

Shape from shading under coplanar light sources

Christian Wöhler

DaimlerChrysler Research and Technology, Machine Perception
P. O. Box 2360, D-89013 Ulm, Germany
E-Mail: christian.woehler@daimlerchrysler.com

Abstract. In this paper image-based techniques for 3D surface reconstruction are presented which are especially suitable for (but not limited to) coplanar light sources. The first approach is based on a single-image shape from shading scheme, combined with the evaluation of at least two further images of the scene that display shadow areas. The second approach allows the reconstruction of surfaces by an evaluation of quotients of images of the scene acquired under different illumination conditions and is capable of separating brightness changes due to surface shape from those caused by variable albedo. A combination of both techniques is suggested. The proposed approaches are applied to the astrogeological task of three-dimensional reconstruction of regions on the lunar surface using ground-based CCD images. Beyond the planetary science scenario, they are applicable to classical machine vision tasks such as surface inspection in the context of industrial quality control.

1 Introduction

A well-known method for image-based 3D reconstruction of surfaces is *shape from shading* (SFS). This technique aims at deriving the orientation of the surface at each pixel by using a model of the reflectance properties of the surface and knowledge about the illumination conditions.

1.1 Related work

Traditional applications of such techniques in planetary science, mostly referred to as *photoclinometry*, rely on single images of the scene and use line-based, integrative methods designed to reveal a set of profiles along one-dimensional lines rather than a 3D reconstruction of the complete surface [1, 5]. In contrast to these approaches, *shape from shading* and *photometric stereo* techniques based on the minimization of a global error term for multiple images of the scene have been developed in the field of computer vision – for detailed surveys on the SFS and photometric stereo methodology see [1, 3]. For a non-uniform surface albedo, however, these approaches require that the directions of illumination are not coplanar [3]. Recent work in this field [6] deals with the reconstruction of planetary surfaces based on shape from shading by means of multiple images acquired from precisely known locations at different illumination conditions, provided that the reflectance properties of the surface are thoroughly modelled. In

[7] shadows are used in the context of photometric stereo with multiple non-coplanar light sources to recover locally unique surface normals from two image intensities and a zero intensity caused by shadow.

1.2 Shape from shading: an overview

For a single image of the scene, parallel incident light and an infinite distance between camera and object the intensity $I(u, v)$ of image pixel (u, v) amounts to

$$I(u, v) = \kappa I_i \Phi(\mathbf{n}(x, y, z), \mathbf{s}, \mathbf{v}). \quad (1)$$

Here, κ is a camera constant, \mathbf{v} the direction to the camera, I_i the intensity of incident light and \mathbf{s} its direction, and Φ the so-called *reflectance function*. A well-known example is the Lambertian reflectance function $\Phi(\mathbf{n}, \mathbf{s}) = \alpha \cos \theta$ with $\theta = \angle(\mathbf{n}, \mathbf{s})$ and α as a surface-specific constant. The product $\kappa I_i \alpha = \rho(u, v)$ is called *surface albedo*. In the following, the surface normal \mathbf{n} will be represented by the directional derivatives $p = z_x$ and $q = z_y$ of the surface function $z(u, v)$ with $\mathbf{n} = (-p, -q, 1)$. The term $R(p, q) = \kappa I_i \Phi$ is called *reflectance map*. Solving the SFS problem requires to determine the surface $z(u, v)$ with gradients $p(u, v)$ and $q(u, v)$ that minimizes the average deviation between the measured pixel intensity $I(u, v)$ and the modelled reflectance $R(p(u, v), q(u, v))$. This corresponds to minimizing the *intensity error term*

$$e_i = \sum_{u,v} [I(u, v) - R((p(u, v), q(u, v)))]^2. \quad (2)$$

Surface reconstruction based on a single monocular image with no constraints is an ill-posed problem as for a given image $I(u, v)$ there exists an infinite number of minima of e_i for the unknown values of $p(u, v)$, $q(u, v)$, and $\rho(u, v)$. A well-known method to alleviate this ambiguity consists of imposing regularization constraints on the shape of the surface. A commonly used constraint is smoothness of the surface, implying small absolute values of the directional derivatives p_x and q_x of the surface gradients p and q (cf. [1, 3]). This leads to an additional error term

$$e_s = \sum_{u,v} [p_x^2 + p_y^2 + q_x^2 + q_y^2]. \quad (3)$$

Solving the problem of surface reconstruction then consists of globally minimizing the overall error function $e = e_s + \lambda e_i$, where the Lagrangian multiplier λ denotes the relative weight of the error terms. As explained in detail in [1–3], setting the derivatives of e with respect to the surface gradients p and q to zero leads to an iterative update rule to be repeatedly applied pixelwise until convergence of $p(u, v)$ and $q(u, v)$ is achieved. Once the surface gradients are determined, the surface profile $z(u, v)$ is obtained by numerical integration of the surface gradients as described in [3]. According to [2], constraint (3) can be replaced by or combined with the physically intuitive assumption of an integrable surface gradient vector field within the same variational framework.

The ambiguity of the solution of the shape from shading problem can only be completely removed, however, by means of *photometric stereo* techniques, i. e. by making use of several light sources. A traditional approach is to acquire $L = 3$ images of the scene (or an even larger number) at different illumination conditions represented by \mathbf{s}_l , $l = 1, \dots, L$. As long as these vectors are not coplanar and a Lambertian reflectance map can be assumed, a unique solution for both the surface gradients $p(u, v)$ and $q(u, v)$ and the non-uniform surface albedo $\rho(u, v)$ can be obtained analytically in a straightforward manner [3].

In many application scenarios, however, it is difficult or impossible to obtain $L \geq 3$ images acquired with non-coplanar illumination vectors \mathbf{s}_l , $l = 1, \dots, L$. For example, the equatorial regions of the Moon (only these appear nearly undistorted for ground-based telescopes) are always illuminated either exactly from the east or exactly from the west, such that all possible illumination vectors \mathbf{s} are coplanar. The illumination vectors are thus given by $\mathbf{s}_l = (-\cot \mu_l, 0, 1)$ with μ_l denoting the solar elevation angle for image l . Identical conditions occur e. g. for the planet Mercury and the major satellites of Jupiter. In scenarios beyond planetary science applications, such as visual quality inspection systems, there is often not enough space available to sufficiently distribute the light sources.

Hence, this paper proposes shape from shading techniques based on multiple images of the scene, including the evaluation of shadows, which are especially suitable for (but not limited to) the practically very relevant case of coplanar light sources, and which do not require a Lambertian reflectance map.

2 Shadow-based initialization of the SFS algorithm

A uniform albedo $\rho(u, v) = \rho_0$ and oblique illumination (which is necessary to reveal subtle surface details) will be assumed throughout this section. Despite this simplification the outcome of the previously described SFS scheme is highly ambiguous if only one image is used for reconstruction, and no additional information is introduced by further shading images due to the coplanarity of the illumination vectors. Without loss of generality it is assumed that the scene is illuminated exactly from the left or the right hand side. Consequently, the surface gradients $q(u, v)$ perpendicular to the direction of incident light cannot be determined accurately for small illumination angles by SFS alone unless further constraints, e. g. boundary values of $z(u, v)$ [1], are imposed.

Hence, a novel concept is introduced, consisting of a shadow analysis step performed by means of at least two further images (in the following called “shadow images”) of the scene acquired under different illumination conditions (Fig. 1a). All images have to be pixel-synchronous such that image registration techniques (for a survey cf. [4]) have to be applied. After image registration is performed, shadow regions can be extracted either by a binarization of the shadow image or by a binarization of the quotient of the shading and the shadow image. The latter technique prevents surface parts with a low albedo from being erroneously classified as shadows; it will therefore be used throughout this paper. A suitable binary threshold is derived by means of histogram analysis in a straightforward




Fig. 1. Shadow-based initialization of the SFS algorithm. (a) Shadow and shading images. The region inside the rectangular box is reconstructed. (b) Surface part between the shadows. (c) Initial 3D profile $\tilde{z}_0(u, v)$ of the surface patch between the shadows.

manner. The extracted shadow area is regarded as being composed of S *shadow lines*. As shown in Fig. 1b, shadow line s has a length of $l^{(s)}$ pixels. This corresponds to an altitude difference $(\Delta z)_{\text{shadow}}^{(s)} = l^{(s)} \tan \mu_{\text{shadow}}$, where the angle μ_{shadow} denotes the elevation angle of the light source that produces the shadow.

Altitude difference $(\Delta z)_{\text{shadow}}^{(s)}$ can be determined at high accuracy as it is independent of a reflectance model. It is used to introduce an additional shadow-based error term e_z into the variational SFS scheme according to

$$e_z = \sum_{s=1}^S \left[\frac{(\Delta z)_{\text{sfs}}^{(s)} - (\Delta z)_{\text{shadow}}^{(s)}}{l^{(s)}} \right]^2. \quad (4)$$

This leads to a minimization of the overall error term $e = e_s + \lambda e_i + \eta e_z$ with η as an additional Lagrangian multiplier, aiming at an adjustment of the altitude difference $(\Delta z)_{\text{sfs}}^{(s)}$ measured on the surface profile along the shadow line to the altitude difference obtained by shadow analysis. For details cf. [8].

This concept of shading and shadow based 3D surface reconstruction is extended by initializing the SFS algorithm based on two or more shadow images, employing the following iterative scheme:

1. Initially, it is assumed that the altitudes of the ridges casting the shadows (solid line in Fig. 1b) are constant, respectively, and identical. The iteration index m is set to $m = 0$. The 3D profile $\tilde{z}_m(u, v)$ of the small surface patch between the two shadow lines (hatched area in Fig. 1b) is derived from the measured shadow lengths (Fig. 1c).
2. The surface profile $\tilde{z}_m(u, v)$ directly yields the surface gradients $p_0(u, v)$ and $q_0(u, v)$ for all pixels belonging to the surface patch between the shadow lines. They are used to compute the albedo ρ_0 , serve as initial values for the SFS algorithm, and will be kept constant throughout the following steps of the algorithm. Outside the region between the shadow lines, $p_0(u, v)$ and $q_0(u, v)$ are set to zero.
3. Using the single-image SFS algorithm with the initialization applied in step 2, the complete surface profile $z_m(u, v)$ is reconstructed based on the shading image. The resulting altitudes of the ridges casting the shadows are extracted from the reconstructed surface profile $z_m(u, v)$. This yields a new profile $\tilde{z}_{m+1}(u, v)$ for the surface patch between the shadow lines.

4. The iteration index m is incremented ($m := m + 1$). The algorithm cycles through steps 2, 3, and 4 until it terminates once convergence is achieved, i. e. $\langle (z_m(u, v) - z_{m-1}(u, v))^2 \rangle_{u,v}^{1/2} < \Theta_z$. A threshold value of $\Theta_z = 0.01$ pixels is applied for termination of the iteration process.

This approach is applicable to arbitrary reflectance functions $R(p, q)$. It mutually adjusts in a self-consistent manner the altitude profiles of the floor and of the ridges that cast the shadows. It allows to determine not only surface gradients $p(u, v)$ in the direction of incident light, as it can be achieved by SFS without additional constraints, but to estimate surface gradients $q(u, v)$ in the perpendicular direction as well. Furthermore, it can be extended in a straightforward manner to more than two shadows and to shape from shading algorithms based on multiple light sources or regularization constraints beyond those described by eq. (3) and (4).

3 Quotient-based photometric stereo

The second approach to SFS under coplanar light sources copes with a non-uniform albedo $\rho(u, v)$ and the very general class of reflectance maps given by $R(\rho, p, q) = \rho(u, v)\tilde{R}(p, q)$. At least two pixel-synchronous images of the scene acquired under different illumination conditions and containing no shadow areas are required. For each pixel position (u, v) , the quotient $I_1(u, v)/I_2(u, v)$ of pixel intensities is desired to be identical to the quotient $R_1(u, v)/R_2(u, v)$ of reflectances. This suggests a quotient-based intensity error term

$$\tilde{e}_i = \sum_{u,v} \left(\frac{I_1(u, v)\tilde{R}_2(u, v)}{I_2(u, v)\tilde{R}_1(u, v)} - 1 \right)^2 \quad (5)$$

which is independent of the albedo (cf. [5] for a quotient-based approach for merely one-dimensional profiles). This error term can easily be extended to $L > 2$ images by computing the $L(L - 1)/2$ quotient images from all available image pairs and summing up the corresponding errors. This method allows to separate brightness changes due to surface shape from those caused by variable albedo. Similar to SFS with a single image and constant albedo, however, the values for $q(u, v)$ obtained with this approach are quite uncertain as long as the illumination vectors are coplanar, so error term (5) should be combined with the shadow-based approach of Section 2 provided that corresponding shadow information is available.

4 Experimental results

Fig. 2 illustrates the performance of the proposed algorithms on a synthetically generated object (Fig. 2a). The shadow-based technique outlined in Section 2 and utilizing intensity error term (2) reveals the surface gradients in image v direction

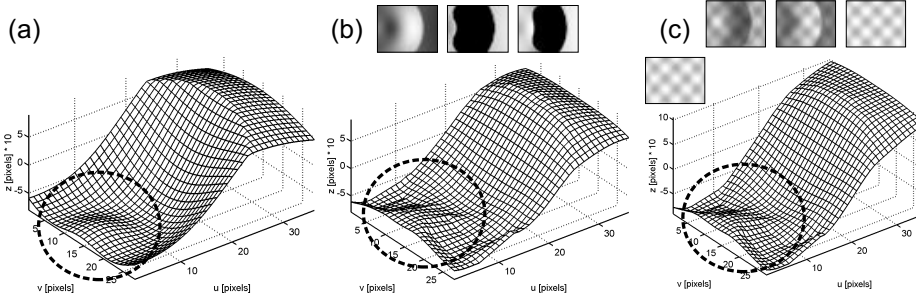


Fig. 2. Surface reconstruction of synthetic images. (a) Ground truth surface profile. (b) One shading and two shadow images (top, $\mu_{\text{SFS}} = 175^\circ$, $\mu_{\text{shadow}}^{(1)} = 4.0^\circ$, $\mu_{\text{shadow}}^{(2)} = 5.0^\circ$) along with surface profile (bottom) obtained according to Section 2. (c) Two shading images and true albedo map (top, $\mu_{\text{SFS}}^{(1)} = 15^\circ$, $\mu_{\text{SFS}}^{(2)} = 170^\circ$) along with the reconstructed surface profile and albedo map (bottom) obtained by using the quotient-based error term (5) instead of single-image error term (2).

(dashed circles). Traditional single-image SFS as outlined in Section 1.2 is not able to extract these surface gradients. As suggested in Section 3, the single-image error term (2) was then replaced by the quotient-based error term (5) for a reconstruction of the same synthetic object but now with a non-uniform albedo (Fig. 2c). Consequently, two shading images are used in combination with the shadow information. As a result, a similar surface profile is obtained, and the albedo is extracted at an accuracy (root mean square error) of 1.1 percent.

For 3D reconstruction of regions on the lunar surface it is possible to use a Lambertian reflectance map because for small parts of the lunar surface, the relative Lambertian reflectance (an absolute calibration of the images is not necessary) differs by only a few percent from those values derived from more sophisticated models such as the Lunar-Lambert function [6] – the presented framework, however, can also be applied to non-Lambertian reflectance models.

The CCD images were acquired with ground-based telescopes of 125 mm and 200 mm aperture. Image scale is 800 m per pixel.

Fig. 3 shows the reconstructed surface profile of the floor of lunar crater Theatetus, generated by the technique outlined in Section 2. Both the simulated shading image and the shapes of the simulated shadows correspond well with their real counterparts. Even the ridge crossing the crater floor, which is visible in the upper left corner of the region of interest in Fig. 1a and in the Lunar Orbiter photograph in Fig. 3d shown for comparison, is apparent in the reconstructed surface profile (arrow). Furthermore, it turns out that the crater floor is inclined from the north to the south, and a very shallow central elevation rising to about 250 m above floor level becomes apparent. This central elevation does not appear in the images in Fig. 1a used for reconstruction, but is clearly visible in the ground-based image acquired at higher solar elevation shown in Fig. 3e (left, lower arrow). The simulated image (right half of Fig. 3e) is very similar to the corresponding part of the real image although that image has not been used for




Fig. 3. Reconstruction result for the western part of lunar crater Theaetetus (see Fig. 1 for original images). (a) Simulated regions of interest. (b) Reconstructed surface profile. The z axis is two-fold exaggerated. (c) Reconstructed surface profile with absolute z values. (d) Lunar Orbiter image IV-110-H2 (not used for reconstruction). (e) Ground-based image, solar elevation $\mu = 28.7^\circ$, real image (left) and simulated image (right). This image has not been used for reconstruction. (f) Reconstruction result obtained with traditional SFS, selecting the solution consistent with the first shadow image.




Fig. 4. Reconstruction result for the region around lunar dome Herodotus ω . (a) Original images. Solar elevation angles are $\mu_1 = 5.0^\circ$ and $\mu_2 = 15.5^\circ$. (b) Albedo map. (c) Reconstructed surface profile. The slight overall bending of the surface profile reflects the Moon's spherical shape.

reconstruction. This kind of comparison is suggested in [2] as an independent test of reconstruction quality. For comparison, traditional SFS as outlined in Section 1.2 yields an essentially flat crater floor and no ridge (Fig. 3f). Here, the uniform surface albedo was adjusted to yield an SFS solution consistent with the first shadow image (for details cf. [8]).

Fig. 4 shows the region around lunar dome Herodotus ω , obtained with the quotient-based approach outlined in Section 3. The images are rectified due to the proximity of this region to the moon's apparent limb. The reconstructed sur-

face profile contains several shallow ridges with altitudes of roughly 50 m along with the lunar dome, whose altitude was determined to 160 m. The resulting albedo map displays a gradient in surface brightness from the lower right to the upper left corner along with several ray structures running radially with respect to the crater Aristarchus.

5 Summary and conclusion

In this paper, shape from shading techniques for 3D reconstruction of surfaces under coplanar light sources are proposed. The first presented method relies on an initialization of the surface gradients by means of the evaluation of a pixel-synchronous set of at least one shading image and two shadow images and yields reliable values also for the surface gradients perpendicular to the direction of illumination. The second approach is based on at least two pixel-synchronous, shadow-free images of the scene acquired under different illumination conditions. A shape from shading scheme relying on an error term based on the quotient of pixel intensities is introduced which is capable of separating brightness changes due to surface shape from those caused by variable albedo. A combination of both approaches has been demonstrated.

In contrast to traditional photometric stereo approaches, both presented methods can cope with coplanar illumination vectors. They are successfully applied to synthetically generated data and to the 3D reconstruction of regions on the lunar surface using ground-based CCD images. The described techniques should be as well suitable for space-based exploration of planetary surfaces. Beyond the planetary science scenario, they are applicable to classical machine vision tasks such as surface inspection in the context of industrial quality control.

References

1. B. K. P. Horn. Shape from Shading. MIT Press, Cambridge, Massachusetts, 1989.
2. B. K. P. Horn. Height and Gradient from Shading. MIT technical report, 1989. <http://www.ai.mit.edu/people/bkph/papers/newsfs.pdf>
3. X. Jiang, H. Bunke. Dreidimensionales Computersehen. Springer-Verlag, Berlin, 1997.
4. L. Gottesfeld Brown. A Survey of Image Registration Techniques. ACM Computing Surveys, vol. 24, no. 4, pp. 325-376, 1992.
5. A. S. McEwen. Albedo and Topography of Ius Chasma, Mars. *Lunar and Planetary Science XVI*, pp. 528-529, 1985.
6. C. Piechullek. Oberflächenrekonstruktion mit Hilfe einer Mehrbild-Shape-from-Shading-Methode. Ph. D. thesis, Technical University of Munich, Munich, 2000.
7. K. Schlüns. Shading Based 3D Shape Recovery in the Presence of Shadows. *Proc. First Joint Australia & New Zealand Biennial Conference on Digital Image & Vision Computing*, Albany, Auckland, New Zealand, pp. 195-200, 1997.
8. C. Wöhler. 3D surface reconstruction by self-consistent fusion of shading and shadow features. Accepted for publication at *International Conference on Pattern Recognition*, Cambridge, UK, 2004.