

Image-Based Lunar Surface Reconstruction

Stephan Wenger, Anita Sellent, Ole Schütt, and Marcus Magnor

Computer Graphics Lab, TU Braunschweig, Mühlenpfordtstraße 23,
D-38106 Braunschweig, Germany

Abstract. For the creation of a realistic 3 meter-sized relief globe of the Moon, a detailed height map of the entire lunar surface is required. Available height measurements of the Moon's surface are too coarse by a factor of 15 for this purpose. The only publicly available source of high-resolution information are photographic images from the *Lunar Orbiter IV* mission in 1967. We present a shape-from-shading approach to plausibly increase the resolution of existing low-resolution height data, based on a single high-resolution photographic mosaic image of the Moon. The presented reconstruction approach is designed to be robust with respect to frequent imperfections of the photographic imagery. Aside from the automatic reconstruction of a complete detailed lunar surface height map, we give a qualitative validation by the reconstruction of lunar surface details from close-up photographs of the Apollo 15 landing site.

1 Introduction

In July 1969, Neil Armstrong and Edwin Aldrin were the first men to land on the Moon during the *Apollo 11* mission. Forty years and another five manned Moon landings later, much of the Moon's surface structure still remains unrevealed. While many international space missions have been carried out since then [1], the most detailed photographs covering much of the lunar surface are still the ones taken by the *Lunar Orbiter* space probes (1966–1967). The available topographic data surprisingly is a lot more sparse for the Moon than, e.g., for the planet Mars.

Our research project was initiated by the constructors of a Moon museum who noticed that the available lunar surface height data was utterly insufficient for the creation of one of the planned exhibits: a 3 meter-sized globe of the Moon with a realistic surface relief. For convincing effect, the necessary resolution of the lunar model height map would have to be about 3 pixels per millimeter on the model, or about 30 000 pixels around the model's equator. For comparison, the resolution of the best existing height data from the *Unified Lunar Control Network 2005* [3], Fig. 3, is on average a factor of 15 lower.

While the *Lunar Orbiter* data seems to be the best available source for high-resolution photographs of the entire lunar surface, displaying most parts of the Moon with a resolution of 60 meters per pixel or better, this imagery is challenging to interpret by a computer. The conventional photographic emulsion film was developed aboard the spacecraft, then digitized in stripes and transmitted




Fig. 1. The *Lunar Orbiter* mosaic [2] was used for the reconstruction of the lunar surface. The mosaic is stitched together from patches with different quality and varying exposure; some parts are entirely missing. Still it is the most comprehensive source of high-resolution shading information of the lunar surface.

to Earth where the stripes were put back together, all that using the technology of the 1960ies. The many snapshots were combined into a mosaic of the entire Moon, Fig. 1 (available at http://webgis.wr.usgs.gov/pigwad/down/Lunar_Orbiter_mosaic.htm). The quality of the mosaic suffers from stains of photographic developer fluid, missing patches, limited dynamic range (saturation) and ex post high-pass filtering, Fig 2. Additionally, there is a variation in the incident angle of the sunlight: in most images the sunlight is incident from the right, about 20 degrees above the horizon, but both angles change for an unknown amount towards the Moon's poles.

Fortunately, the intended purpose of the reconstructed lunar surface height map does not require exact reconstruction, but rather necessitates a certain visual plausibility of the resulting model. In order to be able to distinguish comparatively flat surface features on the model of the Moon, the actual height data would have to be exaggerated anyway, and the reconstruction algorithm needs to qualitatively reproduce the real Moon's surface.

Making use of the existing low-resolution height data, Fig.3, we present a method to automatically reconstruct high-resolution surface detail based on a shape-from-shading approach applied to high-resolution imagery from the *Lunar Orbiter* mission. The algorithm is designed to be robust to the deficiencies of the input images. In addition, it reconstructs the entire surface of the Moon in a global and consistent way without further user interaction.




Fig. 2. The quality of *Lunar Orbiter* photographs suffers from stains of photographic developer fluid (visible in this picture detail), missing patches, limited dynamic range (i.e. over- and underexposure) and ex post filtering.

2 Related Work

The largest control network for the Moon published today is the *Unified Lunar Control Network 2005 (ULCN2005)* [3]. It combines images from the *Clementine* mission and data from an earlier network which had been derived from Earth-based and Apollo photographs, as well as Mariner 10 and Galileo images of the Moon. This network provides a global lunar topographic model that is denser than that provided by *Clementine* laser altimetry (LIDAR) and of similar accuracy. It consists of 272,931 unevenly distributed measuring points, resulting in an average resolution of about 12 kilometers per pixel. At the time being, higher density topographic data is only available in limited areas of the Moon. The Japanese *Kaguya* mission is aiming to acquire height data at a resolution of about 2 km per pixel, but until now only 30 km per pixel data has been published [4].

Up to day, the most comprehensive, high-resolution coverage of the lunar surface is achieved by the monocular images acquired by *Lunar Orbiter* [2]. Using monocular photographs to determine 3D structure has a long tradition in remote sensing, where it is known as photoclinometry, as well as in machine vision, where it is known as shape-from-shading (SFS). Since the first solutions introduced by Rindfleisch [5] and Horn [6], this approach has run through many refinements [7,8]. However, none of these approaches addresses the image imperfections one encounters in *Lunar Orbiter* images.

SFS is an underdetermined problem as it assigns the two directional angles of inclination based on one measured gray scale value. With image acquisition in machine vision growing cheaper and cheaper, today most algorithms for height or depth estimation rely on several images. The most common approach is stereopsis using stereo image pairs acquired from different viewpoints but under the same lighting conditions [9]. Multi-image shape-from-shading unites the advantages of stereo with SFS. It considers the reflection model and uses images acquired under the same lighting conditions from different viewing directions [10]. Image pairs that depict the lunar surface under comparable lighting conditions are rare – especially on the far side of the Moon – and were acquired only in low




Fig. 3. This low-resolution height map from the *ULCN2005* network [3] is the best publicly available height information. We use it to initialize our reconstruction method of the entire moon surface.

resolution during the *Clementine* mission. This information is already included in the *ULCN2005* height data.

Another approach to obtain height information is to consider several images acquired under different lighting conditions: *Clementine* images and ground-based telescopic CCD images were used to reconstruct 3D elevation information for certain Moon regions in Ref. [11], [12] and [13]. Still these methods do not obtain the resolution required for our purpose. In their recent work, Glencross et al. [14] concentrate on perceptually plausible height map reconstruction. As input, they require a pair of one diffusely illuminated and one flashlight illuminated image. Although the *Lunar Orbiter* and *Clementine* images were acquired under different lighting conditions, the *Lunar Orbiter* mosaic is already high-pass filtered in order to account for slowly varying albedo. The *Clementine* images only provide pure albedo measurements. Thus these images cannot be used as a comparable input to constrain the solutions of the SFS problem.

A great part of SFS literature directly calculates height information instead of estimating surface normals [15,16]. In our algorithm, we divide normal estimation and height estimation into two steps adapting the integration algorithm of Smith and Bors [17]. This allows us, on the one hand, to easily incorporate the low resolution height field given with *ULCN2005* and, on the other hand, to weight estimated normal information with a *credibility map*.

3 Algorithm

Our algorithm is designed to deal automatically with the imperfections of the monocular high-resolution *Lunar Orbiter* images. In order to calculate a global




Fig. 4. SFS on Lambertian surfaces determines only the angle between normal n and incident light l . We pick the normal in direction of the image gradient ∇I that is closest to the normal of a flat surface.

height map of the Lunar surface, our algorithm proceeds in two steps. In a first step, we calculate normals wherever information is available. In a second step, a given low-resolution height map, e.g. the *ULCN2005*, is iteratively refined until it closely fits the reconstructed normals. The resulting height map is used as a basis for another reconstruction step, iteratively increasing the resolution.

Since the amount of data being handled during the reconstruction of the whole lunar surface easily exceeds the available memory of recent desktop computers, large height maps are cut into overlapping pieces that are reconstructed separately. The results can be blended without problems – using a suitable continuous weighting function, e.g. a linear ramp – because the long-range coherence of the resulting height map is ensured by the lower resolution height map used as a basis. The same procedure is used to ensure wrap-around continuity at the left and right image borders.

3.1 Normal Estimation

In order to estimate the normal vector for each pixel based on its intensity, we first assume Lambertian reflectance of the lunar surface so that the angle α between the normal vector \mathbf{n} and the incident light direction \mathbf{l} can be computed from the observed intensity I via

$$I = \mathbf{l} \cdot \mathbf{n} = \cos \alpha \quad (1)$$

where $I \in [0, 1]$. While the Moon’s surface is not perfectly approximated by a Lambertian reflector [18], the error introduced by this assumption vanishes in comparison to the error caused by the unknown deviation of the light source from the position at the right side and at an angle of 20 degrees over the horizon.

Knowing α , the normal vector is only restricted to a circle around the light direction vector \mathbf{l} . In order to entirely fix the normal vector, another constraint is needed. For typical lunar geometries, the height gradient (and thus the projection of the normal vector onto the horizontal plane) is likely to be approximately

collinear with the intensity gradient of the image, Fig. 4. This assumption proves reasonable because important height map features like rims and ridges cause strong intensity gradients (as long as they are not parallel to the incident light direction), while variations in the direction of the ridges – which might cause an intensity gradient that violates the assumption – are usually on such large scales that the associated intensity gradient is small, cf. Fig. 6(a). If we use this assumption to further constrain the normal vector, at most two possible normal vectors remain. We select the one that is closer to the normal vector of a flat surface. (For the incident light angle of only about 20 degrees above the horizon, the other possible normal vector would usually represent an almost vertical wall that is highly unlikely.)

Note that the input data is presented in a cylindrical projection. Therefore, the x coordinate has to be scaled by the cosine of the latitude whenever image gradients are calculated in order to maintain the correct length scale throughout the whole map.

3.2 Credibility Map

Because of the challenging input data, some precautions have to be taken in order to compensate for shortcomings of the photographic images. Regions that are saturated or underexposed do not yield any gradient information. They are assigned a credibility of zero. Towards saturated or underexposed regions, the credibility of usable pixels decreases with a Gaussian function to ensure smooth and plausible transitions. Additionally, all image gradients are smoothed using a Gaussian filter.

The following normal integration step then takes care of enforcing continuity between the heights of neighboring pixels.

3.3 Normal Integration

The reconstructed normal vectors have to be *integrated* in order to obtain the final height map. We adapt an iterative algorithm by Smith and Bors [17] which we extend so that it takes the *credibility map* into account. The algorithm iteratively modifies a low-resolution height map, changing the pixel heights in order to approximate the specified normal vectors. In each step i , the height $\tilde{h}_i(x, y)$ dictated by the normal map is computed for each pixel (x, y) from the current heights $h_i(x, y)$ of its neighbors and the x and y normal vector components $n_1(x, y)$ and $n_2(x, y)$ as

$$\tilde{h}_i(x, y) = \frac{1}{4} \sum_{(u,v) \in N} (h_i(x+u, y+v) + un_1(x+u, y) + vn_2(x, y+v)) \quad , \quad (2)$$

where $N = \{(\pm 1, 0), (0, \pm 1)\}$. We weight h and \tilde{h} by a function $\eta_i(x, y) = \frac{\eta_0 c(x, y)}{1 + 2^i / (w + h)}$ which is proportional to the credibility $c(x, y)$ of the corresponding pixel and decreases with iteration i in order to enforce convergence. w and h




Fig. 5. The result of our reconstruction algorithm (top) for large-scale photographic input data (bottom), compared to the initial *ULCN2005* height map (middle). Plausible surface detail has been added where shading information was available; the missing areas of the photograph were recognized as invalid and have therefore remained unmodified. Note how many surface features become recognizable in the reconstruction that were not present in the initial height map.

are the image width and height, respectively, and η_0 was set to 0.2. The heights are then updated by

$$h_{i+1}(x, y) = (1 - \eta_i(x, y))h_i(x, y) + \eta_i(x, y)\tilde{h}_i(x, y) . \quad (3)$$

Particularly, the initial height map doesn't change at all when the credibility is zero, i.e. in regions without any available image data. The iteration stops as soon as the average difference between \tilde{h}_i and h_i falls below a specified threshold that we set to 0.005.

4 Results

The goal of our algorithm is to produce a plausible high resolution height map of the entire lunar surface in a fully automated way. Because of the lack of high resolution ground truth data and the aim of perceptual plausibility rather than accuracy, we give two examples of the typical results of our reconstruction

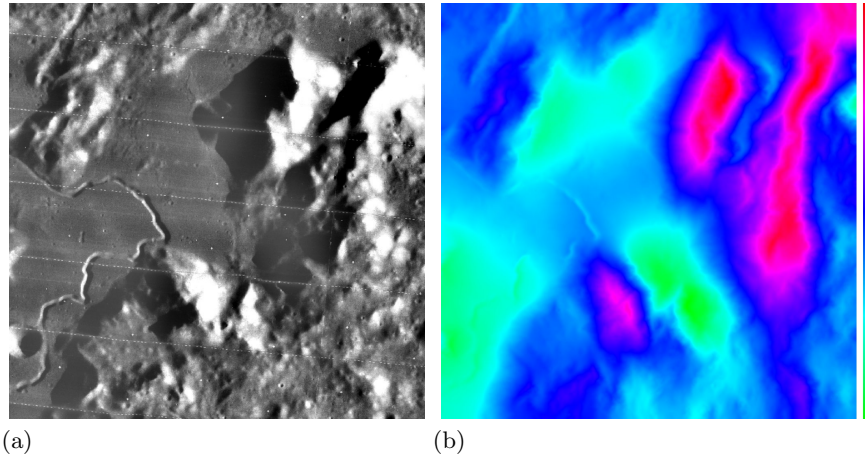


Fig. 6. Input close-up photograph (a) and resulting height field (b) of the Apollo 15 landing site near Hadley rille.

algorithm that can be visually evaluated. Figure 5 shows our reconstruction of an approximately 5000 km by 1000 km patch from the far side of the Moon close to the equator. In regions where the image data is usable, the perceived resolution of the heightmap is increased, while regions for which no suitable data is available remain unaltered. Note how e.g. small craters are added to formerly flat regions.

In the second example we show that our reconstruction algorithm also works on much smaller scales. We reconstruct an approximately 64 km by 64 km region around the Apollo 15 landing site. For this region, images acquired during extravehicular activity are available that permit perceptual validation of reconstructed heights. At this scale, no reasonable height data is available, so the height map was initialized as a plane and updated with the single photographic image shown in Figure 6(a). The reconstructed height map is displayed in Figure 6(b). The human observer easily recognizes the reconstructed surface features of the photographic image. This is even more apparent in the comparison of the rendered height map with an actual photograph of the site shown in Fig. 7. However, due to the little amount of input data to our algorithm, some limitations remain: Of course, there are surface features which cannot be determined metrically correct based on one image acquired with fixed lighting conditions, e.g. the small rille in the left part of the image cannot be reconstructed where it runs parallel to the incident light direction. Also, the saturated regions and shadows close to high mountains cause overshooting effects in some places, but still the result looks plausible to the human observer and reproduces the important geographical features well enough to make the region easily recognizable.




Fig. 7. Apollo 15 surface panoramic photograph (top) and perspective rendering of the reconstructed height map from a similar viewpoint (bottom) allow for visual validation of the presence of important surface features.

5 Conclusion and Discussion

We have presented a shape-from-shading reconstruction method for lunar surface geometry that is based on known low-resolution height data and single high-resolution photographic images. While large-scale coherence of the height data is inherited from the low-resolution data, surface detail is plausibly added based on shading information. The algorithm is robust with respect to the many flaws present in high-resolution lunar surface imagery. It has successfully been used to reconstruct a detailed height map of the entire lunar surface based on *ULCN2005* height data and imagery from the *Lunar Orbiter* mission. In spite of the quality deficits of the *Lunar Orbiter* images, the algorithm strongly increases the perceived resolution and richness of detail of the height map. The reconstruction algorithm is able to detect typical error sources and assigns a lower credibility value to the corresponding regions so that the known height data is left unchanged where no better information is available.

References

1. Kirk, R., Archinal, B.A., Gaddis, L.R., Rosiek, M.R.: Cartography for lunar exploration: 2008 status and mission plans. *European Planetary Science Congress* **3** (2008)
2. United States Geological Survey. http://webgis.wr.usgs.gov/pigwad/down/Lunar_Orbiter_mosaic.htm
3. Archinal, B.A., Rosiek, M.R., Kirk, R.L., Redding, B.L.: Completion of the Unified Lunar Control Network 2005 and topographic model. In: 37th Annual Lunar and Planetary Science Conference. Volume 37. (2006) 2310–2311
4. Araki, H., Tazawa, S., Noda, H., Ishihara, Y., Goossens, S., Sasaki, S., Kawano, N., Kamiya, I., Otake, H., Oberst, J., Shum, C.: Lunar global shape and polar topography derived from Kaguya-LALT laser altimetry. *Science* **323**(5916) (2009) 897–900
5. Rindfleisch, T.: Photometric method for lunar topography. *Photogrammetric Engineering* **32**(2) (1966) 262–27
6. Horn, B.K.P.: Shape from Shading: a Method for Obtaining the Shape of a Smooth Opaque Object from one View. PhD thesis, Department of Electrical Engineering, MIT (1970)
7. Horn, B.K.P.: Height and gradient from shading. *Int. J. of Computer Vision* **5**(1) (1990) 37–75
8. Zhang, R., Tsai, P.S., Cryer, J.E., Shah, M.: Shape-from-shading: a survey. *IEEE T-PAMI* **21**(8) (1999) 690–706
9. Scharstein, D., Szeliski, R.: A taxonomy and evaluation of dense two-frame stereo correspondence algorithms. *Int. J. of Computer Vision* **47**(1) (2002) 7–42
10. Heipke, C., Piechullek, C., Ebner, H.: Simulation studies and practical tests using multi-image shape from shading. *J. of Photogrammetry and Remote Sensing* **56**(2) (2001) 139–148
11. Wöhler, C., Hafezi, K.: A general framework for three-dimensional surface reconstruction by self-consistent fusion of shading and shadow features. *Pattern Recognition* **38**(7) (2005) 965–983
12. Lena, R., Wöhler, C., Bregante, M.T., Fattinanzi, C.: A combined morphometric and spectrophotometric study of the complex lunar volcanic region in the south of Petavius. *J. of the RASC* **100**(1) (2006) 14
13. Wöhler, C., Lena, R., Lazzarotti, P., Phillips, J., Wirths, M., Pujic, Z.: A combined spectrophotometric and morphometric study of the lunar mare dome fields near Cauchy, Arago, Hortensius, and Milichius. *Icarus* **183**(2) (2006) 237–264
14. Glencross, M., Ward, G.J., Jay, C., Liu, J., Melendez, F., Hubbard, R.: A perceptually validated model for surface depth hallucination. In: *ACM Transactions on Graphics*. Volume 27. (2008) 1–8
15. Worthington, P., Hancock, E.: New constraints on data-closeness and needle map consistency for shape-from-shading. *IEEE T-PAMI* **21**(12) (1999) 1250–1267
16. Frankot, R.T., Chellappa, R.: A method for enforcing integrability in shape from shading algorithms. *IEEE T-PAMI* **10**(4) (1988) 439–451
17. Smith, G.D.J., Bors, A.G.: Height estimation from vector fields of surface normals. In: 14th Int. Conf. on Digital Signal Processing. Volume 2. (2002) 1031–1034
18. Wildey, R.L.: The Moon’s photometric function. *Nature* **200**(4911) (1963) 1056–1058