POLAR Stereo Dataset

(Polar Optical Lunar Analog Reconstruction)

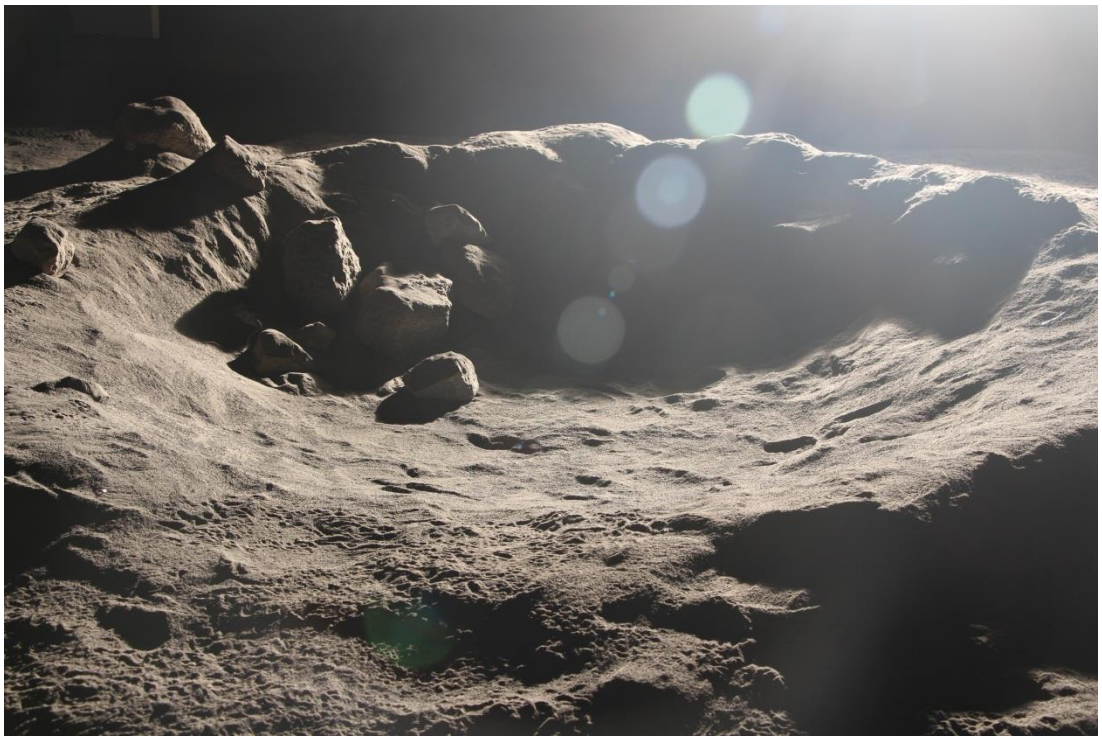

User Documentation V1.4

October 09, 2023

NASA Ames Research Center

Developed by the Intelligent Robotics Group, through the work of:

Dr. Uland Wong (PI, NASA)

Dr. Ara Nefian (SGT, Inc.)

Dr. Larry Edwards (NASA)

Xavier Buoyssounouse (NASA)

P. Michael Furlong (SGT, Inc.)

Dr. Matt Deans (NASA)

Dr. Terry Fong (NASA)

Table of Contents

Description

You have downloaded the NASA Ames Lunar Polar Analog Stereo Dataset for quantitative evaluation of rover stereo imaging in the polar lighting conditions of the Moon. This dataset contains over 2500 HDR stereo pairs and 12GB of sensor data, supplemented with ground truth depth, calibration, and registration information. The rest of this document describes the contents of the dataset and experimental procedure for generating these products.

In generating this dataset, we sought to recreate the surface appearance conditions at the poles of the Moon. These conditions are rarely encountered on Earth, but are common on airless bodies throughout the solar system. During the course of stereo navigation development on a proposed NASA mission to the lunar poles, we were surprised to learn that there was little prior work in the area and almost no publicly available rover-relevant data on this topic. Despite the proliferation of robots and the wealth of terrestrial stereo datasets, few in the robotics community have looked at imaging specific to airless planetary bodies. The objectives of this work were to develop a common framework for stereo vision researchers: (1) understanding and reproduction of polar/airless optical conditions in a laboratory setting, (2) a library of stereo images for algorithm development complete with ground truth, and (3) metrics and standards for comparative evaluation.

Several appearance attributes were considered in producing this dataset. Firstly, the lunar surface is covered in a layer of fine, powdery dust called *regolith*. Regolith is created by the space weathering processes[1] which create jagged, microscopic particles. Optically, this results in characteristic effects such as *heiligenschein* from backscattering and whiteout due to uniform covering. The Lunar reflectance model has been the topic of decades of study [Hapke 2001], and is counter-intuitively neither matte nor specular, despite being strongly view dependent. Secondly, the dominant illumination conditions are oblique[2] sunlight and the absence of scattering atmosphere. Lighting is therefore very harsh, resulting in long cast shadows and high dynamic range (HDR) conditions due to contrasting shadowed and illuminated regions. Lastly, planetary terrain geometry at the rover scale is dominated by fractal-like distributions of rock (positive) and crater (negative) obstacles. Though a principal purpose of stereo imaging is to detect these obstacles, they also introduce occlusions and shadowing, making this task more difficult.

Though this dataset was developed for lunar exploration, it is also applicable to many types of planetary surfaces including: mercury, asteroids, and regolith covered moons like Phobos. While specific environmental parameters do not match these bodies as closely, we believe as of 2016, no other publicly available data of this type has greater fidelity of appearance.

This release is made possible through the work of the Intelligent Robotics Group at the NASA Ames Research Center. The work was funded by the Resource Prospector Mission (through the Advanced Exploration Systems program) and the Game Changing Development (GCD) program. We are extremely

[1] Such as micrometeorite impacts
[2] Low elevation angle

grateful to the Solar System Exploration Research Virtual Institute (SSERVI) for providing laboratory facilities to do regolith testing and operational support.

Experimental Setup

Our experimentation attempts to recreate the polar appearance conditions with highest practical fidelity. The dataset is logically divided into "terrains," which are independent data collection trials organized by a single scene - there are 12 total trials. Each scene is crafted in a 4m x 4xm 0.7m sandbox containing 8 tons of lunar regolith simulant. The terrain plane and negative features are formed from the simulant. Surface undulations are shaped with hand tools according to design. Next, rocks are added to the terrain in order to simulate either planned size-frequency distributions or specific obstacle fields. The terrain and rocks are then dusted with a layer of regolith simulant to produce a "fluffy" top layer of soil, erasing shovel and brush marks and distributing a thin layer on the faces of rocks. Solar simulator lights setup at 4 positions around the terrain create low-angle, high-contrast illumination with the correct angular diameter and a known surface irradiance factor. Near-field lighting and use of blackout material reduces the contribution of atmospheric scattering and diffuse interreflections. All test parameters are summarized in Table 1.

Scene Construction

Regolith**:** Regolith is the lunar dust which covers the vast majority of surfaces on the moon. Regolith is composed of tiny, jagged particles produced by space weathering which have unique effects on light reflecting off the surface. Though it is often compared to sand, optically driving through regolith might be better compared to a blanket of snow or powder. In experimentation, we used JSC-1A regolith simulant to form terrain features. This simulant is the gold standard for regolith from the lunar *mare regions*. Though JSC-1A was designed for mechanical properties, the fidelity of composition and particle size distribution, which are the main factors for bulk appearance, makes it a plausible optical simulant. One source of inaccuracy is that regolith at the poles is thought to represent the "highlands" and not mare regions, with bulk albedo differences of 30% and 12% respectively. While, JSC-1A is widely available to NASA, highlands simulant has been impractical to obtain in large quantities. We believe that testing with mare simulant is still valuable as the bulk BRDFs[3] of the two materials are similar, the lower albedo serves as a worst-case, and there is still significant uncertainty about the polar regions. As a reference, we have secured a small quantity of a highlands regolith simulant, NU-LHT, for testing in a small 1m x 1m sandbox. This data is offered as an addendum download (Terrain 13) to the main dataset.

[3] Bi-directional Reflectance Distribution Functions – both are Hapke materials with similar particle size distributions

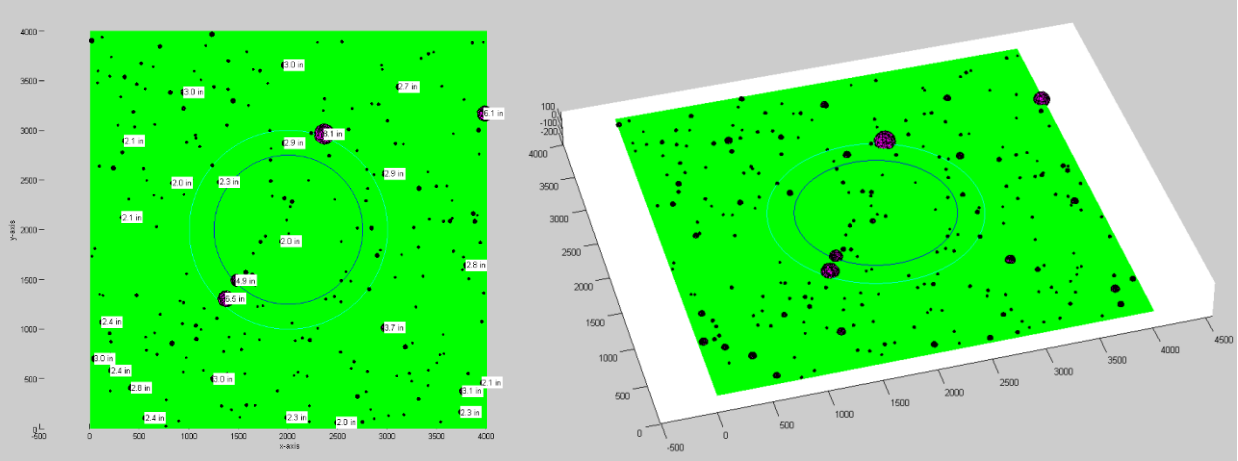

Figure 1. Computer Generated Terrain Model. Note: distribution of rocks > 10cm is precise, while rocks <10cm are mostly randomly placed.

Rocks: We have selected an assortment of rock shapes and sizes to simulate a variety of interesting geometric conditions. These include volcanic-origin rocks with pitted surfaces, angular fracture rocks, and smoother rocks to simulate space weathering. Rock placement in half of the terrains tested specific user-designed obstacle conditions (e.g. a free path in an obstacle field), while in the other half, a random spatial placement has been designated using sampling from a lunar power law size-frequency distribution[4]. Rocks between approximately 0.5-5cm in diameter are not placed in the terrain. This was a decision made to balance experimental setup time with visual/geometric significance of these rocks. Small rocks (<0.5cm) were previously mixed in with the JSC1A. See Table 3 for descriptions of all the conditions used to design each terrain.

Illumination: Solar illumination is provided by a constellation of four lights distributed azimuthally about the sandbox at critical angles (positions are listed in Table 1). Sequential illumination from the lights simulates planetary rotation at the poles. Each source is a 1kW tungsten-halogen light mounted on a stand 1.3 meters tall and placed 4 meters distance from the center of the terrain, giving an apparent elevation angle of ~8 degrees. Circular reflectors are selected such that sources subtend an angular diameter similar to the sun as seen from the moon[5]. "Barn doors" are also installed on each source to minimize stray light to the area of the terrain. Total output power is 80k lumen, and is empirically 16x less intense than direct sunlight[6] dropping to about 25x at the periphery of the sandbox. We believe that a single photographic scalar of 20x suffices for most purposes, such as predicting lunar optical settings.

There are several sources of error arising from the use of our particular illumination setup. The tungsten-halogen bulbs output a continuous spectrum with 3200K perceptual light color which is substantially "warmer" than the illumination at the poles[7]. Though dichroic filters exist to shift the illuminant to a near 6000K "daylight," this would result in a 50% reduction in output. A decision was

[4] Surveyor 7, JPL Technical Report 1968
[5] 0.5 degree diameter
[6] Direct sunlight is about 110K lux
[7] Approximately 5900K

made to use monochromatic cameras[8] and focus on the dynamic range instead of color fidelity. Testing constraints and mitigation of atmospheric scattering necessitated used near-field illumination. Thus, the source angles and intensity distributions vary with position on the terrain and an assumption of uniform parallel rays may not suffice for applications such as photoclinometry. It is possible to backsolve for a more accurate light field given the surface extent of the sandbox[9]. Lastly, the sandbox has a raised lip which casts a linear shadow at the edge of the terrain. This shadow has unrealistic appearance, but otherwise should not affect results. We advise anyone concerned about edge effects to mask results to the central portion of the scene.

Data Collection

The stereo imaging setup was positioned along the periphery of the terrain for data collection. The stereo cameras were rigidly mounted with a high-accuracy[10] Faro x330 LIDAR sensor, which provided ground truth for each stereo pair. Positions of the light sources and desired sensor views were calibrated with a total station (scene spatial coordinates are given in Table 4). However, experimental movement of the sensor platform resulted in only approximate pose. Each scene is imaged across all illumination conditions and from several locations in order to collect more data. A schematic diagram of the experimental setup is given in Figure 2. Dimensional limitations in the laboratory required asymmetric imaging and illumination angles.

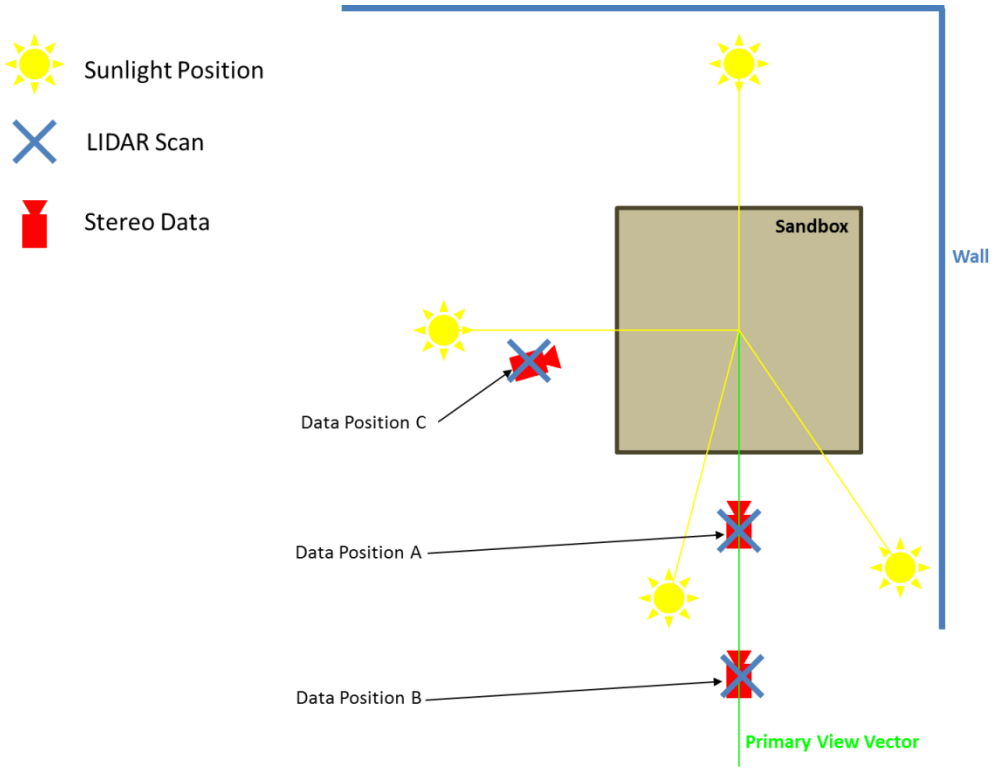

Figure 2. Bird's eye view of the experimental setup with illumination and data collection positions.

[8] Which are more likely for our mission archetype
[9] This may be released in future work
[10] Sub 2mm single shot range accuracy

Imaging: The stereo pair uses a 30cm baseline which was selected for similarity to prior NASA rovers like MER and MSL. The camera poses are determined from a sensor mast height of 1.85 meters. The view angle is fixed at approximately 23 degrees negative pitch to the ground plane. Stereo pairs are captured with a bracket of images from short to long shutter speed (32ms to 8192ms in 1-EV escalations). Optical parameters of experimentation are listed in Table 2.

Figure 3. Data capture rig with stereo cameras, mast illumination, and ground truthing sensor (left). Photo of experimental setup showing relative position of the terrain, lights and sensors (right).

Ground Truthing: Ground truth depth data is collected using a Faro x330 LIDAR which is registered to the stereo imaging pair. The LIDAR uses a rotating mirror to slice a 360 degree vertical plane with the beam. The device is then rotated about the azimuth to provide full 360x180 degree coverage. The regolith testbed is contained within a 90x90 front cone of the scanner. We select scanning parameters such that a single scan of sufficiently high density such that one or more LIDAR readings correspond to a single imaging pixel on average. Because both sensing systems are offset, occlusions are possible between stereo and LIDAR point clouds; however, because of the large scene distances, we believe this to have minimal effect on calculating performance. Scanning and imaging with directional lighting are conducted separately to ensure there is no possible cross-talk.

Table 1. Experimentation Parameters

Parameter (# values)	**Possible values**
Terrain Type (12)	12 different obstacle distributions
Camera Elevation (1)	23° relative to center of terrain
Camera Azimuth (3)	A (0°) ,B (0°),C (280°)
Scene Distance (2)	{1.5, 4.0} meters relative to center of terrain
Rover Lights (2)	{off, on} *35W external illumination carried for scene
Sun Azimuth (5)	{30°, 180°, 270°, 350°}, {NO_SUN}
Sun Elevation (1)	8°
Exposure Time (9)	{32, 64, 128, 256, 512, 1024, 2048, 4096*, 8192*} milliseconds *available only for some trials

Table 2. Imaging Parameters

Parameter	Value
Camera	AVT Manta G-235
Resolution	1936x1216, 5.86um pixel size
Sensor	1/1.2” Monochrome, no IR-cut filter
Bit Depth	12bit (packed in 16bit)
Sensitivity	Linear gamma used, 0 gain
Stereo Baseline	30cm
Lens	Goyo Optical GM12HR41216MCN
Focal Length	12mm
Aperture Setting	F/8.0
Solar Radiance	27600 Lumens
Solar Diameter	0.8° relative to center of terrain
Illumination scalar	1/16x the brightness of direct sunlight
Solar spectrum	Tungsten-Halogen, FEL 1000W, 3200k

Table 3. Terrain Conditions

Scene #	Description
Terrain 1	Mostly large rocks, 15 rocks total.
Terrain 2	A very large rock in the foreground casting a large shadow and a distribution of 16 rocks elsewhere.
Terrain 3	28 randomly placed rocks, sampled from distribution.
Terrain 4	28 randomly placed rocks, sampled from distribution.
Terrain 5	Random, sampled from distribution.
Terrain 6	Featureless, no rocks.
Terrain 7	Split, large rocks on one side and small rocks on the other. There are 19 rocks in this terrain.
Terrain 8	1 Large cluster of rocks with several small rocks, 16 total.
Terrain 9	Random distribution weighting medium sized rocks (~10cm), 20 total.
Terrain 10	Random distribution weighting medium sized rocks (~10cm), 20 total.
Terrain 11	A fresh 1.5m diameter crater with prominent rim, smooth slopes and 9 rocks.
Terrain 12	An old 1.5m crater simulating space weathering, in-fill, and 20 rocks.
Terrain 13*	A bonus setup with a sandbox using the NU-LHT simulant and 3 small rocks. Azimuthal lighting is uncalibrated, but stereo images are ground truthed.

Dataset Organization

The POLAR dataset is organized in an easily navigable directory tree structure. The directories and files have human-readable names which can be used to backsolve the exact experimental parameters used to capture the data.

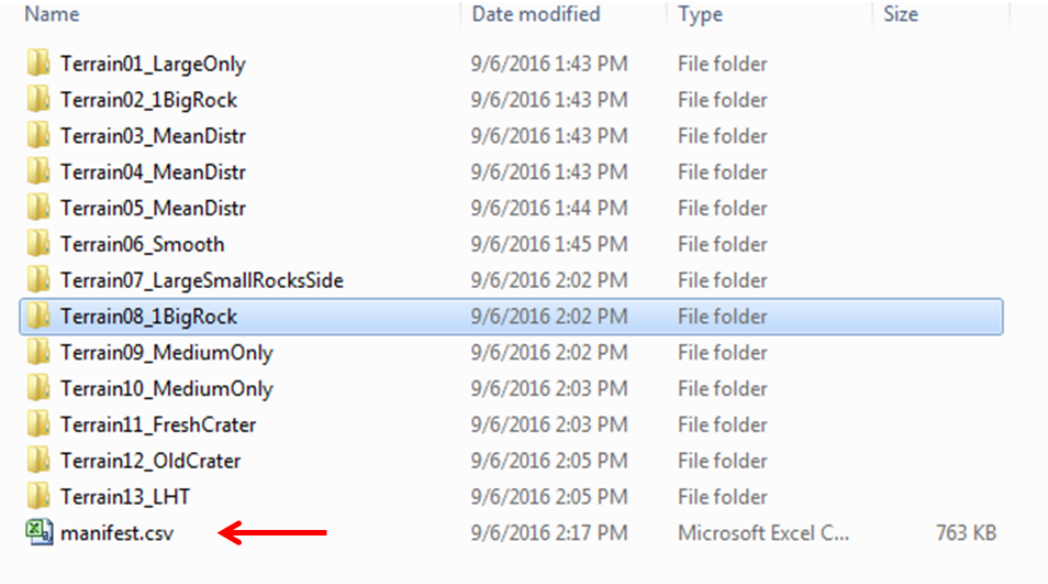

Figure 4. Top Level Directory Structure. The manifest file contains information about the entire contents of the dataset. Data is organized in directories named for the Terrain conditions in which the data was collected.

The root directory of the dataset contains data organized by Terrain type as described in Table 3. There are 12 terrain conditions and one bonus un-registered Terrain of a different surface material (Figure 4). Additionally, we provide a manifest CSV for quick lookup of specific image/parameter pairs in the directory tree. The manifest file is ASCII and can be opened in Excel – each row contains a unique image or ground truth point cloud and the columns contain information such as relative path and experimental parameters (Figure 5). We recommend use of the manifest file for quickly selecting subsets of images to test desired conditions.

	A	B	C	D	E	F	G	H	I	J	K
1	fileid	filename	sensor_ty	terrain_ic	terrain_d	sensor_p	scene_dis	rover_ligl	sun_azim	exposure	uniqueid
2	1	./Terrain01_LargeOnly/Faro_ascii/PosA.xyz	ground_ti	Terrain01	LargeOnly	PosA	-1	-1	-1	-1	-1
3	2	./Terrain01_LargeOnly/Faro_ascii/PosB.xyz	ground_ti	Terrain01	LargeOnly	PosB	-1	-1	-1	-1	-1
4	3	./Terrain01_LargeOnly/Faro_ascii/PosC.xyz	ground_ti	Terrain01	LargeOnly	PosC	-1	-1	-1	-1	-1
5	4	./Terrain01_LargeOnly/PosA_1500_Loff/Sun_180/	stereo_L	Terrain01	LargeOnly	PosA	1500	off	Sun_180	1024	2.5E+08
6	5	./Terrain01_LargeOnly/PosA_1500_Loff/Sun_180/	stereo_L	Terrain01	LargeOnly	PosA	1500	off	Sun_180	128	2.5E+08
7	6	./Terrain01_LargeOnly/PosA_1500_Loff/Sun_180/	stereo_L	Terrain01	LargeOnly	PosA	1500	off	Sun_180	16	2.5E+08
8	7	./Terrain01_LargeOnly/PosA_1500_Loff/Sun_180/	stereo_L	Terrain01	LargeOnly	PosA	1500	off	Sun_180	256	2.5E+08
9	8	./Terrain01_LargeOnly/PosA_1500_Loff/Sun_180/	stereo_L	Terrain01	LargeOnly	PosA	1500	off	Sun_180	32	2.5E+08
10	9	./Terrain01_LargeOnly/PosA_1500_Loff/Sun_180/	stereo_L	Terrain01	LargeOnly	PosA	1500	off	Sun_180	512	2.5E+08
11	10	./Terrain01_LargeOnly/PosA_1500_Loff/Sun_180/	stereo_L	Terrain01	LargeOnly	PosA	1500	off	Sun_180	64	2.5E+08
12	11	./Terrain01_LargeOnly/PosA_1500_Loff/Sun_180/	stereo_R	Terrain01	LargeOnly	PosA	1500	off	Sun_180	1024	2.5E+08
13	12	./Terrain01_LargeOnly/PosA_1500_Loff/Sun_180/	stereo_R	Terrain01	LargeOnly	PosA	1500	off	Sun_180	128	2.5E+08
14	13	./Terrain01_LargeOnly/PosA_1500_Loff/Sun_180/	stereo_R	Terrain01	LargeOnly	PosA	1500	off	Sun_180	16	2.5E+08
15	14	./Terrain01_LargeOnly/PosA_1500_Loff/Sun_180/	stereo_R	Terrain01	LargeOnly	PosA	1500	off	Sun_180	256	2.5E+08
16	15	./Terrain01_LargeOnly/PosA_1500_Loff/Sun_180/	stereo_R	Terrain01	LargeOnly	PosA	1500	off	Sun_180	32	2.5E+08
17	16	./Terrain01_LargeOnly/PosA_1500_Loff/Sun_180/	stereo_R	Terrain01	LargeOnly	PosA	1500	off	Sun_180	512	2.5E+08
18	17	./Terrain01_LargeOnly/PosA_1500_Loff/Sun_180/	stereo_R	Terrain01	LargeOnly	PosA	1500	off	Sun_180	64	2.5E+08
19	18	./Terrain01_LargeOnly/PosA_1500_Loff/Sun_270/	stereo_L	Terrain01	LargeOnly	PosA	1500	off	Sun_270	1024	2.5E+08
20	19	./Terrain01_LargeOnly/PosA_1500_Loff/Sun_270/	stereo_L	Terrain01	LargeOnly	PosA	1500	off	Sun_270	128	2.5E+08

Figure 5. Contents of the manifest file. The file is an ASCII comma separated variables (CSV) format which can be opened in Excel or a text editor. Each row of the file is an entry for one data file, while the columns describe values for each experimental condition pertaining to the file (such as scene distance or lighting condition).

Within each terrain directory, the data is organized by camera position. There are three types of data here: (1) a stereo calibration file, (2) camera pose directories and (3) a single directory with ground truth point clouds (Figure 6). See the next section (File Formats and Details) for details of the files and how to read them.

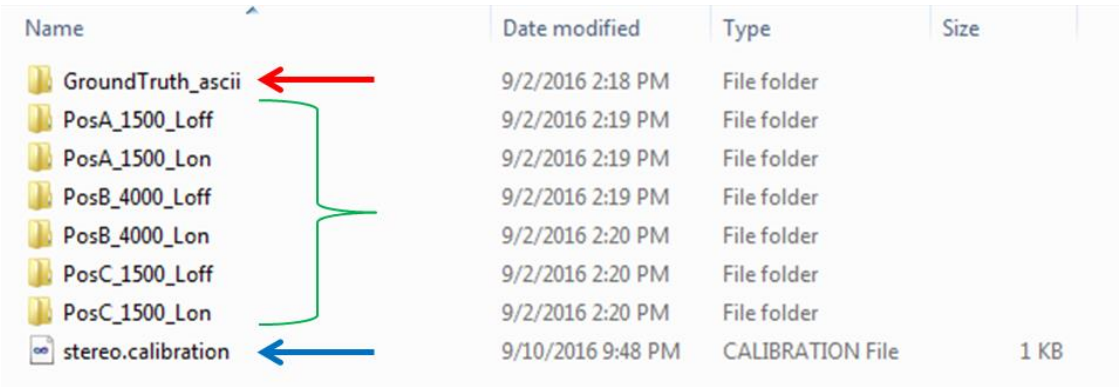

Figure 6. Directory structure within each terrain type. Each terrain directory contains ground truth LIDAR scans (red), stereo data organized by sensor position and active illumination (green), and a calibration parameters file (blue).

Position subdirectories have name strings of the form Pos{A,B,C}_{1500,4000}_{Lon, Loff}. These indicate the position, the scene distance in mm, and whether active flash photography (mimicking rover lights) was used. Further down the tree within each camera position directory are a series of sun angle directories for each of the illumination conditions (Figure 7). Finally, within each of the illumination directories are HDR image stacks for the left and right cameras. The image filenames indicate the

camera (L or R), the exposure time in miliseconds, and a unique identifier serial number. All of the images are taken from the same static location, modulating only shutter time.

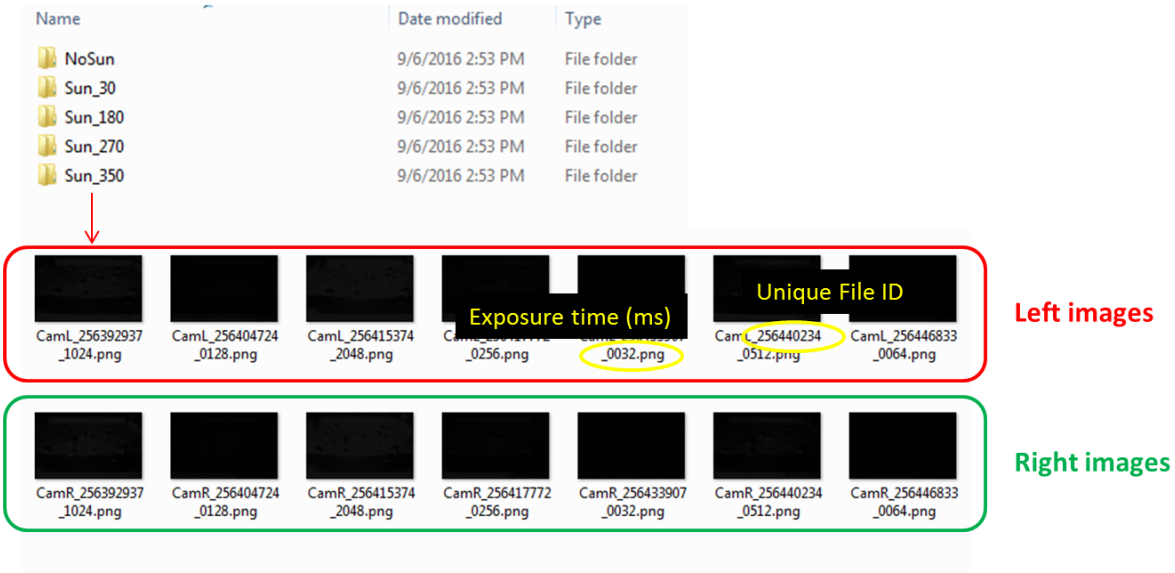

Figure 7. Lighting conditions and stereo image HDR brackets contained in data directory.

File Formats and Details

Image Files

Images are stored in lossless 16-bit PNG format with a single channel (monochrome). The sensor has a 12-bit native depth, so the full range is not used. Instead, valid pixel values are from [0, 4095] inclusive. Anything outside this range should be considered an erroneous reading. Keep in mind the clamped nature of the data when converting or viewing the dataset images in software expecting 8-bits. We provide the raw images from the camera, without rectification. This is so that the end user can do their own rectification/undistortion with a custom interpolation method.

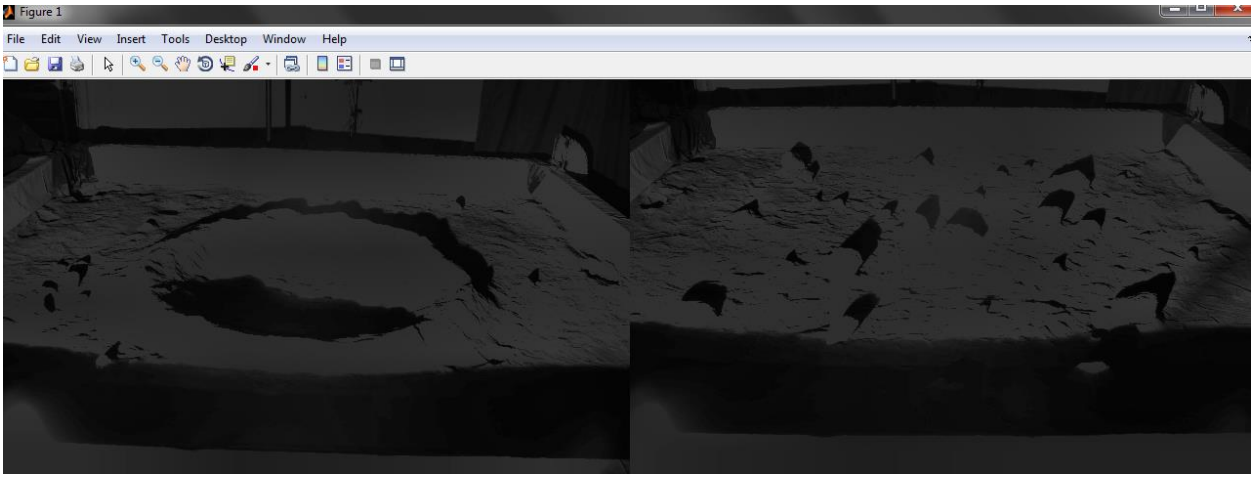

Figure 8. An example of images collected of two different scenes (terrains), but with the same optical parameters and environmental conditions.

Stereo Calibration File

The stereo calibration file is a YAML formatted ASCII file compatible with OpenCV. The file contains is one property per line in the format "PROP_NAME: val1 val2 etc". The properties stored in the file include the camera intrinsics (focal length, resolution, plumb-bob distortion), left-to-right transformation (stored as a rotation matrix and translation vector) and the left camera-to-LIDAR transformation (Figure 9). Matrices are stored in row-major order. The distortion coefficients are stored in the order: [Radial1 Radial2 Tangential1 Tangential2]. While in practice, calibrations were kept across multiple experiments, the calibration file for each terrain must be assumed to be unique.

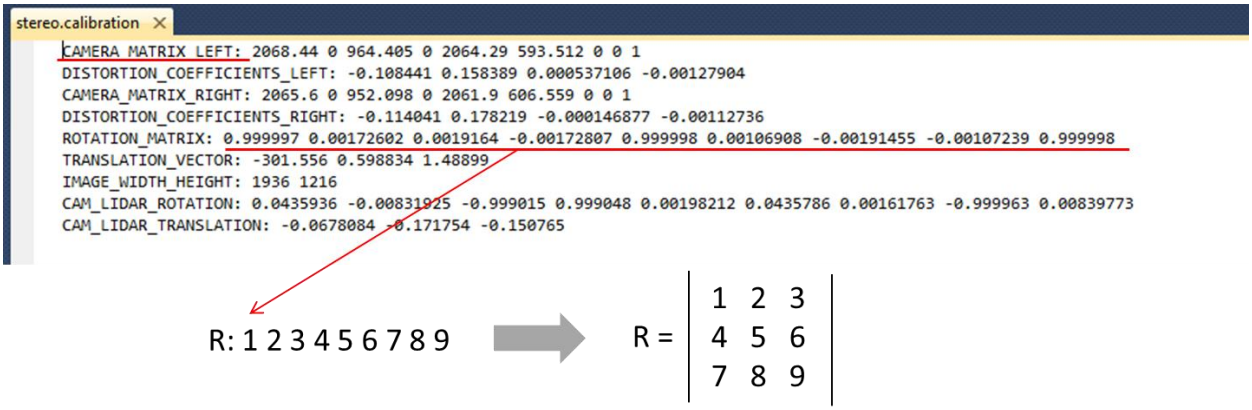

Figure 9. Contents of the Stereo Calibration YAML file.

Ground Truth LIDAR (XYZ) Point Cloud

The GroundTruth_ascii directory contains ASCII (.XYZ) point clouds captured from the LIDAR scanner (Figure 10). Each row is a single point, with 6 space-delimited attributes. The 6 attributes include the (row, column) raster order of the point, the 3D Cartesian coordinates of the point in meters, and the returned LIDAR intensity value.

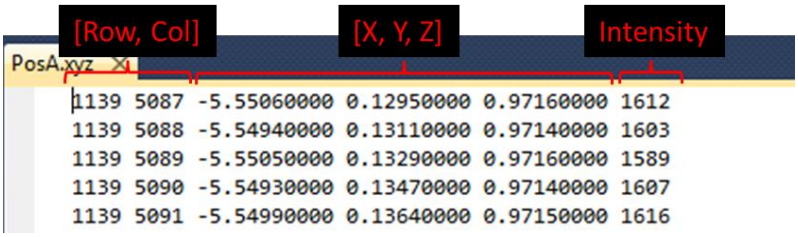

Figure 10. Format of ASCII point cloud files

The point clouds are manually cropped such that only points both within the stereo field of view and inside the square testbed area where scenes are created (Figure 11). These point clouds must be transformed to the left camera frame using the cam-to-LIDAR transformation given in the calibration files. To find the pixel correspondence for each point, (1) transform to the left camera frame, (2) project into pixel space using camera matrix and (3) calculate distorted pixel coordinates using distortion coefficients.

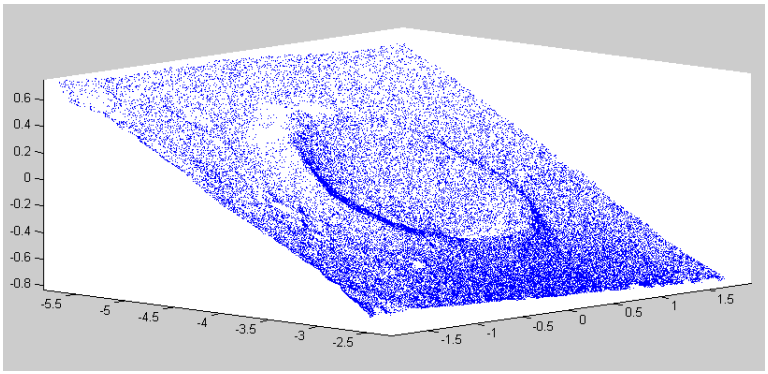

Figure 11. Ground truth point clouds have been cropped so that only point that fall within the 4mx4m sandbox are included. Thus, by projecting these points into the image plane.

LIDAR Organized Point Cloud (optional)

We have provided organized point clouds for convenience as optional downloads (Figure 11). Organized point clouds are NxMx3 images where each pixel's tri-channel value corresponds to the ray location in Cartesian coordinates. The organized point clouds have been projected and undistorted so that they

correspond pixel-wise to the left stereo camera intensity image (rectified). An average is used whenever multiple LIDAR readings map to the same pixel location. It is advised to create your own maps if errors such as mixed-pixel effects are important to your application.

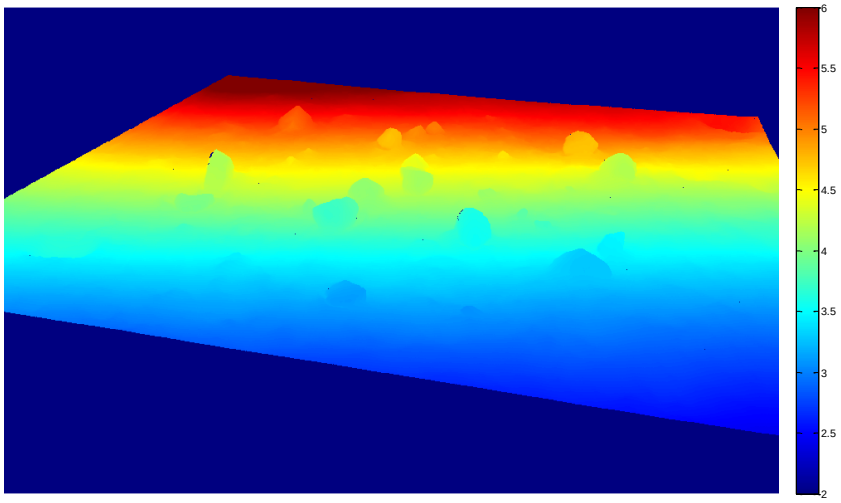

Figure 12. Organized point cloud generated from ground truth LIDAR.

Radiometric Calibration File (optional)

Radiometric calibration files are included for the stereo cameras as optional downloads. Radiometric curves map valid pixel values [0,4095] into a linear floating point range giving the illumination incident on the sensor plane (relative exposure). Figure 13 illustrates that the cameras we use are already quite linear, which is what most users will want. However, we include calibrations as there are some minor deviations which could be important in some HDR processes. The relative exposure between cameras can be directly compared.

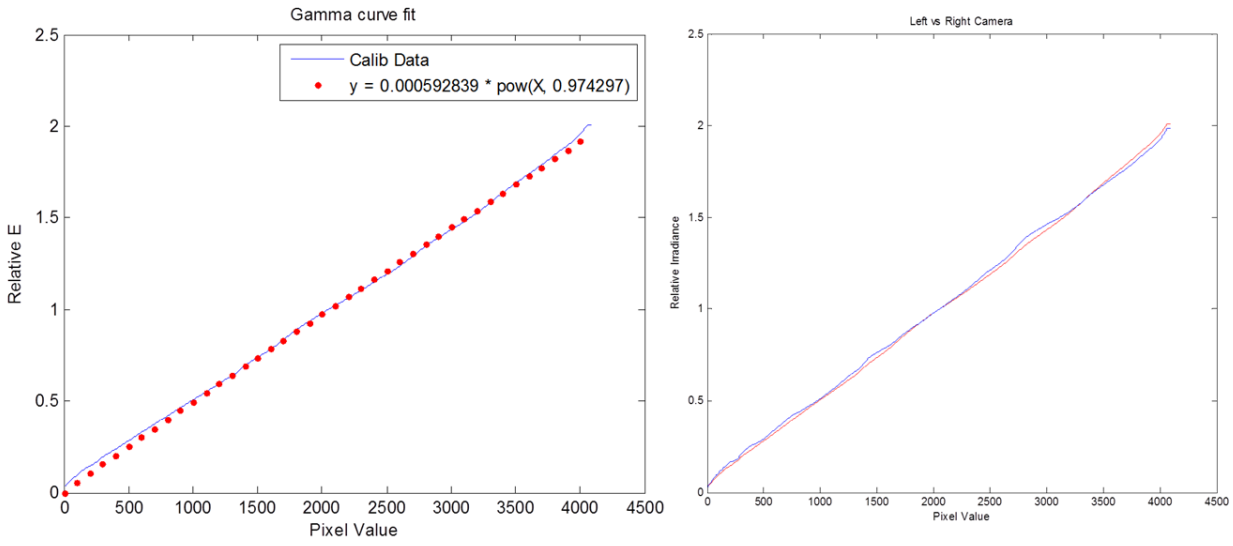

Figure 13. Gamma curve fit to camera response functions (left) and comparison between right and left cameras in the stereo pair demonstrating similar response.

We present the radiometric data as an ASCII formatted lookup table. Each of the 4096 rows corresponds to pixel value, while the columns correspond to the exposure for R, G, B channels (Figure 14). Since the released portion of the dataset uses monochrome cameras, all columns should be the same. If L is the lookup table mapping and P is the pixel value P, the linear exposure E is therefore E = 2^L(P).

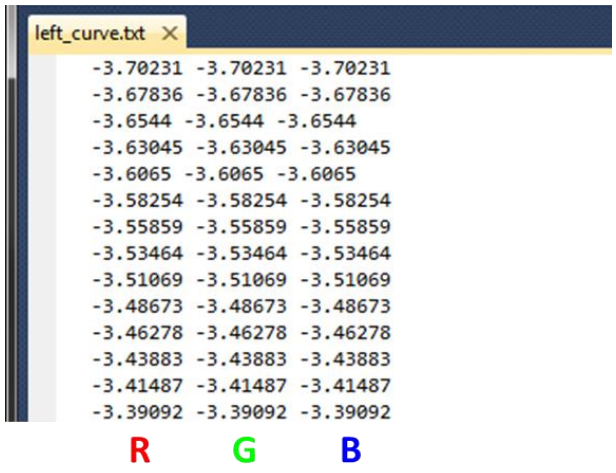

Figure 14. Example radiometric calibration curve file.

Appendix

Survey Coordinate Locations

Cartesian coordinates of the illuminants, sensor views and sandbox corners have been surveyed with a total station. This information may be valuable for those needing precise illumination and view vectors; for example in Shape-from-Shading application. This information is given in the (Table 4, Figure 15) below. Due to our tripod setup, the view positions are not as precise as the other control points (~5cm error). It is possible to backsolve more precise view coordinates using the known control points and PnP or trilateration from the images or organized point maps.

Table 4. Survey Target Points

Surveyed Target	3D Spatial Coordinates (xyz meters)
Sun30	[1.422, -3.594, 0.583]
Sun180	[0.475, 3.544, 0.556]
Sun270	[-4.109, 0.271, 0.591]
Sun350	[-0.896, -4.026, 0.551]
Camera PosA*	[-0.303, -3.541, 1.346], [-0.002, -3.532, 1.344]
Camera PosB	[-0.303, -5.500, 1.345], [-0.002, -5.500, 1.344]
Camera PosC	[-3.540, -0.316, 1.346], [-3.479, -0.611, 1.344]
Sandbox BL Corner	[1.998, -1.982, -0.006]
Sandbox BR Corner	[-1.998, -2.000, 0.011]
Sandbox TL Corner	[-1.986, 1.998, 0.004]
Sandbox TR Corner	[1.986, 1.984, -0.009]
Terrain Midpoint	[0.0, 0.0, 0.0] (defined origin)

* Camera positions are approximate due to the need to shift the sensor around

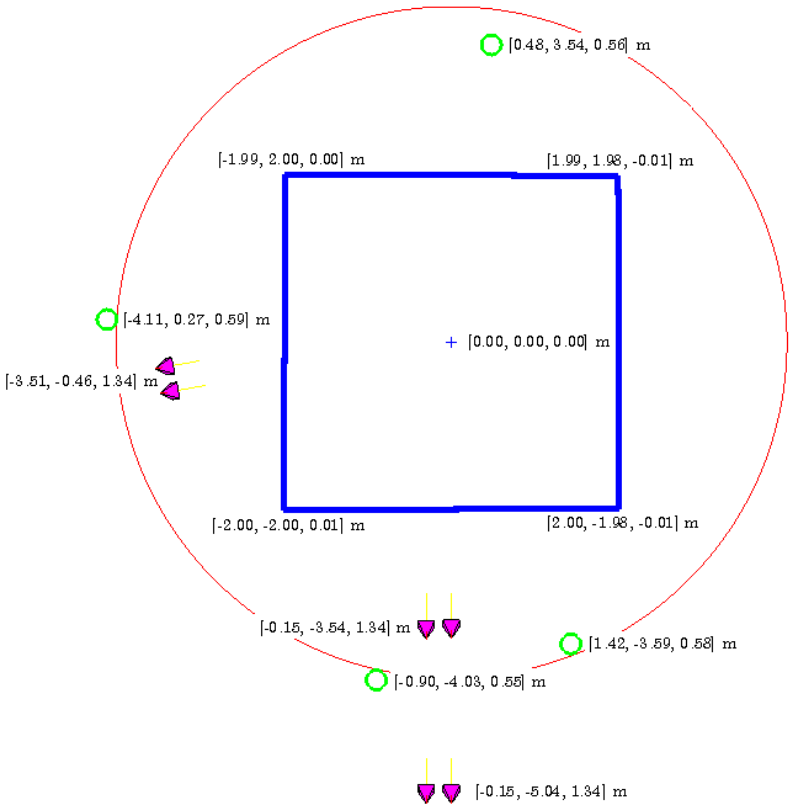

Figure 15. Map of Survey Targets