Strobe Lit High Dynamic Range Stereo Imagery for Dark Navigation

Kyle Husmann[1] and Liam Pedersen[2]
[1]California Polytechnic State University, San Luis Obispo, California, USA
[2]NASA Ames Research Center / Carnegie Mellon University, Moffett Field, California, USA

Abstract

Permanently shadowed lunar craters are high priority targets for rover exploration because the possibility they harbor water ice. Stereo cameras are well-established robotic sensors for navigation, but require ambient illumination to operate.

This paper addresses the challenges of doing stereo vision in a dark, lunar like environment using rover mounted, unstructured LED lighting to illuminate the area. LED spotlights are attractive because of their low power consumption, ruggedness and the large number of times they can be pulsed.

Because of inverse squares illumination drop off with distance, nearby terrain is usually saturated in images and far terrain barely exposed. We combine multiple low dynamic range (LDR) images with different exposures to construct high dynamic range (HDR) images suitable for stereo mapping.

This paper investigates the effects of HDR imaging on stereo coverage and accuracy. We test various lighting geometries, comparing pulsed red LED lighting, a halogen floodlight and sunlight.

Key words: navigation; HDR; stereo; sensors; dark lunar craters.

1. Introduction

Terrestrial radar measurements [1] of the Lunar south polar regions determined the existence of craters with permanently shadowed interiors sufficiently cold to harbor water ice (Figure 1). This is corroborated by the radar echoes and Lunar Prospector neutron spectrometer data that suggest localized concentrations of H_2O up to 40% by weight (collected references in [2]).

The potential for water makes permanently shadowed lunar craters high priorities for exploration. The calculated H_2O abundances are very uncertain, and the local distribution unknown. NASA's LCROSS mission will impact a polar crater in 2009 to attempt a confirmation of H_2O abundance, but the definitive assessment of the water/hydrogen distribution in the crater floor regolith is likely to require exploration with a neutron spectrometer and drill-equipped robotic rover.

Figure 1: Shackleton Crater (diameter 19km), near the lunar South Pole, at 89.54° South latitude, is one of several permanently shadowed craters thought to harbor water ice. (Photo: ESA SMART–1)

Terrain sensing for scientific data collection, hazard detection and localization is one particular problem that must be overcome to navigate a vehicle inside the permanently dark craters. Stereo vision has a proven track record for rover navigation, both terrestrially and in space, particularly the Mars Exploration Rover's Spirit and Opportunity. To use stereo vision in a dark crater we would need either an artificial light source on the rover, or a camera sufficiently sensitive to get usable images from ambient starlight and sunlight reflected from crater rims.

This paper addresses how to obtain usable stereo images to map a lunar like terrain using an unstructured artificial light source mounted close to the cameras (as would be needed for this system to be deployed on a rover). The main difficulty posed by

(a) (b)

Figure 2: Stereo cameras (a) and LED spotlight (b)

this configuration is the rapid (1/distance^2) fall off in illumination intensity resulting in a high illumination dynamic range, and the possible lack of visible surface texture. Visible surface texture results from shadows, differing surface orientations and albedos. Camera's close to, or directly in front of, the light source will not see shadows. Apollo astronauts reported difficulty in seeing lunar surface with the sun directly behind them [3].

We solve the high dynamic range problem by combining multiple low dynamic range (LDR) images of a scene into a single high dynamic range (HDR) image on which we do stereo. We compare the accuracy and the percentage of "Good Pixels" – pixels around which there is sufficient texture to do stereo correlation – of HDR stereo to LDR stereo using sunlight to illuminate the scene. We also investigate various plausible camera–lighting geometries to see which gives best results.

LED spotlights are an attractive light source because they can be pulsed many times, have low power consumption and are rugged solid state devices. Very significant power savings, on the order of 1000×, can be achieved by turning the lights on only when the camera shutter is open. We test and compare a red pulsed 24 W LED spotlight with a 2 kW tungsten-halogen work site light.

The following sections review the basic mathematics of image intensity as a function of distance and other factors, characterize the performance of our hardware on JSC-1AF lunar regolith simulant and extrapolate LED lighting power requirements for rover obstacle avoidance. We review the basics of HDR images and present results obtained at NASA Ames's outdoor "Marscape" test site, which we subsequently analyze.

2. Illumination Basics

The total light energy E incident on a camera pixel looking at an object distance D away from a light source with power P, uniformly emitting light over the solid angle Ω (steradians) is given by

$$E = k_1 \frac{PT\rho}{\Omega f_{\#}^2 D^2} \tag{1}$$

Where k_1 is a constant (proportional to pixel area and optical collection efficiency), T is the exposure time (sec), $f_{\#}$ is the camera F-number and ρ is the object albedo. Assuming a linear sensor response, it follows that the exposure time to obtain a maximum pixel response is given by:

$$T = k_2 \frac{\Omega f_{\#}^2 D^2}{P\rho G} \tag{2}$$

Where G is the (dimensionless) sensor amplifier gain (sometimes given in dB - in which case $G = 10^{(dB/20)}$), k_2 is another constant (proportional to pixel area, optical collection efficiency and sensor conversion efficiency).

Stereo ranging involves matching points between two or more images. This is usually done by maximizing the correlation between one image and shifted windows in the other images. Doing this requires sufficient image texture that a meaningful correlation peak exists. Images must be exposed correctly such that texture variations in the region of interest are retrieved with enough intensity precision to do stereo.

Assuming an exposure time T such that camera pixels attain 50% response for objects at distance D, then the range of distances within which average pixel intensity is between 25% and 75% of full scale deflection is $\sqrt{1/1.5}D$ to $\sqrt{1/0.5}D$ or $0.8D$ to $1.4D$. Interestingly, other factors such as detector sensitivity, optics, or light power do not change this. That is, for a given exposure time, the range of distances within which acceptable exposures result depends only on how close to the sensor extreme ranges we are permitted to go. Increased pixel intensity resolution enables a reduction in the bottom threshold.

3. High Dynamic Range Imaging

High dynamic range (HDR) imaging, either with a logarithmic response sensor, or by combining multiple LDR images, enables both increased dynamic range and increased intensity resolution at the lower limits of light intensity.

HDR images can be obtained using ordinary commercially available cameras. Given multiple LDR images $\{I_i | i = 1...N\}$ of the same scene, each taken with a different exposure time T_i, an HDR image I is given by

$$I(x, y) = \frac{1}{Z} \sum_{i=1}^{N} w_i(x, y) \frac{f^{-1}(I_i(x, y))}{T_i} \tag{3}$$

where w_i are weights, $Z = \sum_{i=1}^{N} w_i$ is a normalizing constant, and f is the sensor response function mapping image plane light intensities to pixel response (Figure 3), which can be estimated from multiple images of a scene at with different integration times. It is important to note that f *must* be flat outside the dynamic range of the camera to get good results (i.e. further exposing already saturated pixels will not result in higher pixel values).

The weights should be chosen to emphasize pixels near 50% of the dynamic range, and should be zero for saturated or under-exposed pixels.

Note that while the LDR images I_i pixel values are typically 8-12 bit integers, the HDR image I pixel intensities are higher precision floating point numbers. Details of our HDR implementation, and further collected references are in [4].

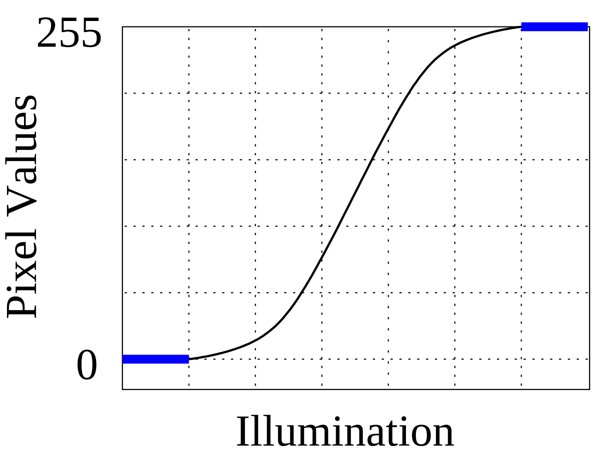

Figure 3: Typical CCD light intensity response curve. Response *must* be flat outside the dynamic range of the sensor (thick blue segments).

4. Hardware Characterization

Our hardware setup consisted of a modified Advanced Illumination SL6404 LED spotlight (Table 1) and a Point Grey Flea camera (Table 2). The trigger for the spotlight was connected through a voltage amplifier to a GPIO on the camera, so that the spotlight would fire in sync with the acquisition of images.

Table 1: Custom Advanced Illumination SL6404 Red LED Spotlight

Property	**Description**
Voltage	24V
Current	1A
Trigger to Pulse Delay	$< 10\,\mu$s
Max Repeat Rate	40 Hz
Pulse Width	5 ms to 40 ms
Field of View	*approx* 15 °

Table 2: Point Grey Flea Camera

Property	**Description**
Focal Length	2.8 mm
Aperture	$f/1.2$

We characterized this hardware combination by determining the exposure time necessary to just barely saturate a sheet covered with JSC-1AF lunar regolith simulant [5], which was placed at different distances from the camera and spotlight in a darkroom. Coriander 2.0 [6] was used to control the camera and determine when the regolith was barely saturated (Using the “Exposure check” feature). Since the camera had no reliable way to set the aperture, we used the maximum aperture of $f/1.2$ for all of our images. Also, we used a gain of 29 dB, the highest gain giving acceptable images.

If we examine T as a function of D^2 (keeping all other parameters constant), the slope A of the resulting linear function will be:

$$A = \frac{k_2 \Omega f_{\#}^2}{P \rho G} \tag{4}$$

For our hardware combination, A was determined to be $0.3\,\mathrm{ms/m^2}$ using a least squares regression fit. Extrapolating from this result, we deduce an exposure time of approximately 30 ms to get a full pixel response for objects at 10 m (Figure 4). Assuming an average 50% response is sufficient to resolve

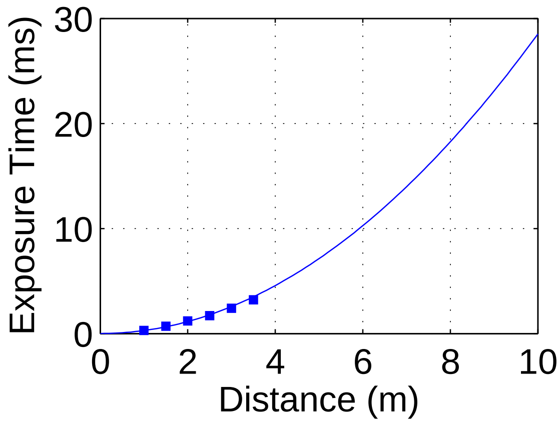

Figure 4: Exposure time as a function of distance for our hardware combination. Extrapolating from our data, we deduce an exposure time of approximately 30 ms to get a full pixel response for objects at 10 m.

enough texture for stereo based navigation, this equates to a 15 ms exposure time. Increasing the lighting angle to a 90° fov requires 36× increase in exposure time. For continuous navigation, a short exposure time is desirable to avoid motion blur. Assuming 10 ms exposure time, 10 m look-ahead and 90° fov would require a 56× increase in peak light power.

5. HDR Experiment

An experiment to test the utility and accuracy of HDR images with artificial lighting for stereo was done at NASA Ames Research Center's Marscape test site (Figure 5). Marscape with its uneven, undulating terrain and occasional rocks and is (mostly) free of vegetation and was the most lunar-like site at our disposal.

The purpose of the experiment was to gauge the effectiveness of HDR for stereo – i.e the fraction of pixels suitable for stereo correlation ("Good Pixels"), the accuracy of HDR stereo and the effect of various camera-lighting geometries on accuracy and effectiveness. As both accuracy and effectiveness depend on many factors besides lighting, we compare HDR images performance against sunlit images of the same scene (sunlight being the stereo camera "gold standard" for unstructured scene illumination).

5.1. Setup and Procedure

A stereo pair of Point Gray Flea cameras (Table 2) were stably mounted on a tripod so as to view the same scene for the duration of the experiment. Initial baseline LDR images were taken during the afternoon with the sun at approximately 90° to the camera direction. Subsequent LDR image sequences were taken after dark using illumination from a 2 kW tungsten-halogen floodlight from Home Depot and the red LED spotlight (Table 1) in various configurations with respect to the cameras (Table 3).

The stereo cameras were calibrated using the Matlab calibration toolbox from [7]. NASA Ames's Vision Workbench package [8] was used to calibrate the camera luminance response curves, and to compute the HDR images from the LDR image sequences.

The camera's were connected to a linux host computer via a IEEE 1394 ("firewire") bus (Figure 6). A rising-edge strobe signal from a camera GPIO pin triggered the LED spotlight flash at the precise moment the camera took an image.

The Point Gray Flea cameras automatically synchronize themselves to other Flea camera's on the same IEEE 1394 bus. By comparing time stamps from acquired images from each camera we ensured that all stereo pairs were acquired within 125 μs of one another.

Figure 5: NASA Ames Reseach Center's Marscape robot test facility.

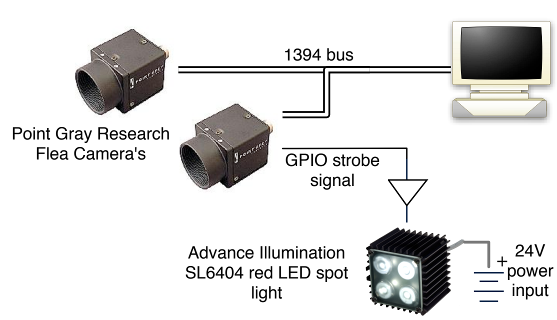

Figure 6: Hardware Diagram

Table 3: Table of lighting configurations. Camera pair is at 1.5 m above the ground.

Configuration	Description
Baseline Sunlit	Ambient Sunlight
LED Spotlight 1	Single LED spotlight placed slightly to the left of cameras, at a height of 1.5m
LED Spotlight 2	Single LED spotlight placed slightly to the left of cameras, at a height of 1.2m
LED Spotlight 3	Single LED spotlight placed slightly to the left of cameras, at a height of 0.8m
LED Spotlight 4	Single LED spotlight 1.5m to the left of cameras, at a height of 0.8m
Floodlight	A floodlight 1m to the right of the cameras

5.2. HDR Effectiveness

Figures 9 and 10 give quantitative representations of how the exposure time and HDR processing effect the number of Good Pixels. Clearly, the HDR algorithm boosts the number of Good Pixels considerably, giving it almost as good coverage as the Baseline Sunlit LDR image pair in the case of the Floodlight configuration.

5.3. Accuracy of stereo reconstruction

We found that the depth reconstruction result of pixels that were successfully matched was consistent across lighting configurations and exposure times. As an example, Figure 11 shows a histogram of the differences in the depths calculated from the Baseline Sunlit image and an HDR image. From this observation, it can be seen that the best parameter for comparing the quality of a given stereo image pair is the percentage of pixels in the image that were successfully matched.

5.4. The Effect of Geometry on Percentage of Good Pixels

Figures 12 and 13 show the effect of lighting geometry on the percentage of Good Pixels found by the stereo algorithm. The configurations that spread light out the most over the scene had the best Good Pixel percentages.

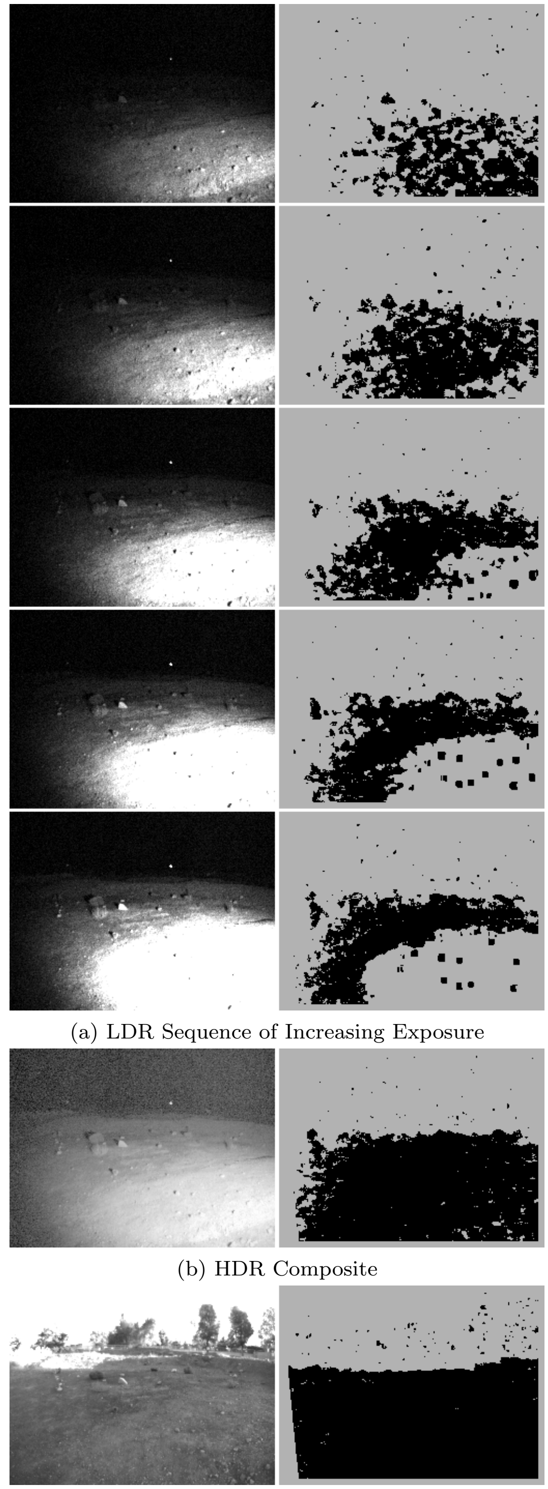

(a) LDR Sequence of Increasing Exposure

(b) HDR Composite

(c) LDR Baseline Sunlit (for reference)

Figure 7: Effect of exposure time and HDR processing on Good Pixels for the Floodlight configuration. Images in the left column are the left frame of the stereo pair, images in the right column are the locations of Good Pixels (black).

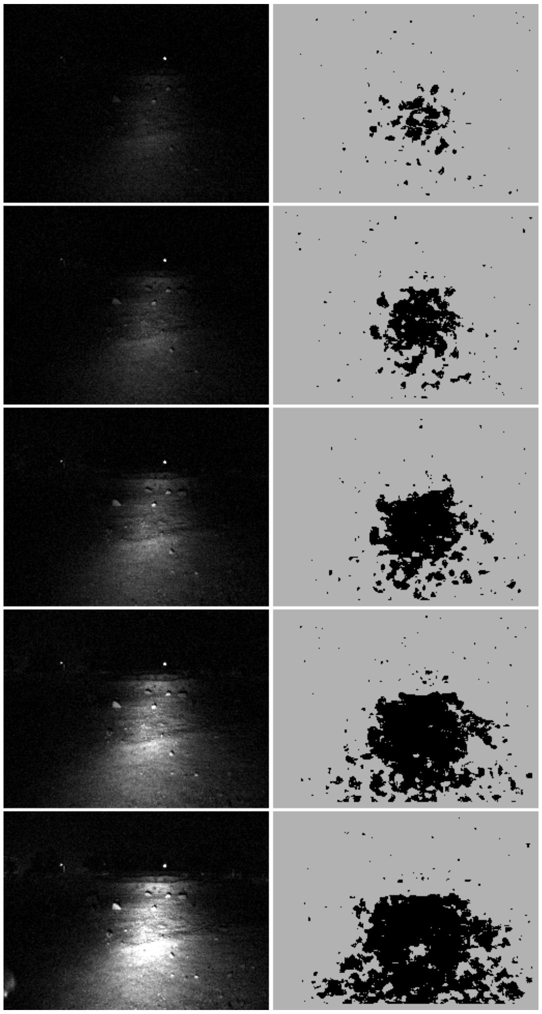

(a) LDR Sequence of Increasing Exposure

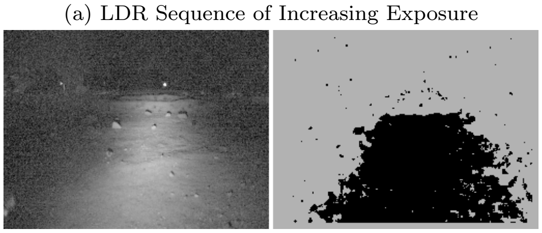

(b) HDR Composite

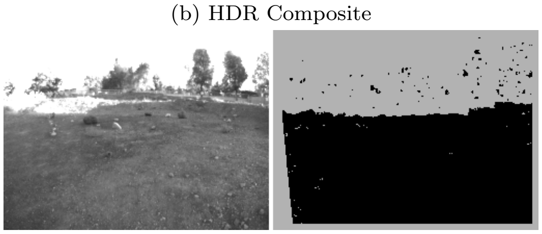

(c) LDR Baseline Sunlit (for reference)

Figure 8: Effect of exposure time and HDR processing on Good Pixels for the LED Spotlight 3 configuration. Images in the left column are the left frame of the stereo pair, images in the right column are the locations of Good Pixels (black).

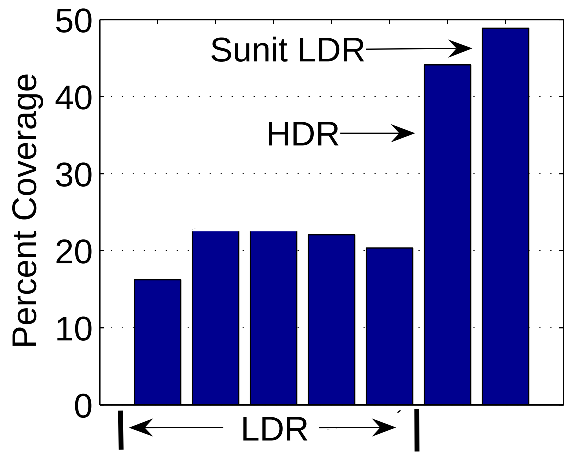

Figure 9: Effect of exposure time and HDR processing on Good Pixels for the Floodlight configuration

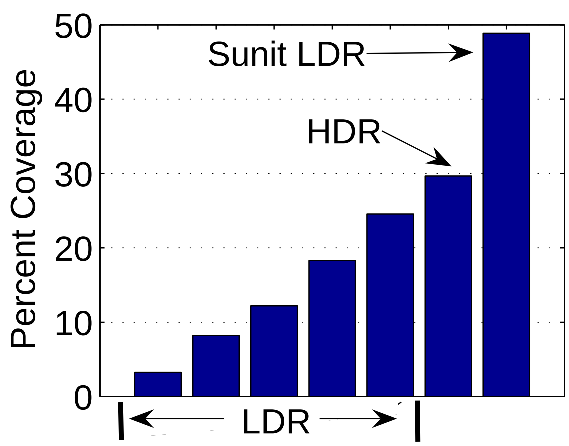

Figure 10: Effect of exposure time and HDR processing on Good Pixels for the LED Spotlight 3 configuration

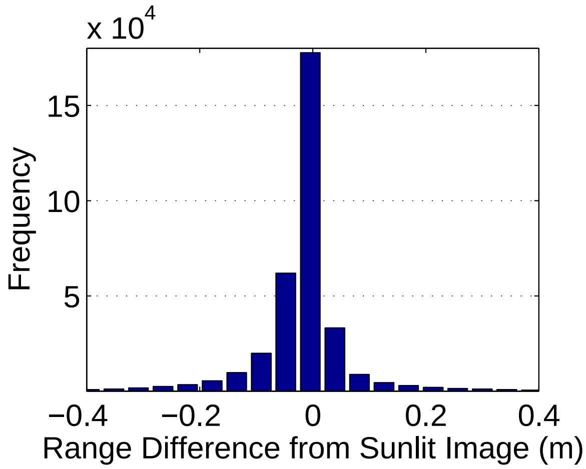

Figure 11: A histogram of the range differences between the range map produced by the HDR pair taken in the Floodlight configuration and LDR pair taken in the Baseline Sunlit configuration. ($\mu = -1.6\,\text{cm}$, $\sigma = 12\,\text{cm}$)

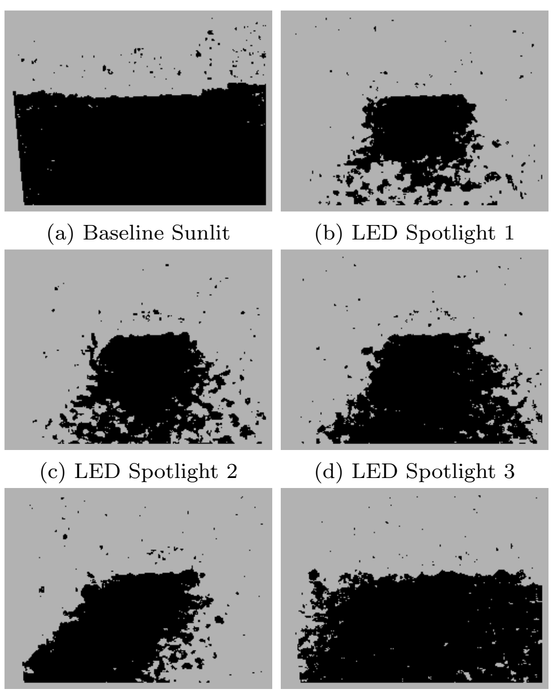

(a) Baseline Sunlit (b) LED Spotlight 1

(c) LED Spotlight 2 (d) LED Spotlight 3

(e) LED Spotlight 4 (f) Floodlight

Figure 12: Effect of geometry on Good Pixels (black)

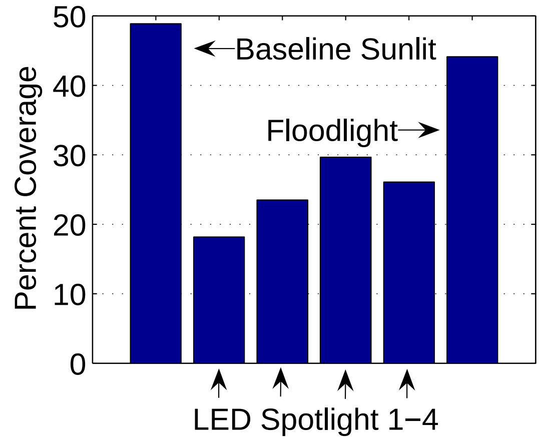

Figure 13: Effect of geometry on Good Pixels

6. Conclusion

HDR processing makes a large improvement on the percentage of pixels with sufficient surrounding texture for stereo correlation when using the diffuse tungsten-halogen floodlight. Using the LED spotlight, the HDR image results in a lesser but not insignificant improvement over the best LDR image. This is partly explained by the fact the highly collimated LED spotlight beam was aimed at a 10m distant point and that it does not emit light uniformly in all directions within it's field of view. HDR processing can be expected to offer a greater improments with a less collimated light source, such as would be needed to navigate a rover (90° FOV being a common requirement).

The range difference between the Good Pixels in the sunlit image and the Good Pixels in the night images was close to zero. This suggests that when there is sufficient image texture data for stereo correlation, ranging accuracy is as not affected by the quality of illumination. In other words, LED lit stereo ranging accuracy is as good (or bad) as sunlit stereo accuracy. Thus, our challenge in getting a good stereo range map is not so much the quality of our texture data, but the quantity of our coverage.

The geometry of the camera-light system has a discernible effect on stereo coverage, with the best results obtained by a light source below and laterally offset from the cameras. This ensures that small rocks and other protuberances cast shadows visible to the cameras.

Whilst our test environment does not truly replicate the properties of lunar regolith (which is not a lambertian reflector amongst other things), it supports the notion that high dynamic range stereo camera ranging with unstructured artificial light is a viable option for operating a rover inside a dark lunar crater.

Continuous locomotion would require specialized camera hardware with a logarithmic response to get HDR images in a single image, along with high *peak* power (1 kW) lighting. Relaxing the requirement for continuous locomotion allows longer exposure times (reducing power needs) and multiple images of the same scene (allowing HDR images from standard cameras).

Acknowledgment

The authors gratefully acknowledge the support of Chin San Han, Michael Broxton, Patrick Mihailich and NASA Ames's Vision Workbench development team. This project was supported by the NASA Exploration Technology Development Program (Human-Robotic Systems Project). Kyle Husmann was supported at NASA Ames by NASA's Robotics Academy Project.

References

1. Margot, J., Campbell, D., Jurgens, R., and Slade, M. Topography of the lunar poles from radar interferometry: a survey of cold trap locations. *Science*, 284:1658–1660, June 1999.

2. Allen, C. Shoemaker crater - going where we can "see". Lcross site selection, NASA Johnson Space Center, NASA Johnson Space Center, Mail Code KT, Houston, TX 77058, 2006.

3. Heiken, G. H., Vaniman, D. T., and French, B. M., editors. *Lunar sourcebook - a user's guide to the moon.* Cambridge University Press, 1991.

4. Hancher, M. D., Broxton, M. J., and Edwards, L. J. The vision workbook: A user's guide to the nasa vision workbench v1.0. Technical report, NASA Ames Research Center, Moffett Field, CA 94035-1000, November 2007.

5. McKay, D. S., Carter, J. L., Boles, W. W., Allen, C. C., and Allton, J. H. Jsc-1: A new lunar soil simulant. In *Proceedings of Space IV: Engineering, Construction and Operations in Space*, volume 2, pages 857–866, 1994.

6. http://damien.douxchamps.net/ieee1394/coriander/index.php.

7. Bouguet, J.-Y. http://www.vision.caltech.edu/bouguetj/calib_doc/index.html.

8. Edwards, L., Hancher, M. D., and Broxton, M. http://ti.arc.nasa.gov/visionworkbench/.