

Image-Directed Sampling for Geometric Modeling of Lunar Terrain

Uland Wong¹, Ben Garney, Warren Whittaker, and Red Whittaker

Abstract Geometric modeling from range scanners can be vastly improved by sampling the scene with a Nyquist criterion. This work presents a method to estimate frequency content *a priori* from intensity imagery using wavelet analysis and to utilize these estimates in efficient single-view sampling. The key idea is that under certain constrained and estimable image formation conditions, images are a strong predictor of surface frequency. This approach is explored in the context of lunar application to enhance robotic modeling. Experimentation on simulated data and in artificial lunar terrain at aerial and ground rover scales is documented. Results show up to 40% improvement in MSE reconstruction error. Lastly, a class of image-directed range sensors is described and a hardware implementation of this paradigm on a structured light scanner is demonstrated.

1 Introduction

Despite proliferation of motion-coupled industrial scanners in field robotics, there remain important applications which require actuated sensors and intentional sampling. These applications must consider the question of: "Where to sample the data?" Planetary exploration is perhaps the best example of the need to plan geometric sampling. Correct photometric and geometric classification of rocks and craters could enhance automated sample collection and manipulation for drilling on science missions. Steered and optically reconfigurable flash LIDAR is of great interest in automated planetary landing for its capability in handling a spectrum of ranges and scales. Lastly, cooperative orbital and ground mapping, where the robot itself is a bore-sighted, Dirac sampler, will likewise benefit from a saliency-based approach to resolving terrain obstacles first detected in aerial sensing [1]. In these applications the importance of intelligent sampling is twofold: (1) the relative expense of range scanning places a premium on sampling well and (2) the requirements of model quality dictate the application.

Uland Wong, Warren Whittaker, Red Whittaker
Robotics Institute, Carnegie Mellon University
email: {uyw, warrenw, red} @ andrew.cmu.edu

Ben Garney
The Engine Company, LLC
email: ben.garney@gmail.com

The quality of 3D reconstruction from point measurements is a function of the density, distribution and order of sampling. Only grid-based or uniform angular sampling strategies – which are commonly used - inform *a priori* the total number of readings required to scan a scene from a single viewpoint. In spite of this, these are poor sampling strategies for reconstruction. Triangulation of the resultant point clouds produces glaring artifacts like polygonal slivers and incorrect connectivity. Aliasing of grid-based approaches creates high gradient planes out of depth discontinuities. Wasteful and duplicate measurement increases sensing time for little information gain. These effects are particularly problematic in applications where the 3D detail is used in object recognition or understanding. One of the compelling motivations of this work stems from the observation that sensors which are not limited by grid sampling are the same (low-throughput) sensors which would most benefit from sampling in a principled manner.

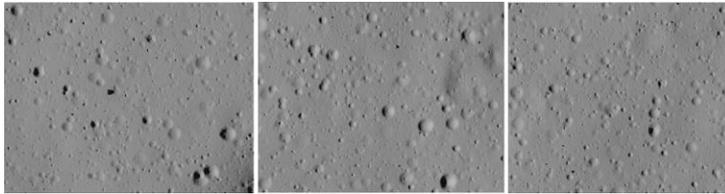


Fig. 1. Image features in diffuse environments, such as the moon, correlate strongly with geometric features. These three simulated lunar terrains are utilized in this work.

Cameras are natural complements to range sensors [2]; the authors advocate that it is possible to utilize color imagery to *direct* range sensing of a scene to avoid these problems. At the core of this idea is the observation that under certain image formation assumptions, the frequency content of color images correlates strongly with that of geometry images. By sampling the scene with Nyquist-informed density distribution - and by utilizing camera information to plan such a sampling - vast improvements can be made in both the quality and efficiency of range sensing.

This paper identifies frequency features for use with the image-directed scanning approach, compares feature performance based on reconstruction metrics and demonstrates new sensor designs and visualization techniques which utilize this paradigm. The advantages of diffuse planetary environments are discussed in tailoring domain-specific salient image features. Lastly, the aforementioned planetary applications are used to test the efficacy of the method.

2 Related Work

Image-direct sampling can be seen as a complementary approach to that detailed in the prior camera/LIDAR range fusion work of the authors [2]. Specific motiva-

tion for this work arises from the observation that the quality of geometric fusion depends as strongly on the quality of the underlying range data as it does on the image-based shape estimation. By affecting the nature of range samples, image-directed sampling can be utilized in tandem with a spectrum of traditional fusion techniques [3] to significantly enhance range models.

Research towards determining and generating optimal point samples is well documented. Work of particular note includes [4], which devised a method to reduce the post-processing time and increase the convergence probability of multi-scan alignment by storing only high-saliency points from scans. Related multi-view extensions of the sampling problem include the body of next-best-view work [5], [6]. Frequency content is also distinguished as a strong predictor of novelty in exploratory map building in [7], specifically for planetary terrains.

Wavelets and other multi-scale frequency features are widely used in data understanding and reconstruction. They include approaches to LIDAR [8] and intensity images [9], [10]. In particular image based approaches have been successfully used in interest seeking for planetary science [11]. Range and intensity fusion with wavelets has also been studied for forest region mapping [12].

While this research draws inspiration from prior work, it distinguishes itself in two important aspects. Prior approaches have enhanced unstructured range data using image content in a post-processing manner, this work tightly couples image analysis in the range *acquisition* phase. Secondly, this work is motivated by analysis of surface reflectance in environmental appearance. The authors believe that proper domain knowledge can constrain image features to physical phenomena.

3 Frequency-Based Modeling

Frequency is one of the most effective predictors of information content in a signal. Continuous time or spatially-varying signals can be decomposed into their frequency components by utilizing a number of different transforms to the frequency domain. While, there is no exact definition of this domain, transforms generally have desirable properties that indicate the change-over-time of the signal.

The frequency components of range models cannot be directly measured, as the only sensors that exist for range modeling are Dirac samplers. ToF and Triangulation based methods all produce point estimates; such discrete samples are subject to aliasing. It is therefore prudent to consider the Nyquist sampling theorem for spatially varying samples [13] when reconstructing a range image. This states that perfect reconstruction is possible from severely sub-Nyquist spatial samples, if the signal has *known* compact support and shift invariance in the frequency domain.

These conditions are impossible to satisfy in practice due to sensor noise and discretization. Thus, only effective use of samples can be made by apportioning low and high information areas with some sample weight distribution and relative regard for Nyquist. This corresponds to oft-opposing objectives for modeling in a

frequency sense: (1) accuracy of volumetric reconstruction and (2) localization of high spatial resolution detail to important areas. The approach taken here is that a range of sample importance priors should be explored for application-specific reasons, rather than attempting theoretical optimality.

3.1 An Image-Directed Approach to Sampling

A prior over the Nyquist frequency distribution is sought, which can guide spatially-varying sample acquisition. Two dimensional range images are common manifold representations of 3D models. The wavelet transform [14] is a simple and elegant representation of the spatial frequency content that can be used for such range images. Sum-squared detail coefficients at each pixel, the wavelet energy, are one possible indicator of the information distribution. Sample priors can be estimated using this metric on data from a complementary camera, which can instantaneously capture high-density intensity with the same perspective as a range image at a low constant time cost.

However, intensity is a generally inseparable composition of three-dimensional geometric, material and illumination properties projected onto a two-dimensional plane [15]. This work considers a subset of intensity images, whether raw or feature-transformed, which correlate significantly with scene geometry. Features from these images are viewed as noisy approximations of high resolution range or geometry cues.

Consider a simplified form of the rendering equation [16], which describes light transport for geometric optics at each scene point:

$$L_o = L_e + \int_{\Omega} f \times L_i(w) \times (w \cdot \hat{n}) dw \quad (3.1)$$

where L_o, L_e, L_i are the light intensities outgoing, emitted and incoming respectively; f is the BRDF; w is the incoming light vector and \hat{n} is the surface normal for calculating incident irradiance. This equation can be reduced to the manageable form of (3.5) using common assumptions of planetary environments, which are *barren, dry, rocky*, and *simple illumination* (point or direct solar) [17]. *Barren* describes geometric smoothness, whereby surface normals can be approximated by geometric samples. *Dry* indicates a lack of participatory media, such as atmosphere or dust, which significantly scatters light. *Rocky* means scene reflectances are Lambertian and albedo change is minimal. Lastly, *simple illumination* assumes an additive set of discrete point sources, which in most cases reduces to a single sun source approximated by a point.

$$E = L_o = L_e + \sum_k f L_{i,k}(w_k) (w_k \cdot \hat{n}) \quad (3.2)$$

$$E = \sum_k \frac{\rho}{\pi} L_{i,k}(w_k)(w_{i,k} \cdot \hat{n}) \quad (3.3)$$

$$E = \sum_k \frac{\rho}{\pi} I_k(w_{i,k} \cdot \hat{n}) \quad (3.4)$$

$$E \propto (w_{i,k} \cdot \nabla x) \quad (3.5)$$

Simplifications of the rendering equation result from applying *simple illumination* (3.2), *rocky* (3.3), *dry* (3.4), and *barren* (3.5) assumptions. Eq. (3.5) relates image intensity (E) to a proportion of the dot product between the source and the gradient of the geometry (i.e. range image). The frequency relationship between a signal and its gradients is given by Frankot, which enables heuristic, but strongly correlated frequency sample planning for artifact reduction [18]. Note that the planetary assumptions above do not allow for cast shadowing, so only camera-centered illumination is strictly valid for real data.

The reconstruction problem presented here is viewed as selecting $x_k = \langle s_1, \dots, s_k \rangle$ from all possible samples, $s_i \in X$, such that the reconstruction error is minimized: $e = \arg \min_s \|t(x_k) - X\|$ under some interpolating function t . The objective is to determine features from intensity images (i) and learn a mapping such that $f(i, X) \rightarrow \hat{x}_k$, where $\hat{x}_k \sim x_k$.

4 Modeling for Planetary Robots

Validation of the image-directed sampling approach is conducted in the specific context of lunar robotics. Private enterprises, such as the Google Lunar X-Prize, have renewed interest in automation of landing and exploration on the moon. There is particular emphasis on creating high quality maps and models of the moon using robots as precursors to humans. Experiments were conducted with simulated aerial and ground-robot data to demonstrate scale robustness.

4.1 Lunar Aerial Mapping

Aerial-scale lunar terrain was procedurally generated in Blender™, utilizing randomized fractal crater and rock distributions taken from the Surveyor 6 mission [19]. The data includes independent overhead depth and RGB values for each voxel in the scene at 5vox/m. Three different scenes were generated and are shown in Fig. 1.

The datasets simulate lunar sensing during terminal descent, at a scale 50m above the ground and were generated for related research. An automated lander,

for example, might identify geometric hazards (rocks and craters) and divert to a suitable landing spot. Image-feature based methods have been proposed to identify these hazards, including detection of shadows [20]. However, in this scenario a gimbaled altimeter or boresight flash LIDAR can also work in tandem with image-based methods to acquire a full 3D model.

Simulation of the sampling process includes selecting samples (voxels) from the highest resolution depth map, triangulating the subsamples and linearly interpolating to create a depth map at the native resolution. This depthmap is then compared against the ground truth digital elevation map to produce reconstruction error scores. Artifacts such as sensor noise are not considered in this analysis. Several scanning strategies were compared on this dataset. They include:

Uniform Grid – An $N \times M = k$ element uniform grid sampling across the scene, rounded to voxel edges.

Uniform Random – Random sampling of k elements in the scene with each voxel receiving uniform weight.

Gradient Weighted – Random sampling weighted by the 2-norm of the partial image derivatives, a precursor to image-based edge detection. Calculated using the matlab command `gradient`:

$$w = \|\nabla_x I + \nabla_y I\|_2 \quad (4.1)$$

where $\nabla_x I$ and $\nabla_y I$ are the image partial derivatives in the x and y direction respectively.

Wavelet Weighted – Random sampling weighted by the sum-square of the wavelet coefficients, normalized by subregion size. The wavelet transform is generated with DB2 wavelet using the command `wavedec`.

$$w = \sum_{j=1}^N \frac{1}{2^j} (D_j)^2 \quad (4.2)$$

where j is the pyramid level and D_j are the detail coefficients at level j .

Entropy Weighted – Random sampling weighted entropy in a 9x9 image neighborhood centered about the query pixel, calculated with the matlab command `entropyfilt`. This transform is frequently used as a texture cue.

Three separate metrics were used to compare the reconstruction error between sampled depth maps and ground truth. These metrics each emphasize a different requirement of the model.

Mean Squared Reconstruction Error (MSE) – Penalizes for any large difference between the reconstruction and the ground truth.

Gradient Weighted Reconstruction Error (GRAD) – Penalizes for incorrect reconstruction at depth discontinuities.

Saliency Weighted (SAL) – Assigns higher weight to in reconstruction salient features such as rocks and craters correctly. The saliency map is generated using using the Saliency Toolbox [21].

Random samples were generated 20 times using each strategy and the resultant scores were averaged. In the case of the grid strategy, the grid was shifted horizontally and vertically so that sampling did not always occur in the same voxels or on the edge of the scene. Optimal linear coefficients were found using `fminsearch`. The experiments were performed with fractional sampling density of the ground truth resolution in log spaced increments, i.e. $k = \frac{1}{2}, \frac{1}{4}, \frac{1}{8}, \dots, \frac{1}{2048}$.

Table 1. Performance of Sampling Strategies on Simulated Aerial Terrain

Scene #1	Random	Grid	Gradient	Entropy	Wavelet
MSE	1	1.02	0.88	0.87	1.11
GRAD	1	1.1	1.10	1.13	1.12
SAL	1	0.93	0.93	0.88	1.07
Scene #2	Random	Grid	Gradient	Entropy	Wavelet
MSE	1	1.01	0.87	0.85	0.99
GRAD	1	1.05	1.12	1.16	1.12
SAL	1	0.82	0.96	0.95	1.05
Scene #3	Random	Grid	Gradient	Entropy	Wavelet
MSE	1	0.91	0.83	0.80	1.21
GRAD	1	0.98	1.02	1.06	1.01
SAL	1	1.02	0.82	0.79	1.31

The results of the sampling simulation are detailed in Table 1. The scores are given as the mean ratios of the error between random sampling and the strategy in question, weighted by inverse sample density.

$$score = \frac{1}{\sum k} \sum_{k \in S} \frac{T}{k} \cdot \frac{e_{rand|k}}{e_{strategy|k}} \quad (4.3)$$

Thus, scores represent an improvement multiplier over random sampling. Re-weighting by inverse density accounts for performance across the entire logarithmic sampling densities of interest while discounting minute reconstruction noise due to lack of selection replacement at the highest density trials.

It is noted that in sparse samplings ($k < 15000$), wavelets performed the best across all metrics (graphs of all curves are omitted for succinctness, see Fig.3). This is particularly promising when high reductions in the amount of geometry acquired are necessary. Wavelets perform more modestly (20% improvement) given higher numbers of points, as sample selection becomes constrained in the

downsample simulation. Gradient and entropy approaches both perform relatively poorly. Gradient weighting fails because the sharpest edges occur in shadowed craters, not depth discontinuities. A similar problem occurs with entropy, which utilized a fixed scale and neighborhood. While the wavelet transform is also susceptible to overweighing shadow features, the multi-scale capability also captures the crater and rock features. The grid approach comes out ahead of other methods in deficiently sparse reductions. While grid sampling is subject to bias and aliasing, there is the possibility for large holes in random sampling. A better pseudo-random sampling approach for all these features might first compute a Delaunay triangulation within similarly valued regions to ensure adequate density.

4.2 Lunar Ground Mapping

Unlike aerial sensing, which is difficult to scale for lab experimentation, high-fidelity robot scale terrain can be constructed with readily available materials. This section describes experimental verification of the approach on lunar-like terrain constructed in a 2m x 1.3m moonyard. Macro-scale rocks were placed on an undulating bed of garden lime with size and position determined by lunar rock distribution data. The moonyard was then dusted with CMU1, a nonhazardous, optical regolith simulant developed at Carnegie Mellon and consisting primarily of sieved coal dust and powered lime. Micro craters were created by ballistic placement of pebbles and further dusting. Fig. 2 shows the ideal distribution of features generated in simulation and the actual placement of these features as-built.

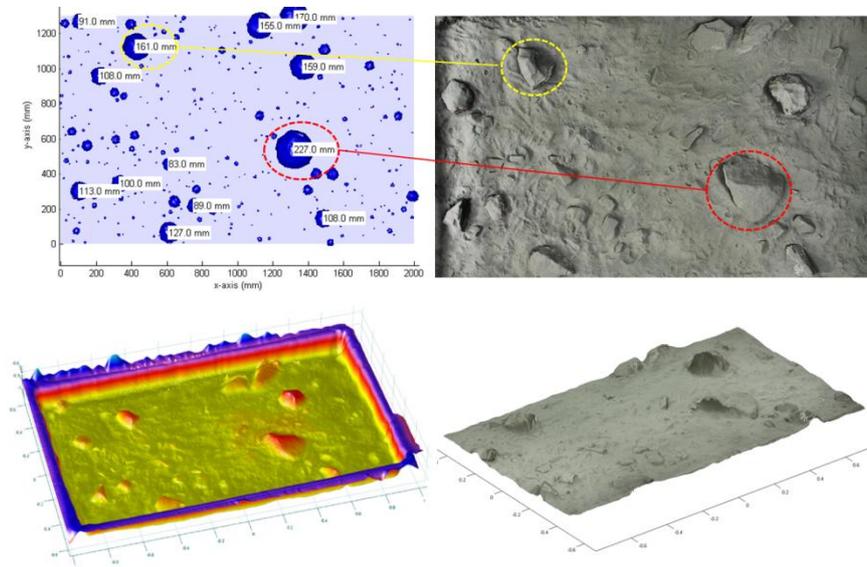


Fig. 2. A simulated lunar rock distribution utilizing Surveyor data (top left) and realization in the construction of an artificial moon scene (top right). The bottom row illustrates a z-colored mesh model of the scene generated with LIDAR scanning (bottom left) and color-mapping of the terrain from DSLR images (bottom right).

The terrain was placed in a light simulator frame with calibrated sensor and sun-scale light source positions for data collection (an overhead sun position was used). Range data was collected with a survey-grade phase-shift LIDAR by scanning from multiple positions to ensure adequate resolution and elimination of range shadows. Likewise, HDR imagery was collected from multiple locations utilizing DSLRs. The raw LIDAR point cloud, consisting of 5million points, was colored with RGB imagery, cropped to the edge of the sandbox and averaged into a 971x1674 voxel (1.6M samples) orthographic digital elevation map. Fig. 5 in the next section shows the data acquisition setup.

As the scanner does not have independent sampling control, it was necessary to scan at maximum resolution, voxelize into a uniform grid and sub-select samples post measurement for testing. Quantization into lower resolution voxels enables independent readings for each position without interpolation and provides resilience against the natural angular biases of the sensor.

Table 2. Performance of Sampling Strategies on Lunar Terrain

	Grid	Random	Gradient	Entropy	Wavelet
MSE	1.1	1	1.09	1.15	1.33
GRAD	0.82	1	1.21	1.32	1.42
SAL	0.96	1	1.2	1.30	1.50

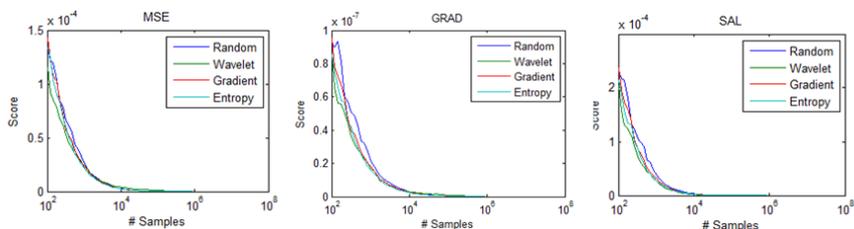


Fig. 3. Comparison of selected strategies over all subsample amounts and metrics on lunar terrain model. Visually, wavelets are the best performers (green).

The sampling strategies described in the previous section were tested and the same metrics described are used in this analysis without change. The results of the experiment in artificial lunar terrain are summarized in table 2. The wavelet-based sampling is clearly the best performer in this experiment (~40% improvement). Curiously, all of the strategies performed significantly better than in the simulated aerial imagery. It is believed that this effect is due to the strong cast shadows in the simulated imagery, while this experiment (though using real data) featured so f-

ter shadows and few large, negative features. While results are promising, it must be noted that this single dataset cannot be representative of the entire spectrum of lunar appearance and geometry and additional testing is required.

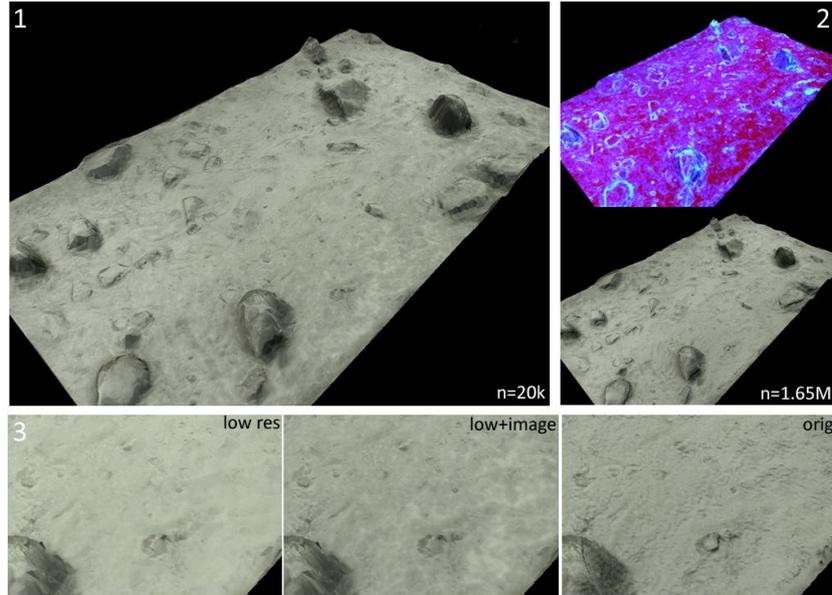


Fig. 4. Lunar terrain model generated with image-directed sampling. (1) An extremely sparse sampling rendered with roughness-modulated Oren-Nayar BRDF retains useful geometry and is visually appealing. (2) The frequency map of the terrain generated using wavelet decomposition. (3) Detail of roughness visualization demonstrates that a data deficient model can retain much of the optical appearance of the highest resolution model.

Visual results of terrain reconstruction using wavelet-weighted, image-directed samples are illustrated in Fig. 4. Fig. 4.1 shows a rendering of the terrain with 20 thousand sample points, representing an 80x density reduction from the ground truth model of 1.6 million points (Fig. 4.2, bottom). As a result of well-placed samples, the macroscopic geometry is surprisingly well-formed for such a sparse mesh; there are few glaring differences when compared to the ground truth. The downsampled model is rendered with an inferred roughness from the wavelet decomposition. High frequency features estimated in the image are "painted" onto the sparse geometry by modulating the roughness term in an Oren-Nayar vertex shader proportional to the energy (Oren-Nayar is close to the true BRDF of lunar soil) [22]. A threshold can be applied to ensure major geometric features remain metrically true. This approach to visualization conveys the high resolution surface characteristics to the viewer while preserving mesh compactness with minimal computation. Fig. 4.3 further illustrates this technique. With roughness inference turned off, the surface is an aliased collection of triangles, but rendering with roughness infuses many of the characteristics of the true surface.

5 A Class of Image Directed Scanners

In the prior sections, a discrete grid based approach is compared to random sampling approaches. Significant improvement was shown when random sampling is combined with image frequency estimates. However, it can be argued that this comparison is unfair. Few existing range sensors can acquire random, independent samples. Most of the applications that are considered in this work would use steered beam sensors if taken at the current state of sensor art. The authors believe that the approach is still beneficial when considering motion-constrained devices with trajectory planning and the possibility of utilizing region-based iterative enhancement in conjunction.

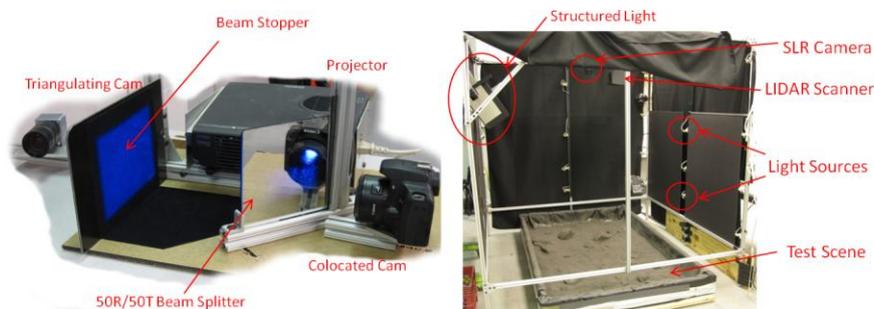


Fig. 5. An image-directed scanner consisting of a high resolution camera optically co-located with a DLP projector (left). Experimental setup utilized in this paper showing mounting positions of all the sensors (right).

However, it is of interest to consider the class of sensors that provide pixel-independent sampling, as they are capable of implementing the approach in the purest form. Perhaps the most obvious and common of these are structured light sensors based on digital projectors. Though, foundational work has been conducted in camera-integrated smart laser scanners [23].

A structured light scanner was utilized in developing a proof-of-concept hardware implementation of image-directed scanning. Fig. 5 illustrates the optical configuration of this device. A high resolution DSLR camera is optically co-located with the projector through the use of a half-silvered mirror. The system is calibrated such that each pixel in the camera corresponds to an exact outgoing ray of projected light. This configuration enables the system to sample the exact location viewed by the camera without shadowing. Finally, a second camera is optically offset from the other two devices and measures depth via triangulation. This setup is closely related to that presented in [24], which is a co-located system for augmented reality. Image-directed sampling is implemented on this scanner as follows. The projector first lights all pixels so that the scene can be imaged by the DSLR and analyzed. Afterwards, scanning proceeds in either of the two modes.

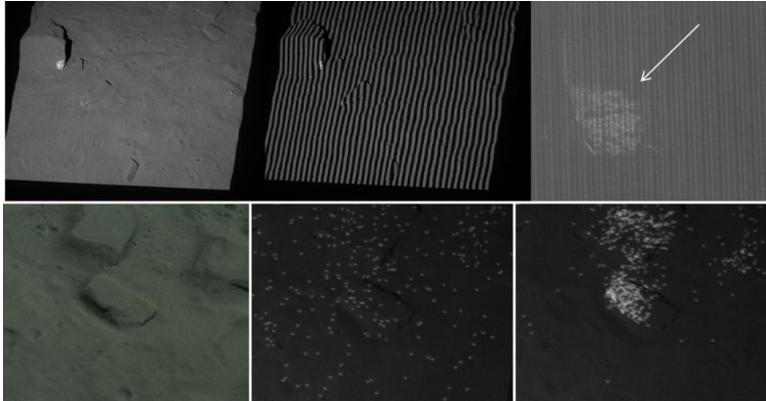


Fig. 6. Linear gray code scanning of the lunar terrain causes errors at the finest resolution due to differences in underlying scene albedos (top). Image cues can direct the scanner to resample salient regions such as rocks a pixel at a time (bottom). Samples are shown simultaneously illuminated for illustrative purposes; scanning occurs one dot at a time.

The first mode is *ambiguity enhancement*. In temporally-coded stripe scanners, the highest resolution stripes often result in ambiguities or errors (illustrated in Fig 6). Among the sources for these errors include the albedos and complex reflectance of materials present in the scene to the resolving power of the camera lens. Gray codes can alleviate this problem, but result in lower effective resolution when such errors occur [25]. Image-directed sampling can be applied to identify regions which require dense, accurate samples, which can then be resampled one pixel at a time (reducing errors). The bottom row of Fig. 6 illustrates salient features (rocks) that require more samples. The locations of these samples are all illuminated by the projector for graphical clarity, but are scanned individually. Results on the moon terrain showed 25% improvement in reconstruction MSE by rescanning $\sim 10\%$ of gray code pixels using a single dot. Though the time cost of this extra scanning as implemented is large, significant reductions are possible with parallelism.

The second mode of operation explored is *data reduction*. Temporal coding in structured light scanning often precludes motion while scanning is in progress. Spatial codes, which utilize color or pseudorandom dot distributions, enable all depth information to be captured in a single frame. However, these techniques do not lend themselves to tightly-packed high resolution samples [25]. Image frequency can be utilized to plan a sampling of the scene with optimally placed dot locations. Reconstruction quality is thus improved while maintaining sparse samples. While not implemented at the time of this writing, using clique coloring in only high density regions and simple spatial pattern distortion estimation in low density regions can greatly reduce the number of colors to be discriminated.

Micro mirror arrays, like that used in the projector unit, are much higher resolution than beam sensors and even flash LIDAR. More importantly they require no actuation. The authors envision future intelligent sensors coupling low-resolution,

low-throughput range receivers, micromirrors and co-located color cameras to determine optimal samples in the same vein our sensor.

6 Conclusion

An approach to image-directed range scanning utilizing wavelet analysis was described in this paper. Lunar application was targeted, where the particulars of surface appearance constrain image formation and correlates intensity with geometry. Validation of the approach was conducted for lunar aerial and ground mapping through simulation and terrestrial recreation respectively. Preliminary results are promising: reconstruction improvements of 40% can be achieved over grid-based scanning techniques from a single viewpoint. Simple, but effective visualization for sparse meshes was demonstrated for combining image-estimated texture with the geometric portion of the lunar surface BRDF. Lastly, a hardware implementation on a pixel-independent structured light sensor was explored.

In the immediate future, we hope to analyze the effects of sensor noise on the approach. Noise in both the scanner and the camera can contribute to inaccurate prediction and wasted samples. Though this problem was mitigated in these experiments by limiting the maximum regional density that could be sampled, the sensors were low noise to begin with. The ability to handle a larger number of different albedos could also enhance the robustness of the method. Possibilities to crosscheck albedo changes (which contribute to frequency content in images) and shape changes include the use of solar-motion and probabilistic albedo estimation techniques. Similarly, while it was shown that this approach is defeated with strong shadows, the shadows can be mitigated with active illumination.

Acknowledgments The authors acknowledge Kevin Peterson, Heather Jones and Jason Koenig for use of lunar model data.

References

1. Jones, H. Wong U., Peterson, K., Koenig, et al. Complementary Flyover and Rover Sensing for Superior Modeling of Planetary Features. Submitted for publication, 2012.
2. Wong, U. Garney, B. Whittaker, W., Whittaker, R. Camera and LIDAR Fusion for Mapping of Actively Illuminated Subterranean Voids. In Proc. Field and Service Robotics, 2009.
3. Diebel, J., Thrun, S. An Application of Markov Random Fields to Range Sensing. Neural Information Processing Systems, 2005.
4. Martinez, J., Reina, A., Mandow, A. Spherical Laser Point Sampling with Application to 3D Scene Genetic Registration. In Proc. ICRA, 2007.

5. Shahid, K., Galina, O. Intelligent LIDAR scanning region selection for satellite pose estimation. *Computer Vision and Image Understanding*, Vol. 107, Issue 3, Sept. 2007. Pp 203-209.
6. Mark, L. H., Okouneva, G., Saint-Cyr, P., Ignakov, D., English, C. Near Optimal Selection of Views and Surface Regions for ICP Pose Estimation. In *Adv. in Visual Computing*, 2010.
7. Thompson, D., Smith, T., Wettergreen, D. Information-Optimal Selective Data Return for Autonomous Science and Survey. In *Proc. ICRA 2008*.
8. Wei, H., Bartels, M. Unsupervised Segmentation Using Gabor Wavelets and Statistical Features in LIDAR Data Analysis. In *Proc. Pattern Recognition (ICPR)*, 2006.
9. Loupias, E., Sebe, N., Bres, S., Jolion, J. Wavelet-based Salient Points for Image Retrieval. In *Proc. 4th International Conference on Advances in Visual Information Systems*, 2000.
10. Achanta, R., Hemami, S., Estrada, F., Susstrunk, S. Frequency-tuned Salient Region Detection. In *Proc. CVPR*, 2009.
11. Dunlop, H. Thompson, D., Wettergreen, D. Multi-scale Features for Detection and Segmentation of Rocks in Mars Images. In *Proc. CVPR*, 2007.
12. Wang, Z., Boesch, R., Ginzler, C., Color and lidar data fusion: application to automatic forest boundary delineation in aerial images. *Int. Arch. of the Photogrammetry Remote Sensing and Spatial Information Sciences*, 2007.
13. Aldroubi, A., Grochenig, K. Nonuniform Sampling and Reconstruction in Shift-Invariant Spaces. *SIAM Review*: 43(4), 2001.
14. Daubechies, I. *Ten Lectures on Wavelets*. Soc. for Industrial and Applied Mathematics, 1992.
15. Horn, B. K. P. *Robot Vision*. The MIT Press, 1986.
16. Kajiya, J.. The rendering equation. In *Proc. SIGGRAPH*, 1986.
17. Wong, U. *Lumehancement: Exploiting Appearance for Planetary Modeling*. PhD Dissertation, Carnegie Mellon University, 2012.
18. Frankot, R., Chellappa, R. A Method for Enforcing Integrability in Shape from Shading Algorithms. *IEEE PAMI*: 10(4), 1988.
19. Heiken, G., Vanniman, D., French, B. *The Lunar Sourcebook*. 1991.
20. Hata, S., Sumioka, K. Detection of obstacles on the moon. *IEEE Symposium on Industrial Electronics*, 2004.
21. Walther, D., Koch, C. Modeling attention to salient proto-objects. *Neural Networks*, 2006.
22. Oren, M., Nayar, S. K. Generalization of Lambert's Reflectance Model. *SIGGRAPH*, 1994.
23. Cassinelli, A., Zerroug, A., Ishikawa, M. Camera-less Smart Laser Projector. *ACM SIGGRAPH 2010 Emerging Technologies*.
24. Fujii, K., Grossberg, M., Shree, N. A Projector-Camera System with Real-Time Photometric Adaptation for Dynamic Environments. In *Proc. CVPR*, 2009.
25. Salvi, J., Pages, J., Battle, J. Pattern Codification Strategies in Structured Light Systems. In *Pattern Recognition*, 2004. Pp 827-849.