Multi-view shape-from-shading for planetary images with challenging illumination.

Oleg Alexandrov[1] and Ross A. Beyer[1,2], [1]NASA Ames Research Center, Mail Stop 269-3, Moffett Field, CA, [2]Carl Sagan Center at the SETI Institute, Mountain View, CA (oleg.alexandrov@nasa.gov, ross.a.beyer@nasa.gov)

Introduction. Shape-from-shading (SfS), also known as photoclinometry, is a set of techniques for recovering surface relief based on variations in light intensity recorded in images [1, 2, 3]. We are interested in applying SfS to create high quality terrains from satellite images starting with an initial guess terrain from either stereogrammetry or LIDAR (for example, LOLA or MOLA).

We consider realistic camera models, such as pinhole and pushbroom, rather than just an orthographic projection used in many publications, multiple input images, each with its own exposure value, and terrain-dependent albedo. We model and mitigate the fact that cameras can have significant errors in position and orientation. Our approach allows the use of a wide variety of reflectance models, but we found that the Lunar-Lambertian [4] model performed well for our Lunar Reconnaissance Orbiter Camera images. We also model shadows via a shadow threshold, as well as regions occluded from the Sun. Our algorithm was shown to work for very large areas and tens of images acquired at high latitudes (about 87°) and hence low illumination angles, with little or no overlap among themselves, and copious amount of shadows.

Modeling. We model SfS as a minimization problem. We assume that we have one or more image views of a terrain, and for each image we know the Sun position and the position and orientation of the camera. The cost function to be minimized is

$$\int\!\!\int \sum_k \left[I_k(\phi)(x,y) - T_k A(x,y) R_k(\phi)(x,y)\right]^2 + \mu \left\|\nabla^2 \phi(x,y)\right\|^2 + \lambda \left[\phi(x,y) - \phi_0(x,y)\right]^2 \, dx\, dy.$$

It is made of three terms, with weights $\mu > 0$ and $\lambda > 0$ attached to the last two. They represent, respectively, the brightness, smoothness, and initial DEM deviation constraints. Here, ϕ is the terrain height above the datum, ϕ_0 is the initial terrain guess, $I_k(\phi)$ is the k-th camera image interpolated at pixels obtained by projecting into the camera points from the terrain, T_k is the k-th image exposure, A is the terrain-dependent albedo, $R_k(\phi)$ is the reflectance computed from the terrain for k-th image, and $\left\|\nabla^2\phi\right\|^2$ is the sum of squares of all second-order partial derivatives of ϕ.

The variables of optimization are the terrain and albedo at each point (x, y), and, for each image, its exposure, camera position, and camera orientation. We start with initial guesses for these parameters, and they can be independently kept fixed or varied.

Initialization. We used stereo or the gridded LOLA dataset to initialize the algorithm. The initial albedo $A(x, y)$ is set to 1. The values of the weights μ and λ are on the order of 0.06 and 10^{-3}. The initial exposures can be derived as the values that minimize the above equation.

Handling errors in camera positions/orientations. Such errors were causing the algorithm to not converge. They were mitigated using a coarse-to-fine terrain scale multi-resolution approach.

Dealing with shadows. Polar regions on the Moon can have very deep shadows. We model these by marking image pixels as valid or invalid based on a per-image shadow threshold. We adjust the cost function at the shadow boundary to avoid slope discontinuities in the computed terrain. If a ground point cannot be seen from the Sun because of being occluded by other higher terrain, we set its reflectance to 0. We do not model, as yet, points occluded from the camera.

Numerical solver. We use Google's Ceres solver to minimize the cost function.

Validation. We tested our algorithm on Lunar images, both at low and high latitudes, with one or more images, and with and without modeling shadows. We used images from the LRO Narrow Angle Camera. We validated our results in two ways: (a) run SfS with images artificially made coarser by resampling by a factor of 10, and validate with a terrain obtained from stereo using the original images, and (b) run SfS on full resolution (1 m/pixel) images, and validate with the sparse LOLA dataset. In both situations we show that SfS increases the terrain accuracy.

SfS on very large terrains. Our algorithm scales to using tens of images of a scene with dimensions of over $10,000 \times 10,000$ pixels, with the images barely overlapping in places and with vast areas in shadows. Handling this required many algorithm improvements. Particularly, we solve for camera positions and orientations by using not the entire region, which is infeasible, but a set of representative cropped areas. Later the solved

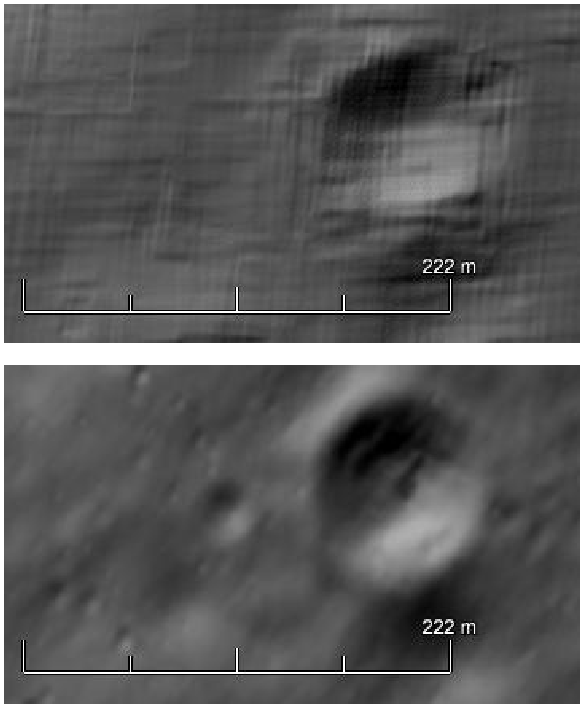

Figure 1: A Lunar terrain at 1 m/pixel before SfS (top) and after SfS (bottom).

camera parameters were kept fixed, and we computed the terrain by decomposing it into overlapping tiles and then merging the results.

SfS on other planets. We tried SfS on Mars with HiRISE images. A challenge with these images is the fact that the solar incidence angle is relatively constant in all of them, hence the illumination conditions do not vary much, making the SfS problem rather undetermined. We did not model atmospheric scattering. Mars in addition has a large variety of minerals on its surface, and hence more diverse (and complex) reflectance properties. We also experimented with Charon, the moon of Pluto, using images from the New Horizons flyby mission. We obtained plausible results, but did not have enough data to validate them. Overall, our solution needs more tuning before it can work for other bodies as well as it does for the Moon.

Software. Our SfS implementation is released, with detailed documentation and usage examples, as part of the NASA Ames Stereo Pipeline, a collection of open-source software for stereogrammetry and geodesy.

Acknowledgments. We thank the Resource Prospector project at NASA Ames Research Center for providing support for this research and our colleagues Ara V. Nefian and Uland Wong, with whom we had many insightful discussions.

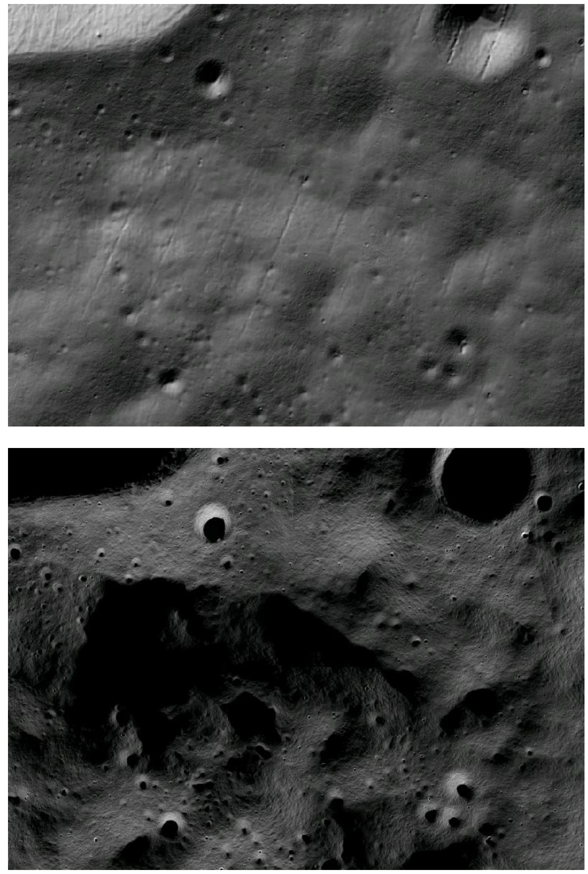

Figure 2: SfS result on a terrain of size $13,000 \times 10,000$ pixels at 1 m/pixel and 87° latitude North (top). The linear artifacts come from the initial guess LOLA gridded terrain that could not be removed by SfS in areas shadowed in all images. The orthoimage (bottom) was obtained by finding at each pixel the maximum image intensity value over all orthoprojected input images.

References

[1] R. Kirk. A fast finite-element algorithm for two-dimensional photoclinometry. *Ph.D. Thesis*, 1987.

[2] Ruo Zhang, Ping-Sing Tsai, James Edwin Cryer, and Mubarak Shah. Shape-from-shading: a survey. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 21(8):690–706, 1999.

[3] Arne Grumpe, Fethi Belkhir, and Christian Wöhler. Construction of lunar dems based on reflectance modelling. *Advances in Space Research*, 53(12): 1735–1767, 2014.

[4] Alfred S McEwen. Photometric functions for photoclinometry and other applications. *Icarus*, 92(2): 298–311, 1991.