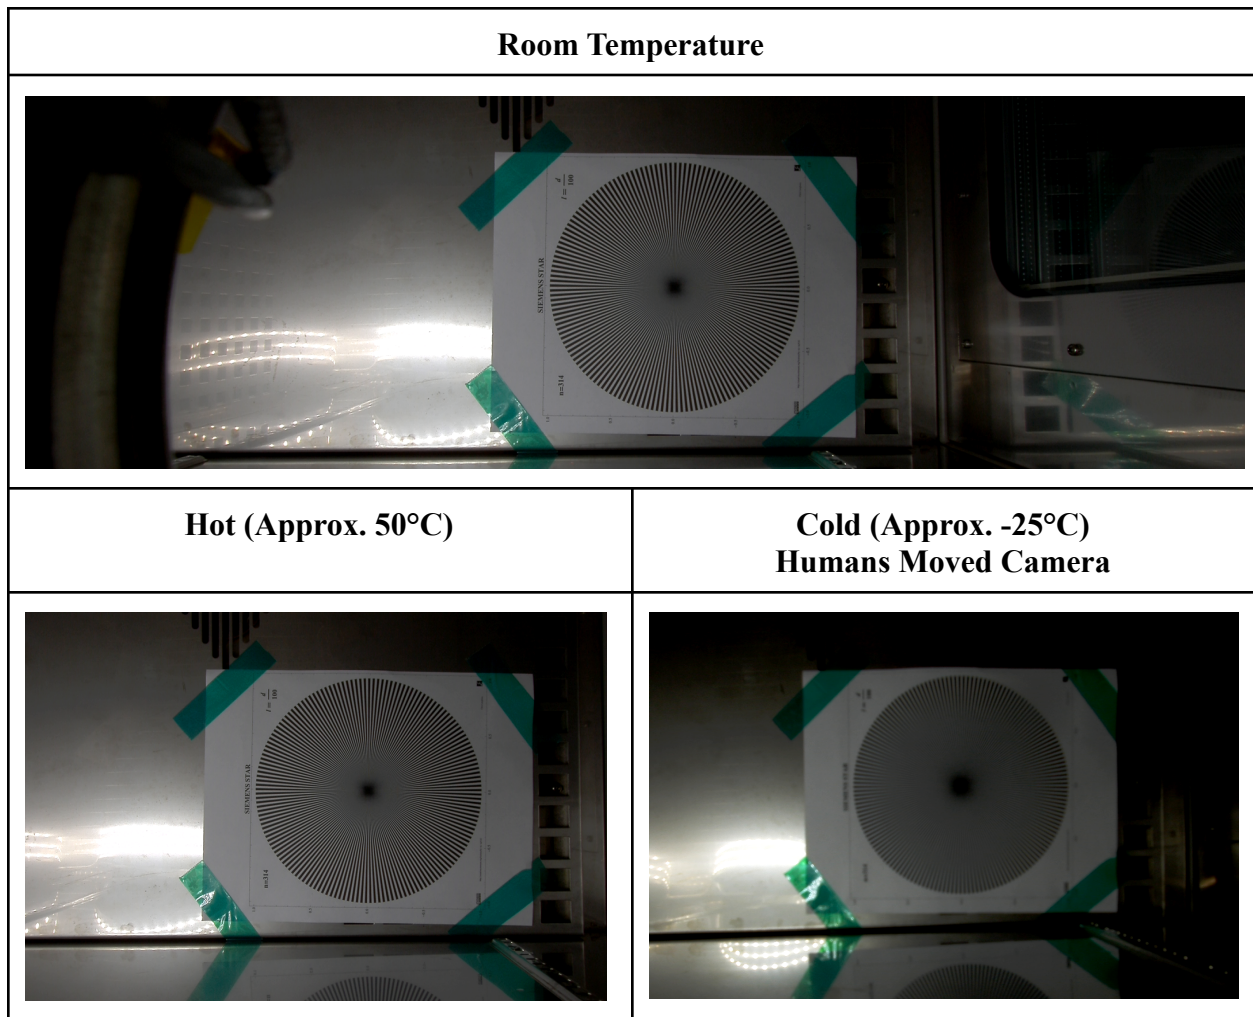


Figure 7.7. Camera Sensor and Lens Subsystem Across Thermal Range. The scene is a focus target in front of the camera sensor and lens subsystem. Focus is maintained from room temperature (top) to 50°C (bottom left). Focus at -25°C (bottom right) is slightly blurred; however, the hypothesis is this was caused by condensation in the thermal chamber, which is not present in space. Images are cropped to the focus target, and the bottom right image was taken after moving the camera.

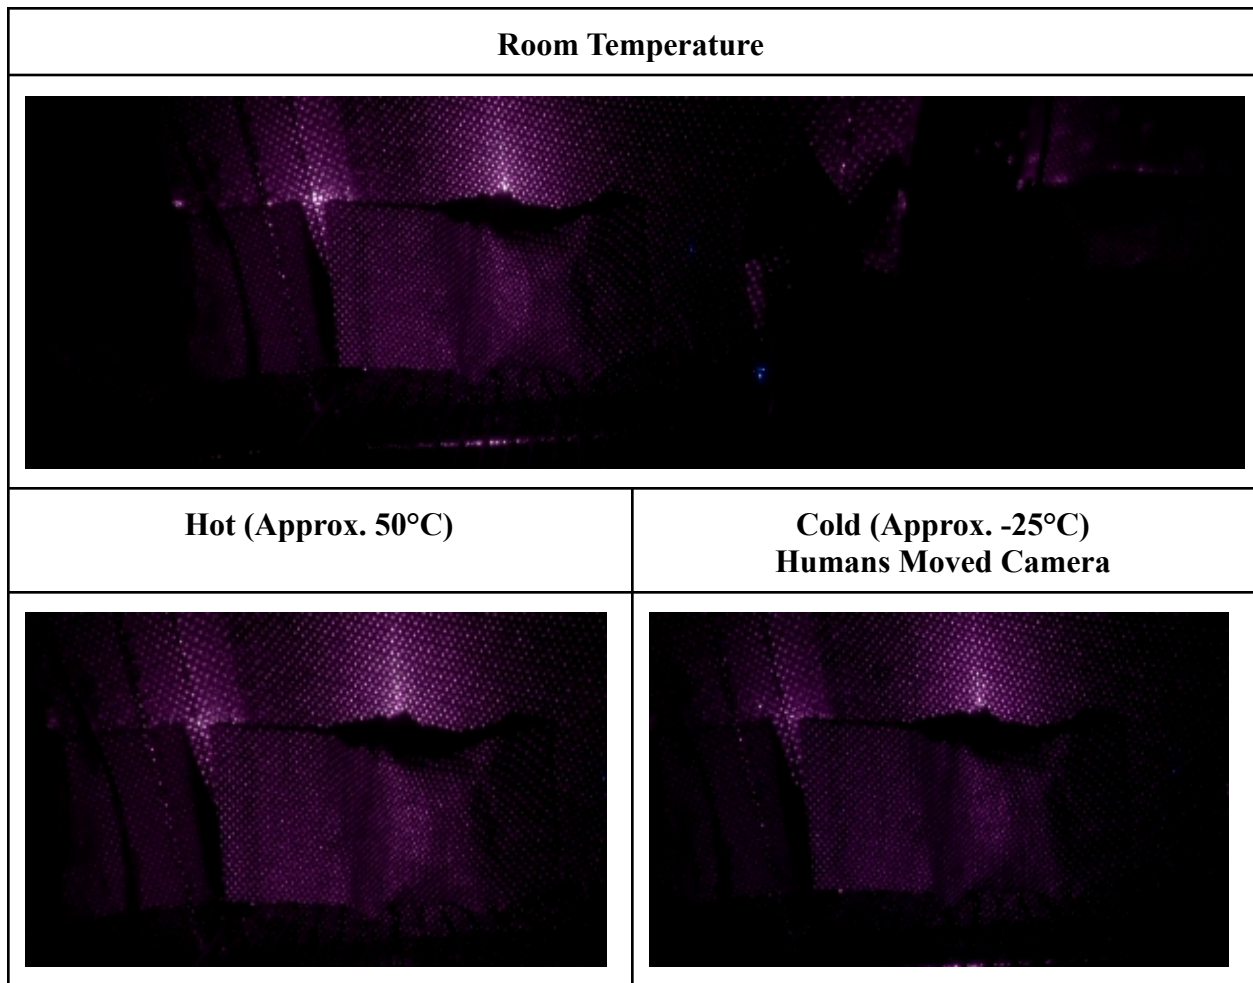


Figure 7.8. Laser Dots Across Thermal Range. The scene is a dented aluminum plate illuminated by a single dot projector. Dot intensity is maintained from room temperature (top) to 50°C (bottom left). Dot intensity at -25°C (bottom right) is slightly less but not so different as to disqualify the dot projectors from use.

Result 7.3: The cameras, lenses, and laser dot projectors function across the expected thermal range. Functionality is better at warmer temperatures than cooler.

7.4 FUNCTIONALITY ON DARK REGOLITH

Lunar regolith, though thought to be nominally gray at the poles [32], can be near-black or quite bright. Because the darkest lunar regolith has a low albedo, or ratio of reflected light to incident light [7], a concern was that the darkest regolith would reflect insufficient illumination back to the cameras. The absolute worst case for the perception system is imaging in total darkness, with only laser dot projectors as illumination. Because sunlight is orders of magnitude brighter than rover illumination, if the perception system can perceive dark regolith with only rover

illumination, it will perceive the dark regolith in sunlight. After creating CMU-2 (Figure 7.9), a regolith simulant of albedo 0.06, the darkest on the Moon [56], terrain dusted with CMU-2 was imaged in darkness. On CMU-2, the perception system had a viewing horizon of 1 meter (Figure 7.10), a non-ideal but sufficient viewing horizon for planetary micro-rover safeguarding.




Figure 7.9. Creating CMU-2 Dark Lunar Regolith. CMU-2 is a 3:2 by volume mixture of coal dust to cement powder. The coal dust was sieved to isolate powder and remove larger grains. Left: To measure the albedos of CMU-2 and cement powder, the regolith simulants were dusted onto a matte acrylic board and illuminated by an all-purpose light for jumpstarting cars. Lux readings were taken above various points on the covered board. The incident lux was measured by holding the lux meter close to the board, facing the light. The reflected/outgoing lux was measured by keeping the meter in the same position, but flipping it over and angling it such that it could measure reflected light from the dusted board. Right: The measured lux readings suggested an albedo of 0.10 for the cement powder and 0.06 for CMU-2. The initial choice of a 3 coal dust : 2 cement powder ratio for CMU-2 was informed by the ratios of lime, cement powder, and coal dust for CMU-1, a lunar equatorial regolith simulant [57].

Result 7.4: A 3:2 by volume ratio of coal dust to cement powder, named CMU-2 by this work, approximated an albedo of 0.06, that of the darkest regolith on the Moon.

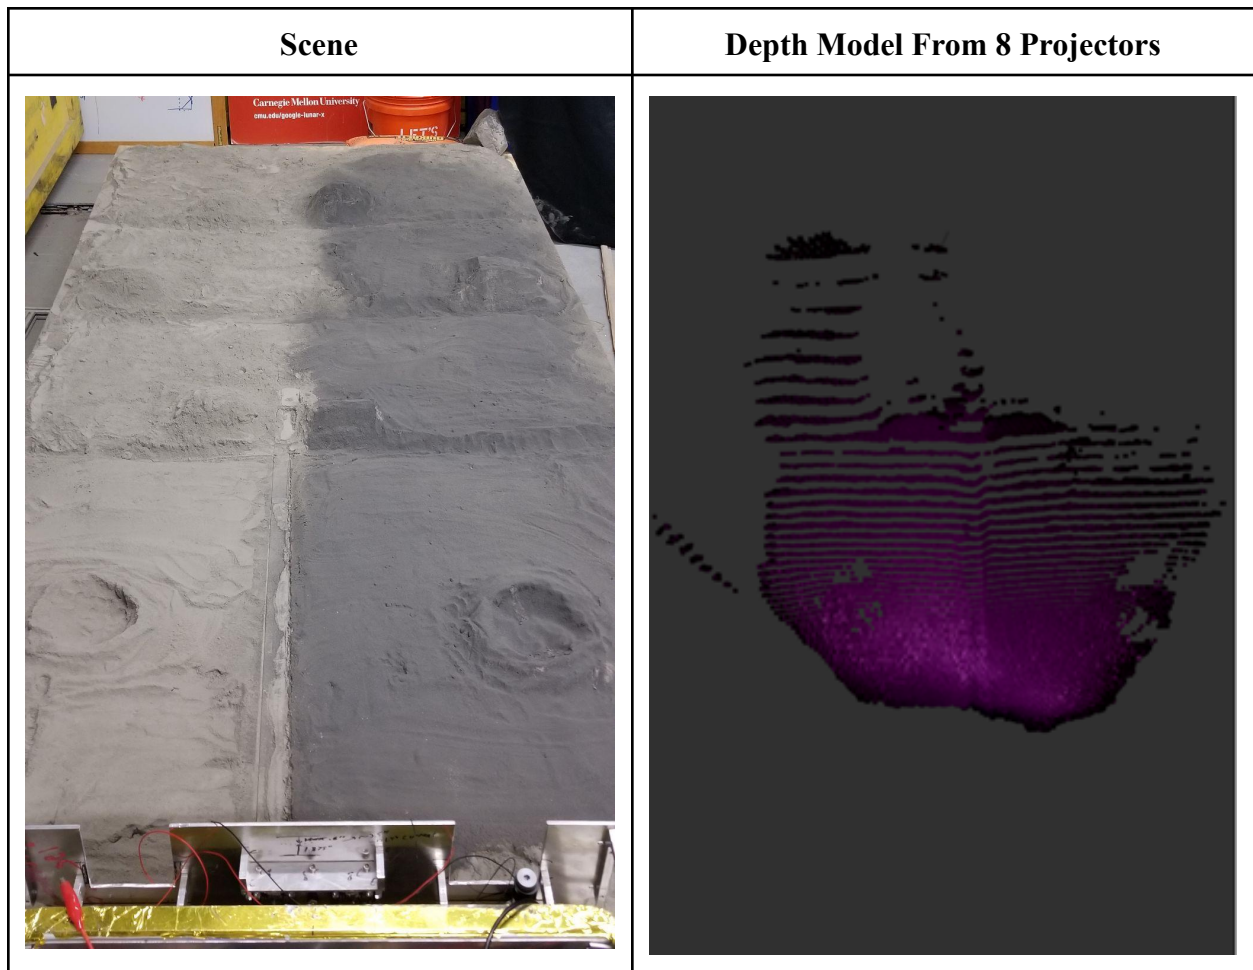


Figure 7.10. Depth Models on Dark Regolith. A terrain (left) with two optical simulants including two craters and three sets of ridges was used to measure the increase in viewing horizon when increasing the number of projectors. The set of craters and each set of ridges were sculpted at approximately half-meter increments measured away from the rover. The brighter optical simulant, cement powder, represented slightly darker than expected optical properties of lunar polar regolith while the darker simulant, CMU-2, represents the absolute darkest possible on the moon. On both types of regolith, the viewing horizon using 8 projectors (right) is at least 1 meter, a sufficient viewing horizon for a micro-rover [72].

Result 7.5: The perception system discussed in this work is space-ready. It survived rocket launch loads, functioned in the vacuum and thermal conditions of space, and successfully computed depth models on the darkest lunar regolith color possible.

8. High Dynamic Range for Continuous Polar Roving

Analysis thus far has assumed a rover that is traveling with the sun at a 90° heading relative to the rover. However, as a lunar polar micro-rover navigates terrain, its bearing with respect to the sun will vary significantly. As a rover turns toward the sun, the sun will start to blind the cameras. As a rover turns away from the sun, the terrain it images will appear brighter because the opposition effect causes terrain directly opposite the sun to appear nearly as brilliant as the sun itself [13]. In both scenarios, regions of images will become oversaturated and exhibit “whiteout.” The heading at which image whiteout becomes detrimental to rover perception, known as a heading limit, is a parameter for path planning algorithms. Planned paths that would cause a rover to turn past a heading limit are avoided. This section investigates what this work’s perception system’s heading limits are and how using High Dynamic Range (HDR) principles to merge an additional image of 1 millisecond exposure to the nominal image of 16 millisecond exposure could mitigate whiteout.

8.1 SIMULATING SUNLIGHT AT DIFFERENT HEADINGS

An extension to the experiment casting an approximation of lunar polar sunlight on lunar analog terrain (Section 4) was moving the light to different headings relative to the cameras. The terrain was imaged at heading increments of 30 degrees, and the images captured informed heading limit analysis.




Figure 8.1. “Sunlight” at Various Headings. Left: An Arrimax 18/12 cinematic daylight casts solar intensity illumination on lunar analog terrain. The lines highlighted in yellow were strings taped on the ground at thirty degree increments. Right: A diagram representing the headings tested. The black lines correspond to the yellow string in the left image. The Arrimax 18/12 light was shifted to simulate the rover turning thirty and sixty degrees into the sun (30° and 60° in diagram) and thirty and sixty degrees away from the sun (-30° and -60° in diagram).

8.2 HEADING LIMITS WITHOUT SHORT-EXPOSURE HDR

After analyzing depth models from images taken at the chosen (Section 4) 16 millisecond exposure at the various headings, heading limits for the perception system were determined to be slightly less than 30 degrees into and away from the sun (Figure 8.2). The worst case “whiteout” was exhibited at a 60 degree turn into the sun.

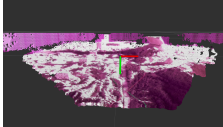
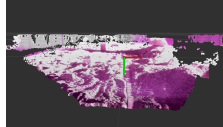
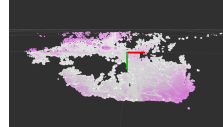
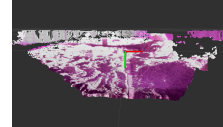
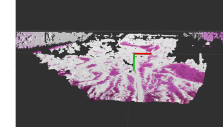
Rover Perpendicular to Sun	Rover Facing 30° Into the Sun	Rover Facing 60° Into the Sun	Rover 30° Away From the Sun	Rover 60° Away from the Sun
				

Figure 8.2. Depth Models at Various Headings and 16 ms Exposure. The depth models have data in regions of color and have no solution in the gray regions. All headings except for perpendicular to the direction of sunlight exhibit whiteout, regions of oversaturation in images that translate to “no solution” gray regions in the depth models. However, whiteout is small in the 30° degree turns (region of gray in top right of depth models), suggesting that a heading limit would be slightly less than 30 degrees. Note that a 60° turn into the sun exhibits the worst whiteout.

Result 8.1: Heading limits for the perception system are slightly less than 30° turns into or away from the sun. Of the headings tested, the worst whiteout occurs at a 60° turn into the sun.

8.3 SHORT-EXPOSURE HDR

High Dynamic Range (HDR) algorithms commonplace in photography merge images captured at many exposures, creating well-illuminated images even at night and visible images even in glaring sunlight [1]. However, traditional HDR requires a camera to be stationary and merges images at many exposures, including exposures hundreds of milliseconds long. Stopping so that long-exposure images do not exhibit motion blur drastically slows a rover mission. Though the delay might be acceptable on multi-year missions, for weeklong polar micro-roving missions, relying on traditional, stationary HDR images critically limits exploration range.

Merging images with a combined exposure time less than the time for a single image to exhibit motion blur would enable a rover to navigate terrain without stopping. Because of the brilliant sunlight on the Moon, merging an image exhibiting “whiteout” with an underexposed image that

perceives the regions of whiteout uses two images of relatively short exposure (Figure 8.3), hence the term Short-Exposure HDR.


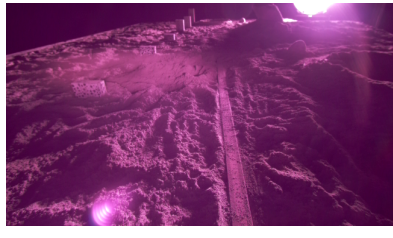
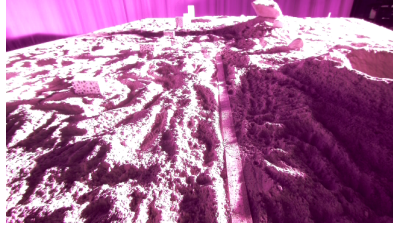
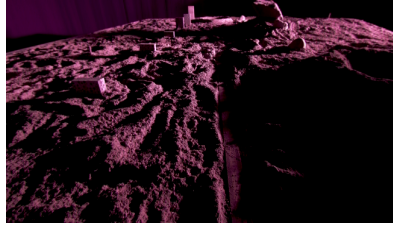
	16-Millisecond Image	1-Millisecond Image
Rover Facing 60° Into the Sun		
Rover Perpendicular to Sun		

Figure 8.3. Image Exhibiting Whiteout and Underexposed Image. Images from the left camera taken at an angle of 60° turned into the sun (top, extreme condition), and with the rover perpendicular to direction of illumination (bottom, nominal condition). Left: Images taken at the nominal 16-millisecond exposure. Right: Images taken at a 1-millisecond exposure.

Because 16 milliseconds is the exposure time at which the least motion blur is expected as a micro-rover navigates terrain [35], an ideal merge might be a 15-millisecond image and a 1-millisecond image. To actually implement Short-Exposure HDR on a rover, a modified camera driver that alternates between capturing 1-millisecond and 15-millisecond images would be necessary. Further, the exact exposure times for the regular and underexposed images will depend on how long a camera takes to read its sensor.

For the perception system in this work, the sensor [58] has a maximum frame rate of 240 Hz, which suggests that even a 1-millisecond exposure may take 4.1 milliseconds (1/240 seconds) to process. As such, without extensive workarounds, the cameras can only provide images for a 12-millisecond/1-millisecond merge. The only potential concern with lowering exposure time to 12 milliseconds might be loss of viewing horizon in darkness relative to that seen in Section 4. However, given that the 25% decrease in exposure duration is less than the 33% increase in projected intensity from 6 to 8 projectors (Section 5.4), it is likely that the viewing horizon seen in Section 4 will be maintained. Alternatively, a 16-millisecond/1-millisecond merge could be accepted if a rover is willing to accept minor motion blur or slow down.

Because a 16-millisecond is brighter and therefore exhibits more regions of “whiteout” than a 12-millisecond image, the rest of this section demonstrates Short-Exposure HDR on a 16-millisecond/1-millisecond merge to prove that severe “whiteout” can actually be eliminated. A simple image merging algorithm picked pixels from the 1-millisecond image only when all three color channels of the brighter image were oversaturated i.e. had a pixel value of 255. A more advanced merging algorithm might be preferable on a flight rover. Nevertheless, this simple merging algorithm generated Short-Exposure HDR Images that significantly increased depth model completeness in images exhibiting whiteout (Figure 8.5).

(Continued on next page)

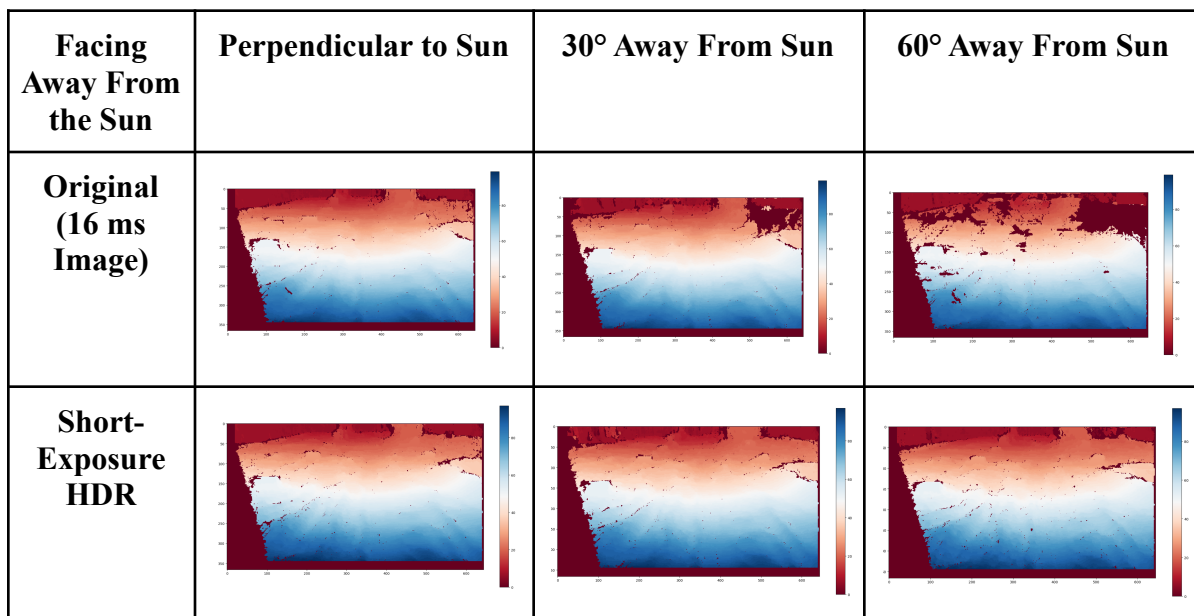
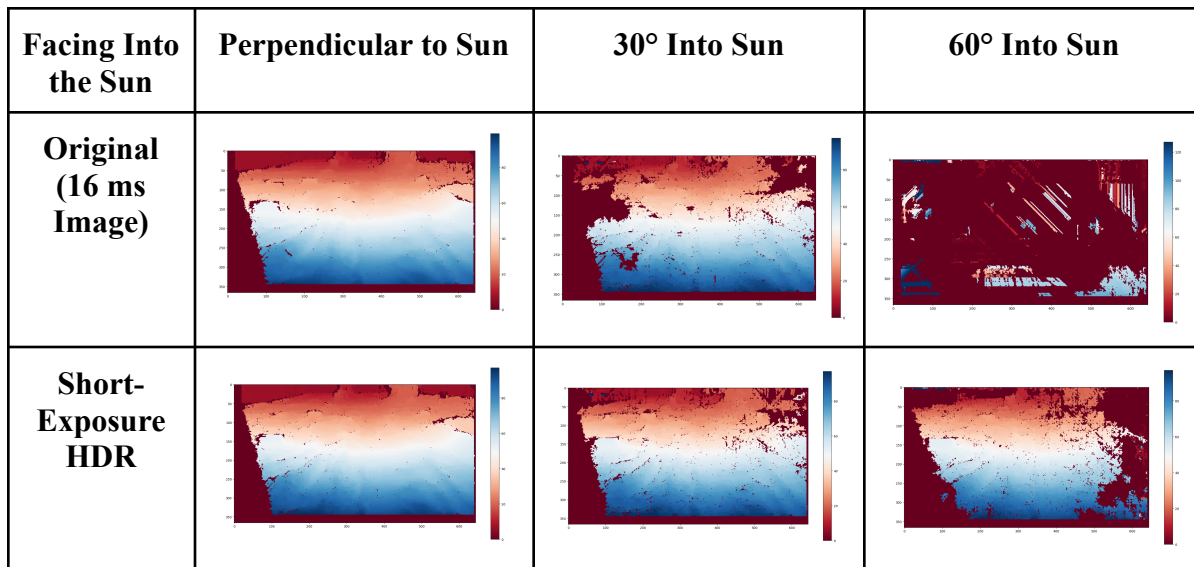


Figure 8.5. Short-Exposure HDR Depth Mitigates Whiteout. Disparity maps colorized above show bright colors where a depth model was calculated and the darkest red where no solution was found. Points closer to the cameras were bluer and those farther away were more red. The original disparity map with the rover perpendicular to the direction of illumination (top left disparity map) can be considered the “true” disparity map. Without Short-Exposure HDR, the disparity maps at non-perpendicular headings have large regions of no solution. In comparison, the disparity maps calculated from Short-Exposure HDR images are more complete and better match the “true” disparity map.

Result 8.2: Short-Exposure HDR increases depth model completeness in images otherwise exhibiting whiteout.

8.4 SHORT-EXPOSURE HDR MAINTAINS DEPTH MODEL CORRECTNESS IN NON-WHITEOUT IMAGES

Given that Short-Exposure HDR mitigates whiteout, a remaining question was whether the completeness and correctness of a depth model in non-whiteout conditions is maintained. There are two non-whiteout cases: when a rover is perpendicular to the direction of sunlight and when a rover is imaging in darkness with laser dots as illumination. In both cases, depth model completeness and correctness are maintained (Figure 8.5).

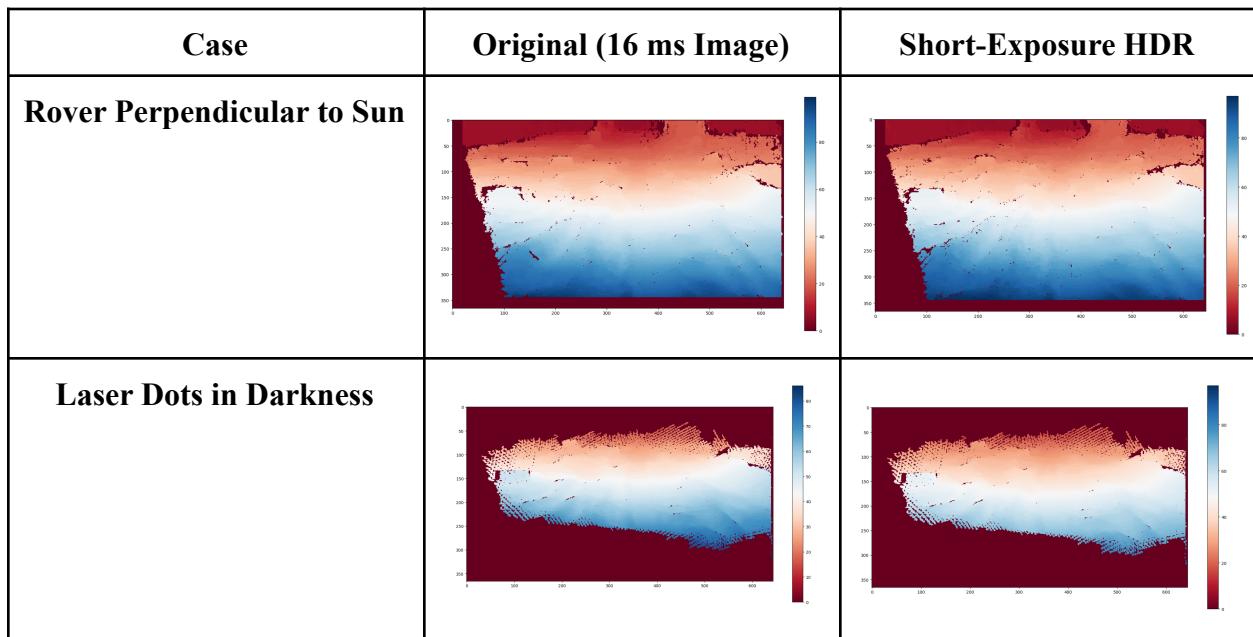


Figure 8.6. Short-Exposure HDR Maintains Depth Model Completeness in Non-“Whiteout”

Conditions. Disparity maps colorized above show bright colors where a depth model was calculated and the darkest red where no solution was found. Points closer to the cameras were bluer and those farther away were more red. Both when a rover is perpendicular to the direction of sunlight and in darkness illuminated by laser dots, the completeness of and distances in the original depth model (center column) and the depth model from Short-Exposure HDR images (right column) were identical.

Result 8.3: Short-Exposure HDR maintains depth model completeness and correctness in non-whiteout conditions.

Result 8.4: Short-Exposure HDR, which merges a ≤ 1 -millisecond exposure image with a 16-millisecond image, can eliminate whiteout and enable a rover to perceive terrain at any heading relative to the sun. With a modified camera driver, Short-Exposure HDR could run live on the perception system hardware presented in this work and should be incorporated in future micro-rover missions.

9. Conclusion

9.1 SUMMARY

This research addressed stereo algorithms, lensing, exposures, material appearance, and terrain illumination unique to lunar poles. Additional considerations included processing on space-relevant computing, achieving the short cycle times essential for achieving continuous rover motion, and space-readiness verification of camera hardware. The principles, issues, and investigations pertained to image-based stereo in extreme lighting. The solutions here are distinct from a purely academic inquiry in that they are manifested into a rigorous flight code and hardware that is enabling navigation and autonomy on the Moon and for that reason meet a gold standard of robustness.

The research discovered a lens aperture, exposure time, source of illumination in darkness, and stereo algorithm that both enable perceiving terrain in scenes of brilliant sunlight and absolutely black shadows characteristic of lunar poles and are simple enough for use on a micro-rover. Considerations, design, implementation, and testing of illumination source and geometry were discussed. Perception research agenda additionally identified a preferred stereo algorithm and implemented stereo to flight-code rigor. The resulting perception system was verified as space-ready and will guide Carnegie Mellon's autonomous MoonRanger micro-rover that will map ice near the lunar south pole in 2023.

Images taken while the rover faces the sun or faces away from the sun exhibit whiteout and cannot be analyzed. Hence, the rover cannot plan or drive headings toward or away from the sun. The research determined the range of acceptable headings. The research also investigated how Short-Exposure High Dynamic Range (HDR) could merge an image of 1 millisecond exposure and to the nominal image of 16 millisecond exposure to mitigate image whiteout, thereby removing heading limits in future lunar polar missions.

9.2 REGARDING THE RESEARCH'S PERCEPTION SYSTEM

The research's perception system suggests the following:

1. Stereo algorithms sufficiently model sunlit and dot-illuminated lunar polar terrain.
2. No stereo is optimal for all lighting variants encountered on the Moon, but libSGm is identified as efficacious.
3. Micro-rover cameras and illumination sources can function in space-relevant vacuum and thermal scenarios and survive launch-relevant vibration loads.

4. Multiple infrared laser dot projectors can additively increase incident illumination on terrain and rover viewing horizon in darkness. Infrared laser dot projectors are superior sources of illumination compared to wattage-equivalent LEDs.
5. Infrared laser dot projectors should be placed above cameras, angled slightly downwards and verged slightly outward.
6. A shorter depth-of-field from a larger aperture lens was acceptable given that the portion lost was extremely small and would be imaged clearly at an earlier time step.
7. Lunar polar micro-rover perception of terrain might be limited to rover headings within 30° relative to the perpendicular to the direction of illumination.

9.3 REGARDING ROBOTIC PERCEPTION FOR EXPLORING LUNAR POLES

The research concludes the following relating to robotic perception for exploring lunar poles:

1. Micro-rover constraints such as low height, myopic field of view and low-wattage terrain illumination are challenging but do not preclude perception suitable for the lunar poles.
2. A flight-quality stereo implementation can successfully run on rover computing, not just the significant computing and camera resources available on Earth.
3. A micro-rover perception and illumination system can generate accurate depth models of terrain in brilliant sunlight and in darkness without changing camera hardware, camera parameters, or stereo algorithm used, or over-taxing micro-rover power constraints.
4. A lunar polar micro-rover can perceive and model terrain without stopping. Therefore, an autonomous lunar polar rover could sustain continuous motion, exploring the poles for the entire duration of a mission.
5. Merging a nominal image that exhibits whiteout with an underexposed image that perceives the regions of whiteout can exponentially increase the region of terrain for which depth is calculable. Such a Short-Exposure HDR algorithm should be included in future lunar polar rover perception systems.

This research is a foundation for all small lunar polar rover perception yet to come. Future rovers should leverage Short-Exposure HDR so that there is less burnout in brightness and hence they can orient at more headings relative to the sun. They might also advance camera technology to enable a modification of Short-Exposure HDR that captures images of shorter exposures within the nominal image of 16 milliseconds of exposure. Merging these multiple images would completely eliminate whiteout at any orientation relative to the sun. Adding images to create the longer, 16 millisecond image would keep the total exposure at 16 milliseconds, completely avoiding motion blur artifacts in merged images without degrading the intensity of images in dot-illuminated darkness. Future rovers for longer-duration missions might also trade continuous motion for power savings by using less illumination and the traditional stationary HDR and step-stop-drive rover navigation paradigm.

10. Acknowledgements

I would first like to acknowledge and thank my mentor, Dr. Whittaker. You have been a pillar in my education. Without you as a mentor, I do not know where I would be right now. You are always there to answer any question or jump in to help a test. You have advised, coached, and encouraged me throughout this thesis and my undergraduate research. I do not know how to express the impact you have had on my education at Carnegie Mellon.

To Dr. Wettergreen, your positivity and attention to detail has encouraged me to keep improving on solutions. Your advice comes right before I would be overwhelmed, and you are always ready to take a deeper look at a question or review a proposed answer.

To Chuck Whittaker, you are always there for whatever test needs to be run or question I would like to ask. You give and give your time and positivity.

To Dr. Jones, you find solutions when there seems to be none. Thank you for always being there. Thank you for your advice.

To Srini Vijayarangan, thank you for your expertise, your advice, and for always being available to answer a technical question.

To Tenzin Crouch and William Lee-Moore, you have shown me what true leadership is. Tenzin, thank you for being such an open listener to random ideas. William, thank you for staying positive as much of this research took longer than we had hoped.

To Dan Scher, thank you for creating the set for the lunar lighting test and for helping run it. You did not have to pour so much of your time into it, but you did.

To Tim Angert, thank you for encouraging me to learn things outside of my area of expertise, and for teaching me as well.

To Calvin Boyle and Paulo Fisch, thank you for your tireless attitude and the simulations and design that created the camera hardware and informed hardware testing.

To Serene Feng and Samyank Jain, thank you for working with me on vacuum and thermal tests of camera hardware.

To Sebastian Bernal, Saksham Khurana, and Ben Kolligs, thank you for believing in the dot projector solution and for designing the PCBs to run the laser dot projectors.

To Maggie Hansen, Abby Breitfield, and Dream Vo thank you for your positive attitude when teaching me how to use the updated Morphin and for staying for the many hours of the lunar lighting test.

To NASA, thank you for the opportunity to work on such a groundbreaking rover by supporting the LSITP contract 80MSFC20C0008 MoonRanger.

To my parents and brother, thank you for your unconditional support.

11. References

- [1] Gin, Dan. “What is HDR photography and how can I shoot it with my camera?,” <https://www.digitaltrends.com/photography/what-is-hdr-photography/>
- [2] Kumar, V., et. al. “Formulation Of Micro-Rover Autonomy Software For Lunar Exploration,” iSAIRAS, October 2020.
- [3] Nikon. “Getting Started: How to Change a DSLR Lens,” <https://www.nikonusa.com/en/learn-and-explore/a/tips-and-techniques/getting-started-how-to-change-a-dslr-lens.html>
- [4] Whittaker, W. L. “Space Robotics Development,” Carnegie Mellon University 16-865, Spring 2021.
- [5] Smith, G. H. “The Mars rover camera lenses,” <https://support.zemax.com/hc/en-us/articles/1500005577202-The-Mars-rover-camera-lenses>.
- [6] Teza, J. “Solar simulation and illumination,” Presentation for Space Robotics Development.
- [7] “Albedo effect,” <https://www.npolar.no/en/fact/albedo/>.
- [8] Gkioulekas, I. “Photometric stereo and shape from shading,” 16-385 Computer Vision, March 2020.
- [10] “The Difference Between Lux and Lumen: What is Brightness?,” <https://blog.betterautomotivelighting.com/the-difference-between-lux-and-lumen-what-is-brightness>.
- [11] Wong, U., et. al. “POLAR Stereo Dataset,” https://ti.arc.nasa.gov/dataset/IRG_PolarDB/.
- [12] Wong, U., et. al. “Polar Optical Lunar Analog Reconstruction (POLAR) Stereo Dataset. NASA Ames Research Center,” May 2017.
- [13] Hapke, B. W., et. al. “The opposition effect of the moon: the contribution of coherent backscatter,” National Library of Medicine, DOI: 10.1126/science.260.5107.509, April 1993.

- [14] Allan, M. "NASA Resource Prospector Simulator "rpgsim" - ROScon 2018," <https://www.youtube.com/watch?v=w-ylrw0zdqM&t=66s>.
- [15] Allan, M. and Chen, I. "Gazebo renders the moon," ROScon, September 2018.
- [16] NASA Mars Exploration Rovers. "Pancam," <https://mars.nasa.gov/mer/mission/instruments/pancam/>.
- [17] Beyer, R. A., et. al. "VIPER Visible Imaging System," 53rd Lunar and Planetary Science Conference, March 2022.
- [18] NASA Science. "2020 Mission Perseverance Rover," <https://mars.nasa.gov/mars2020/spacecraft/rover/cameras/#NavCams-Engineering-Cameras>.
- [19] NASA. "VIPER's Mission Operations," <https://www.nasa.gov/viper/lunar-operations>.
- [20] Indian Space Research Organization. "Chandrayaan 2 - About Pragya Rover," <https://www.isro.gov.in/gslv-mk-iii-m1-chandrayaan-2-mission/chandrayaan-2-about-pragya-rover>.
- [21] Chakorborty, N. "Hours Before Landing On Moon, ISRO Reveals How Chandrayaan-2 Rover Pragyan Works," <https://english.newsnationtv.com/science/news/hours-before-landing-on-moon-isro-introduces-to-pragyan-chandrayaan-2s-rover-236803.html>.
- [22] Whittaker, W. L. "Space Robotics," Carnegie Mellon University 16-861, Fall 2021.
- [23] "Yutu," <https://solarsystem.nasa.gov/missions/yutu/in-depth/>.
- [24] Wang, J., et. al. "Computer Vision in the Teleoperation of the YUTU-2 Rover," ISPRS Annals of the Photogrammetry, Remote Sensing and Spatial Information Sciences, Volume V-3-2020, 2020, pp.595-602, DOI: 10.5194/isprs-annals-V-3-2020-595-2020, August 2020.
- [25] "Tank on the Moon," <https://www.youtube.com/watch?v=JUkAVZMX8QU>.
- [26] Jet Propulsion Laboratory. "A Description of the Rover Sojourner," <https://mars.nasa.gov/MPF/rover/descrip.html>.
- [27] Amos, J. "China's Zhurong Mars rover takes a selfie," <https://www.bbc.com/news/science-environment-57441757>.

- [28] NASA Science. “The Cameras on the Mars Curiosity Rover,” <https://mars.nasa.gov/msl/spacecraft/rover/cameras/#navcams>.
- [29] Wall, M. “After a year on Mars, NASA's Perseverance rover is on course for big discoveries,” <https://www.space.com/mars-perseverance-rover-one-year-landing-anniversary>.
- [30] Grunnet-Jepsen, A., et. al. “Projectors for Intel® RealSense™ Depth Cameras D4xx,” https://www.intelrealsense.com/wp-content/uploads/2019/03/WhitePaper_on_Projectors_for_RealSense_D4xx_1.0.pdf.
- [31] ARRI. “ARRIMAX 18/12,” <https://www.arri.com/en/lighting/daylight/m-series/arrimax-18-12>.
- [32] Jones, H. Discussion on Lunar Regolith, Summer 2021.
- [33] “Lens calculator - Calculating the depth of field (DOF),” <https://www.vision-doctor.com/en/optical-calculations/calculation-depth-of-field.html>.
- [34] Whittaker, W. L. “Space Robotics,” Carnegie Mellon University 16-861, Fall 2019.
- [35] Vijayarangan, S. Discussion on rover motion blur, Fall 2021.
- [36] “BELICE-850 Dot-Pattern Infrared Illuminator for 3D Stereoscopic Imaging,” https://ams.com/documents/20143/36005/BELICE-850_DS000618_2-00.pdf/7ee69839-86b4-aab1-6999-e5920567e469
- [37] “QBHP684-IR3BU,” Infrared LED, <https://www.arrow.com/en/products/qbhp684-ir3bu/qt-brightek-corporation>.
- [38] “gpu_stereo_image_proc,” https://github.com/whill-labs/gpu_stereo_image_proc.
- [39] “stereo_image_proc,” http://wiki.ros.org/stereo_image_proc.
- [40] “stereo_dense_reconstruction,” https://github.com/ut-amrl/stereo_dense_reconstruction.
- [41] “Disparity-Map,” <https://github.com/1kc2/Disparity-Map>.
- [42] NVIDIA. “Stereo Disparity Estimator,” https://docs.nvidia.com/vpi/algo_stereo_disparity.html.

- [43] Kumar, V. “Vision-Based Algorithms for Lunar Polar Micro-Rover Autonomy,” Carnegie Mellon University, May 2020.
- [44] Whittaker, W. L. “Space Robotics Development,” Carnegie Mellon University 16-865, Spring 2020.
- [45] Lee-Moore, W. Code for image rectification, Fall 2021.
- [46] Lin, A. Code for calculating points from disparity map, Winter 2021.
- [47] Lee-Moore, W. Code for filtering disparity maps, 2021.
- [48] Gkioulekas, I. “Stereo,” 16-385 Computer Vision, February 2020.
- [49] “msg,” <http://wiki.ros.org/msg>.
- [50] “stereo_msgs,” http://wiki.ros.org/stereo_msgs.
- [51] “CloudCompare,” <https://www.danielgm.net/cc/>.
- [52] Lee-Moore, W. Suggestion to use results of previous execution of flight code as “ground truth” for point cloud testing.
- [53] “FARO,” <https://www.faro.com/en>.
- [54] Whittaker, C. FARO Scan of Blocks World, Fall 2021.
- [55] Fisch, P. R. Expected Thermal Ranges for MoonRanger Components, July 2021.
- [56] NASA DSNE Revision H. “Cross-Program Design Specification For Natural Environments (DSNE),” August 2020.
- [57] Wong, U. “Lumenhancement: Exploiting Appearance for Planetary Modeling,” CMU-RI-TR-12-12, Robotics Institute, Carnegie Mellon University, April, 2012
- [58] Sony. “IMX274,” https://www.sony-semicon.co.jp/products/common/pdf/IMX274LQC_Flyer.pdf.

- [59] “Q matrix for the reprojectImageTo3D function in opencv,” <https://stackoverflow.com/questions/27374970/q-matrix-for-the-reprojectimageto3d-function-in-opencv>.
- [60] Ghosh, S. “stereo-calibration,” <https://github.com/sourishg/stereo-calibration>.
- [61] ROS. “camera_calibration,” http://wiki.ros.org/camera_calibration.
- [62] “Siemens Star n=314,” https://claclaclack.files.wordpress.com/2015/05/siemensstern_n314_cl_v3.pdf.
- [63] Kasson, J. “Printable Siemens Star targets,” <https://blog.kasson.com/lens-screening-testing/printable-siemens-star-targets/>.
- [64] “How to Convert images to NumPy array?,” <https://www.geeksforgeeks.org/how-to-convert-images-to-numpy-array/>.
- [65] ”numpy.amn” <https://numpy.org/doc/stable/reference/generated/numpy.amn.html>.
- [66] “Plotting a 2D heatmap with Matplotlib,” <https://stackoverflow.com/questions/33282368/plotting-a-2d-heatmap-with-matplotlib>.
- [67] “Colorbar,” https://matplotlib.org/stable/gallery/color/colorbar_basics.html.
- [68] “Matplotlib Colorbar Range,” <https://www.delftstack.com/howto/matplotlib/range-of-colorbar-matplotlib-python/>.
- [69] “reading csv file in opencv,” <https://answers.opencv.org/question/55210/reading-csv-file-in-opencv/>.
- [70] “saving and working with depth images in python,” <https://github.com/IntelRealSense/realsense-ros/issues/1312>
- [71] “How to interpret 4 bytes as a 32-bit float using Python,” <https://stackoverflow.com/questions/37093485/how-to-interpret-4-bytes-as-a-32-bit-float-using-python/37093610>.
- [72] Whittaker, W. L. “Space Robotics,” Carnegie Mellon University 16-865, Spring 2022.