Computer Science Department of The University of Auckland
CITR at Tamaki Campus (http://www.tcs.auckland.ac.nz**)**

CS-TR-157 / CITR-TR-14 September 1997

Shading Based 3D Shape Recovery in the Presence of Shadows

Karsten Schlüns*

Abstract

Shadows usually cause various problems in three-dimensional shape recovery and measurement methods. In particular shading based approaches such as shape-from-shading or the photometric stereo method produce no or wrong results if the shadows are not treated appropriately. We show how information extracted from shadows can be employed to reduce the problems caused by them. This is done for multiple light-source photometric stereo. Unlike other published work, we formulate sufficient conditions to recover locally unique surface normals from two image irradiances (intensities) and a zero-irradiance caused by a shadow. We also distinguish between self-shadows and cast-shadows. Moreover we show how much information is obtainable by using the shadow analysis.

* The University of Auckland, Tamaki Campus, Computing and Information Technology Research, Computer Vision Unit, Auckland, New Zealand

Shading Based 3D Shape Recovery in the Presence of Shadows

Karsten Schlüns

Computing and Information Technology Research – CITR
Computer Science Department, Tamaki Campus
University of Auckland
New Zealand
E-mail:karsten@cs.auckland.ac.nz

Abstract

Shadows usually cause various problems in three-dimensional shape recovery and measurement methods. In particular shading based approaches such as shape-from-shading or the photometric stereo method produce no or wrong results if the shadows are not treated appropriately. We show how information extracted from shadows can be employed to reduce the problems caused by them. This is done for multiple light-source photometric stereo. Unlike other published work, we formulate sufficient conditions to recover locally unique surface normals from two image irradiances (intensities) and a zero-irradiance caused by a shadow. We also distinguish between self-shadows and cast-shadows. Moreover we show how much information is obtainable by using the shadow analysis.

Keywords*: photometric stereo method (PSM), 3D shape recovery, self-shadows, cast-shadows.*

1. Introduction and Related Work

Shading based methods calculate surface normals and/or depth values by employing reflection and illumination models. These methods can be classified by the number of light sources they use. An image is taken separately for each light source, i.e. the light sources are sequentially switched on and off. During the image acquisition process the position of the optical sensor is not changed. In general white light is used but under certain assumptions image acquisition can be done in parallel (all light sources switched on) if colored light is taken.

For diffuse reflecting object surfaces (Lambertian surfaces) a surface normal can be determined to be locally unique if the considered surface point is illuminated from three light sources [1, 2, 6, 7]. If only two non-zero image irradiances (intensities) are observed two surface normals are consistent with the pair of image irradiances. This paper addresses the question of how this ambiguity can be locally resolved.

There are different possible ways to reduce/ overcome this problem that we *do not* intend to take:

- using more than three light sources,
- reducing the angles between the illumination directions,
- independently applying one-source methods and two-source methods, or
- integrating other vision modules (e.g., binocular stereo).

Instead, we analyze the shadows in the images to overcome this problem. Other authors [4, 5] have observed the usefulness of shadow information as well. Unlike these publications we will give sufficient conditions when all visible surface normals of an illumination configuration (definitions see below) are uniquely recoverable. These publi-

cations assume convex objects and hence they do not have to distinguish between self-shadows and cast-shadows. This is done for shadows caused by one light source.

2. Shadow Analysis

We assume the standard geometry of shape-from-shading. Let the object surface be explicitly given in a left-handed sensor coordinate system XYZ by $Z = Z(X,Y)$. The image xy axes coincide with the XY axes of the sensor coordinate system. The view direction is $\mathbf{v} = (0,0,-1)$ and the viewer is looking in the positive Z direction. Hence a surface normal can be defined as

$$\mathbf{n} = (\partial Z(X,Y)/\partial X, \partial Z(X,Y)/\partial Y, -1).$$

We define a surface normal $\mathbf{n}$ as being *visib*le if $\measuredangle(\mathbf{v},\mathbf{n}) < 90°$. Unit surface normals are denoted by $\mathbf{n}°$. We define

$$V = \left\{ \mathbf{n}° \in \Re^3 \,\middle|\, \mathbf{n}°^T \mathbf{v} > 0 \right\}$$

as the set of visible surface normals.

The illumination directions are denoted by $\mathbf{s}_1, \mathbf{s}_2, \mathbf{s}_3 \in \Re^3$. They point towards the light sources. We call a surface normal $\mathbf{n}$ as being *illuminated* by a light source with illumination direction $\mathbf{s}$ if $\measuredangle(\mathbf{s},\mathbf{n}) < 90°$ holds, otherwise the surface normal is *unilluminated*.

It is assumed that the object surface has a positive albedo and that the specular reflection component (if it exists) was eliminated beforehand by using an appropriate technique [3].

The light source λ generates a self-shadow on an object point with surface normal $\mathbf{n}$ if

$$\measuredangle(\mathbf{s}_\lambda, \mathbf{n}) \geq 90°.$$

Let

$$S_\lambda = \left\{ \mathbf{n}° \in \Re^3 \,\middle|\, \mathbf{n}°^T \mathbf{s}_\lambda \leq 0 \ \wedge \mathbf{n}° \in V \right\}$$

be the set of all visible surface normals that are unilluminated by light source λ (the self-shadow set) and let

$$\partial S_\lambda = \left\{ \mathbf{n}° \in \Re^3 \,\middle|\, \mathbf{n}°^T \mathbf{s}_\lambda = 0 \ \wedge \mathbf{n}° \in V \right\}$$

be the boundary of S_λ. A cast-shadow is produced if no light reaches the object point although $\measuredangle(\mathbf{s}_\lambda, \mathbf{n}) < 90°$ holds because the light rays are blocked by an object. Let

$$S_{c\lambda} = \left\{ \mathbf{n}° \in \Re^3 \,\middle|\, \mathbf{n}°^T \mathbf{s}_\lambda > 0 \ \wedge \mathbf{n}° \in V \right\}$$

be the set of all visible surface normals that are potentially in a cast-shadow of light source λ (the potential cast-shadow set). An example showing a simple two-dimensional surface with self-shadow and cast-shadow produced by light source k is depicted in Fig. 1. Light source i does not produce any shadows, since we assume that $\mathbf{v}$ and $\mathbf{s}_k$ are parallel. The intervals U_{sk} and U_{ck} describe the self-shadow region and the cast-shadow region, respectively.

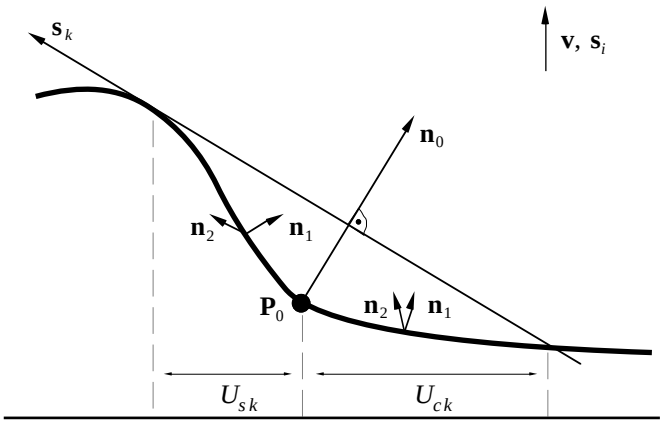

Figure 1: A self-shadow region U_{sk} and a cast-shadow region U_{ck} on a two-dimensional surface and two consistent surface normals pairs $(\mathbf{n}_1, \mathbf{n}_2)$

Let us assume that the light source k causes a shadow on the object surface. Using the measured pair of positive and normalized image irradiances (E_i, E_j) for a considered surface point **P** a pair of consistent surface normals $(\mathbf{n}_1, \mathbf{n}_2)$, with

$$\mathbf{n}_{1,2} = u_{i,j} \cdot \mathbf{s}_i + u_{j,i} \cdot \mathbf{s}_j \pm w \cdot \mathbf{s}_i \times \mathbf{s}_j$$

can be derived locally (pointwise) [1], hence the scalars $u_{i,j}$, $u_{j,i}$ and w can be calculated from the pair of normalized image irradiances (E_i, E_j) and the illumination directions $\mathbf{s}_i, \mathbf{s}_j$. "×" denotes the cross product.

2.1 Resolving the Ambiguity

The ambiguity in the locally recoverable surface normals is implicitly resolved for both shadow types and any configuration of illumination directions if the vectors $\mathbf{n}$, $\mathbf{s}_i$ and $\mathbf{s}_j$ are coplanar. However this only holds for a small subset of visible surface normals. Now we show how to overcome this problem.

At first we assume that the illumination direction $\mathbf{s}_k$ of the shadow-making light source k is oriented perpendicular to illumination directions $\mathbf{s}_i$ and $\mathbf{s}_j$, hence $\mathbf{s}_k = \mathbf{s}_i \otimes \mathbf{s}_j$. The special cross product denoted by "⊗" for two three-dimensional vectors **a**, **b** and the view direction **v** is defined as follows:

$$\mathbf{a} \otimes \mathbf{b} = \begin{cases} \mathbf{a} \times \mathbf{b} & , \text{ if } (\mathbf{a} \times \mathbf{b})^T \mathbf{v} \geq 0 \\ -\mathbf{a} \times \mathbf{b} & , \text{ otherwise.} \end{cases}$$

This definition assures that the resulting vector of the cross product is always visible or is perpendicular to the view direction.

Assuming this illumination configuration the pair of consistent surface normals $(\mathbf{n}_1, \mathbf{n}_2)$ has the property that one surface normal is always illuminated and the other one is unilluminated with respect to light source k, see Fig. 2. If self-shadows and cast-shadows can be distinguished the ambiguity in the surface normals is locally resolvable. In self-shadow regions the surface normal $\mathbf{n} \in (\mathbf{n}_1, \mathbf{n}_2)$ is the right candidate if $90° \leq \measuredangle(\mathbf{s}_k, \mathbf{n}) \leq 180°$ holds. In cast-shadow regions $\mathbf{n} \in (\mathbf{n}_1, \mathbf{n}_2)$ is the correct result if **n** is illuminated with respect to light source k.

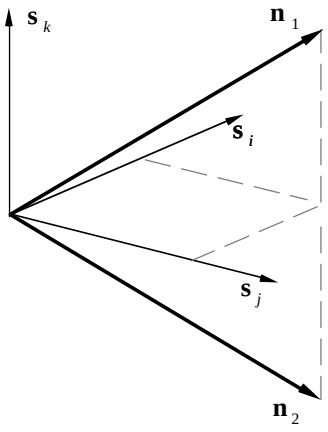

Figure 2: Illumination configuration with $\mathbf{s}_k = \mathbf{s}_i \otimes \mathbf{s}_j$ and a pair of consistent surface normals $(\mathbf{n}_1, \mathbf{n}_2)$

This property should hold for all light sources. It follows that the illumination directions have to be pairwise orthogonal. This configuration is equivalent to the Standard Illumination Configuration SIC (see Appendix) with SIC-angle $\alpha = arctan(\sqrt{2})$.

Employing the trigonometry of spherical triangles it can be shown that about 84% of the visible surface normals in self-shadow regions are uniquely recoverable for $\alpha = arctan(\sqrt{2})$. If the additional information inherent in the irradiance triples is not taken into account then only 25% of the visible surface normals can be recovered (a quarter of the visible hemisphere of the Gaussian sphere).

It can also be shown that the ambiguity is still resolvable under less restrictive assumptions on the illumination configurations. From the combinatorial point of view three cases can occur for the considered surface point **P**.

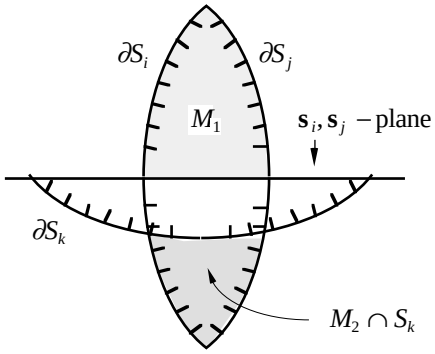

Figure 3: Example of a resolvable ambiguity in surface normals

I Both consistent surface normals of the pair $(\mathbf{n}_1, \mathbf{n}_2)$ are not elements of S_k.

II Exactly one surface normal of the pair $(\mathbf{n}_1, \mathbf{n}_2)$ is an element of S_k.

III Both consistent surface normals of the pair $(\mathbf{n}_1, \mathbf{n}_2)$ are elements of S_k.

Case (I) cannot occur since we have assumed that the surface point **P** is an element of a self-shadow region. In case (II) the ambiguity can be resolved, see example in Fig. 3. The set M_1 is defined by

$$M_1 = S_{ci} \cap S_{cj} \cap \left\{ \mathbf{n}^\circ \in \Re^3 \,\middle|\, \mathbf{n}^{\circ T} (\mathbf{s}_i \otimes \mathbf{s}_j) > 0 \right\}$$

and M_2 is defined by

$$M_2 = S_{ci} \cap S_{cj} \cap \left\{ \mathbf{n}^\circ \in \Re^3 \,\middle|\, \mathbf{n}^{\circ T} (\mathbf{s}_i \otimes \mathbf{s}_j) < 0 \right\}.$$

Case (III) means that both surface normals are unilluminated with respect to light source *k*. Fig. 4 illustrates an example for case (III).

Here, the ambiguity cannot be resolved. Case (III) can only occur if

$$M_1 \cap S_k \neq \varnothing \;\wedge\; M_2 \cap S_k \neq \varnothing.$$

It can be inferred that for the set of illumination configurations

$$\mathbf{s}_k = \omega \cdot \mathbf{s}_i + \mu \cdot \mathbf{s}_j + \nu \cdot \mathbf{s}_i \otimes \mathbf{s}_j \;,$$
$$\text{with } \omega, \mu \geq 0 \text{ and } \nu > 0$$

case (III) can never occur. If this condition should hold then for each of the three light sources the

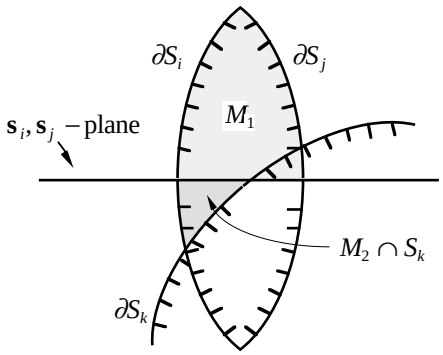

Figure 4: Example of an unresolvable ambiguity in surface normals

possible illumination configurations are given by the following set of inequalities:

$$\measuredangle(\mathbf{s}_i, \mathbf{s}_j) \leq 90° \;\wedge\; \measuredangle(\mathbf{s}_j, \mathbf{s}_k) \leq 90°$$
$$\wedge\; \measuredangle(\mathbf{s}_i, \mathbf{s}_k) \leq 90°.$$

A complementary analysis can be done for cast-shadows.

3. Quantified Improvement

Now we show the quantified improvement if the additional information inherent in the shadows is taken into account. The quantity of surface normals is represented by areas on the Gaussian sphere. Fig. 5 shows the ratios in per cent of recoverable surface normals (by using the shadow information) to the whole set of visible surface normals on the Gaussian sphere for different SIC-angles α with respect to self-shadows. The squares depict the ratio of recoverable areas if the shadows are neglected.

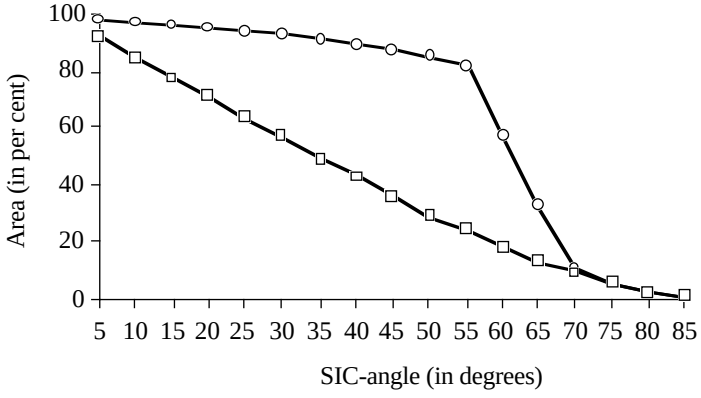

Figure 5: Areas on the Gaussian sphere where surface orientations can be recovered without analyzing the shadows (squares) and with analyzing self-shadows (circles)

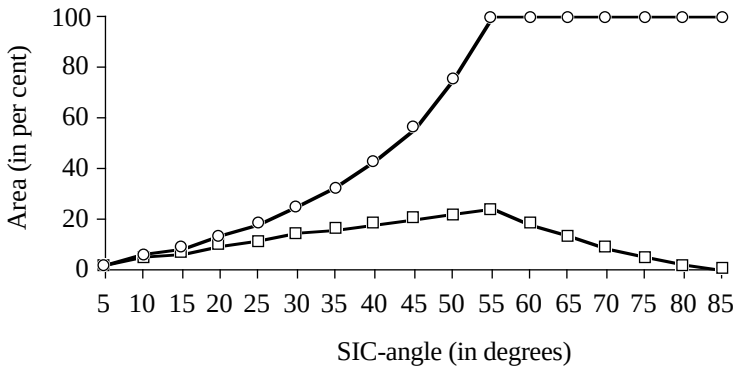

Figure 6: Areas of regions on the Gaussian sphere where surface normals can be recovered using the cast-shadow analysis (see text)

The larger the angles between the illumination directions the smaller the set of recoverable surface normals. The percentages of recoverable surface orientations when analyzing shadows are shown as circles. The difference between the curves is high.

Analogous results for cast-shadows are shown in Fig. 6. The recoverable areas are given with respect to the area of visible surface normals (squares) and with respect to surface normals that are potentially part of cast-shadows (circles). This distinction is necessary because cast-shadows depend not only on the surface orientation but depend also on the position of the surface point. Without making assumptions on the surface geometry it is not possible to give specific data.

4. Experiments

We have performed experiments with synthetic and real images. Here we depict the results for a real Beethoven plaster statue. Fig. 7 shows two of three input images for the 3D shape recovery. The light source in the right image causes large shadow regions on the statue.

Fig. 8 shows the preliminary result of the recovered surfaces of the Beethoven statue with (left) and without (right) making use of the shadow information. As it can be seen from the figure the effect of shadow analysis is quite marked. Remaining errors come from insufficient shadow type classification, and pixels containing shadows from two light sources.

Figure 7: Two irradiance images of a real Beethoven plaster statue, both light sources generate self-shadows and cast-shadows in the images

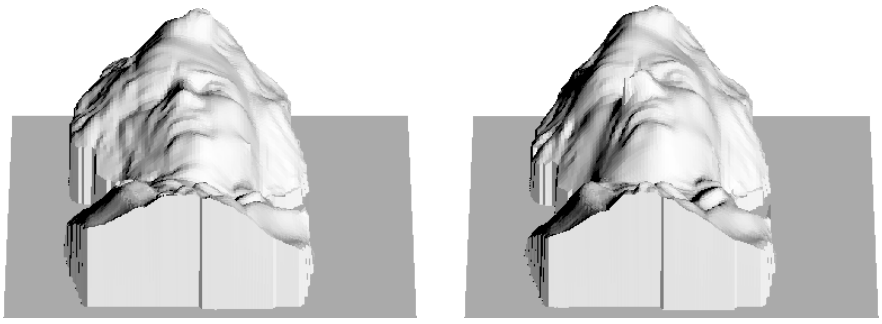

Figure 8: Recovered surfaces of the Beethoven statue with (left) and without (right) making use of the shadow information

5. Conclusions and Future Work

We have analyzed the effect of a shadow analysis on shading based shape recovery. It can be shown that for commonly employed illumination configurations the self-shadow problem can be solved uniquely. We have given first results of a preliminary implementation for shadow classification and shadow analysis. In future work we will implement more robust shadow classification algorithms to better distinguish between the shadow types.

Appendix

A set of illumination directions $\mathbf{s}_1, \mathbf{s}_2$ and $\mathbf{s}_3$ is called a Standard Illumination Configuration (SIC) if the angles between each pair of these illumination directions are equal and the angles between the viewing direction $\mathbf{v}$ and each of $\mathbf{s}_1, \mathbf{s}_2$ and $\mathbf{s}_3$ are equal, too.

We call the angle between the illumination directions and $\mathbf{v}$ the *SIC-angle* α, with

$$\alpha \in (0°, 90°).$$

The set of gradients of the illumination directions $\mathbf{s}_1, \mathbf{s}_2$ and $\mathbf{s}_3$ is

$$\mathbf{G}(\mathbf{s}_1) = (\tfrac{\sqrt{3}}{2}, \quad \tfrac{-1}{2}) \cdot tan(\alpha) \ ,$$

$$\mathbf{G}(\mathbf{s}_2) = (\tfrac{-\sqrt{3}}{2}, \quad \tfrac{-1}{2}) \cdot tan(\alpha) \ \text{ and}$$

$$\mathbf{G}(\mathbf{s}_3) = (0, \ 1) \cdot tan(\alpha) \ .$$

In the discussed analysis we consider SICs with

$$\alpha \leq arctan(\sqrt{2}) \approx 54.74° .$$

In applied work it is not useful to take larger SIC-angles, since the area of the object which is illuminated by all three light sources would be too small.

6. References

[1] B.K.P. Horn: Robot Vision, The MIT Press, Cambridge, 1986.

[2] R. Klette, A. Koschan und K. Schlüns: Computer Vision: Räumliche Information aus digitalen Bildern (in German), Vieweg Verlag, Braunschweig, 1996.

[3] K. Schlüns und M. Teschner: Fast Separation of Reflection Components and its Application in 3D Shape Recovery, Proc. 3rd Color Imaging Conference, Scottsdale, Arizona, 7.-10.11. 95, pp. 48-51, 1995.

[4] K. Shinmoto, S. Kaneko und T. Honda: A Method for Determining 3-D Surface Orientation of Objects with Textures Based on Photometric Stereo, Proc. 10th Int. Conf. on Computer-Aided Engineering, Palermo, Italy, pp. 445-453, 1994.

[5] F. Solomon und K. Ikeuchi: Extracting the Shape and Roughness of Specular Lobe Objects Using Four Light Photometric Stereo, IEEE Trans. on PAMI, Vol. 18, No. 4, pp. 449-454, 1996.

[6] R.J. Woodham: Photometric Method for Determining Surface Orientations from Multiple Images, Optical Engineering, Vol. 19, No. 1, pp. 139-144, 1980.

[7] R.J. Woodham: Gradient and curvature from the photometric-stereo method, including local confidence estimation, J. Opt. Soc. Am. A, Vol. 11, No. 11, pp. 3050-3068, 1994.