

Curvature Weighted Evidence Combination for Shape-from-Shading

Fabio Sartori and Edwin R. Hancock

Department of Computer Science
University of York, York YO10 5DD, UK

Abstract. This paper describes a new curvature consistency method for shape-from-shading. Our idea is to combine evidence for the best surface normal direction. To do this we transport surface normals across the surface using a local estimate of the Hessian matrix. The evidence combination process uses the normal curvature to compute a weighted average surface normal direction. We experiment with the resulting shape-from-shading method on a variety of real world imagery.

1 Introduction

The recovery of surface shape from shading patterns is an under constrained problem. Various constraints, including surface smoothness and occluding boundary positions can be used to render the process computationally tractable. In principal, curvature consistency constraints may also improve the quality of the recovered surface information. Despite proving alluring, there have been relatively few successful uses of curvature information in shape-from-shading, and these have overlooked much of the richness of the differential surface structure. For instance, Ferrie and Lagarde [4] have used the Darboux frame consistency method of Sander and Zucker [10]. This method uses a least-squares error criterion to measure the consistency of the directions of neighbouring surface normals and principal curvature directions. This method is applied as a postprocessing step and does not directly constrain the solution of the image irradiance equation. Worthington and Hancock have used method which enforces improved compliance with the image irradiance equation [11]. They constrain the surface normals to fall on the irradiance cone defined by Lambert's law. Subject to this constraint, the field of surface normals is smoothed using a method suggested by robust statistics. The degree of smoothing employed depends on the variance of the Koenderink and Van Doorn shape-index [8]. This is a scalar quantity which measures topography, but is not sensitive to finer differential structure, such as the directions of principal curvature.

The observation underpinning this paper is that although there is a great deal of information residing in the local Darboux frames, there has been relatively little effort devoted to exploiting this in shape-from-shading. In particular, we aim to use the field of principal curvature directions as a source of constraints that can be used to improve the shape-form-shading process. Our approach is

an evidence combining one. As suggested by Worthington and Hancock [11], we commence with the surface normals positioned on their local irradiance cone and aligned in the direction of the local image gradient. From the initial surface normals, we make local estimates of the Hessian matrix. Using the Hessian, we transport neighbouring normals across the surface. This allows us to accumulate a sample of surface normals at each location. By weighting this sample by the normal curvature of the transport path and the resulting brightness error, we compute a revised surface normal direction.

2 Shape-from-Shading

Central to shape-from-shading is the idea that local regions in an image $E(x, y)$ correspond to illuminated patches of a piecewise continuous surface, $z(x, y)$. The measured brightness $E(x, y)$ will depend on the material properties of the surface, the orientation of the surface at the coordinates-ordinates (x, y) , and the direction and strength of illumination.

The *reflectance map*, $R(p, q)$ characterises these properties, and provides an explicit connection between the image and the surface orientation. Surface orientation is described by the components of the surface gradient in the x and y direction, i.e. $p = \frac{\partial z}{\partial x}$ and $q = \frac{\partial z}{\partial y}$. The shape from shading problem is to recover the surface $z(x, y)$ from the intensity image $E(x, y)$. As an intermediate step, we may recover the needle-map, or set of estimated local surface normals, $\mathbf{Q}(x, y)$.

Needle-map recovery from a single intensity image is an under-determined problem [1,2,7,9] which requires a number of constraints and assumptions to be made. The common assumptions are that the surface has ideal Lambertian reflectance, constant albedo, and is illuminated by a single point source at infinity. A further assumption is that there are no inter-reflections, i.e. the light reflected by one portion of the surface does not impinge on any other part.

The local surface normal may be written as $\mathbf{Q} = (-p, -q, 1)^T$, where $p = \frac{\partial z}{\partial x}$ and $q = \frac{\partial z}{\partial y}$. For a light source at infinity, we can similarly write the light source direction as $\mathbf{s} = (-p_l, -q_l, 1)^T$. If the surface is Lambertian the reflectance map is given by $R(p, q) = \mathbf{Q} \cdot \mathbf{s}$. The image irradiance equation [6] states that the measured brightness of the image is proportional to the radiance at the corresponding point on the surface; that is, just the value of $R(p, q)$ for p, q corresponding to the orientation of the surface. Normalising both the image intensity, $E(x, y)$, and the reflectance map, the constant of proportionality becomes unity, and the image irradiance equation is simply $E(x, y) = R(p, q)$. Lambert's equation provides insufficient information to uniquely determine the surface normal direction.

3 Differential Surface Structure

In this paper we are interested in using curvature consistency information to constrain the recovery of shape-from-shading. Our characterisation of the differential structure of the surface is based on the Hessian matrix which can be

computed from the currently available field of surface normals, or Gauss-map, in the following manner

$$\mathcal{H}_o = \begin{bmatrix} \frac{\partial}{\partial x}(\mathbf{Q}_o)_x & \frac{\partial}{\partial x}(\mathbf{Q}_o)_y \\ \frac{\partial}{\partial y}(\mathbf{Q}_o)_x & \frac{\partial}{\partial y}(\mathbf{Q}_o)_y \end{bmatrix} = \begin{bmatrix} h_{11} & h_{12} \\ h_{21} & h_{22} \end{bmatrix} \quad (1)$$

where the directional derivatives are extracted using first differences of the surface meshes in the approximation of the pixel lattice. The two eigenvalues of the Hessian matrix are the maximum and minimum curvatures:

$$K_o^{max} = -\frac{1}{2}(h_{11} + h_{22} - S) \quad , \quad K_o^{min} = -\frac{1}{2}(h_{11} + h_{22} + S) \quad (2)$$

where $S = \sqrt{(h_{11} - h_{22})^2 + 4(h_{21}h_{12})}$. The eigenvector associated with the maximum curvature K_o^M is the principal curvature direction. On the tangent-plane to the surface, the principal curvature direction is given by the 2-component vector

$$\mathbf{e}_o^{max} = \begin{cases} (h_{12}, -\frac{1}{2}(h_{11} - h_{22} + S))^T & h_{11} \geq h_{22} \\ (\frac{1}{2}(h_{11} - h_{22} - S), h_{21})^T & h_{11} < h_{22} \end{cases} \quad (3)$$

In this paper we are interested in using the local estimate of the Hessian matrix to provide curvature consistency constraints for shape from-shading. Our aim is to improve the estimation of surface normal direction by combining evidence from both shading information and local surface curvature. As demonstrated by both Ferrie and Lagarde [4] and Worthington and Hancock [11], the use of curvature information allows the recovery of more consistent surface normal directions. It also provides a way to control the over-smoothing of the resulting needle maps. Ferrie and Lagarde [4] have addressed the problem using local Darboux frame smoothing. Worthington and Hancock [11], on the other hand, have employed a curvature sensitive robust smoothing method. Here we adopt a different approach which uses the equations of parallel transport to guide the prediction of the local surface normal directions.

Suppose that we are positioned at the point $\mathbf{X}_o = (x_o, y_o)^T$ where the current estimate of the Hessian matrix is \mathcal{H}_o . Further suppose that \mathbf{Q}_m is the surface normal at the point $\mathbf{X}_m = (x_m, y_m)^T$ in the neighbourhood of \mathbf{X}_o . We use the local estimate of the Hessian matrix to transport the vector \mathbf{Q}_m to the location \mathbf{X}_o . The first-order approximation to the transported vector is

$$\mathbf{Q}_m^o = \mathbf{Q}_m + \mathcal{H}_o(\mathbf{X}_m - \mathbf{X}_o) \quad (4)$$

This procedure is repeated for each of the surface normals belonging to the neighbourhood R_o of the point o . In this way we generate a sample of alternative surface normal directions at the location o .

We would like to associate with the transported surface normals a measure of certainty based on the curvature of the path $\Gamma_{o,m}$ from the point m to the

point o . The normal curvature at the point o in the direction of the transport path is approximately

$$\kappa_{o,m} = (\mathbf{T}_{o,m} \cdot \mathbf{e}_o^{max})^2 (K_o^{max} - K_o^{min}) + K_o^{min} \quad (5)$$

where $\mathbf{T}_{o,m} = \frac{\mathbf{X}_m - \mathbf{X}_o}{\|\mathbf{X}_m - \mathbf{X}_o\|}$ is the unit vector from o to m .

4 Statistical Framework

We would like to exploit the transported surface-normal vectors to develop an evidence combining approach to shape-from-shading. To do this we require a probabilistic characterisation of the sample of available surface normals. This is a two-component model. Firstly, we model the data-closeness of the transported surface normals with Lambert's law. We assume that the observed brightness E_o at the point \mathbf{X}_o follows a Gaussian distribution. As a result the probability density function for the transported surface normals is

$$p(E_o | \mathbf{Q}_m^o) = \frac{1}{\sqrt{2\pi}\sigma_E} \exp \left[-\frac{(E_o - \mathbf{Q}_m^o \cdot \mathbf{s})^2}{2\sigma_E^2} \right] \quad (6)$$

where σ_E^2 is the noise-variance of the brightness errors. The second model ingredient is a curvature prior. Here we adopt a model in which we assume that the sample of transported surface normals is drawn from a Gaussian prior, which is controlled by the normal curvature of the transport path. Accordingly we write

$$p(\mathbf{Q}_m^o) = \frac{1}{\sqrt{2\pi}\sigma_k} \exp \left[-\frac{1}{2\sigma_k^2} k_{o,m}^2 \right] \quad (7)$$

With these ingredients, the weighted mean for the sample of transported surface normals

$$\hat{\mathbf{Q}}_o = \sum_{m \in R_o} \mathbf{Q}_m^o p(E_o | \mathbf{Q}_m^o) p(\mathbf{Q}_m^o) \quad (8)$$

where R_o is the index set of the surface normals used for the purposes of transport. Substituting for the distributions,

$$\hat{\mathbf{Q}}_o = \frac{\sum_{m \in R_o} \mathbf{Q}_{o,m} \exp \left[-\frac{1}{2} \left(\frac{(E_o - \mathbf{Q}_m^o \cdot \mathbf{s})^2}{\sigma_E^2} + \frac{\kappa_{o,m}^2}{\sigma_k^2} \right) \right]}{\sum_m \exp \left[-\frac{1}{2} \left(\frac{(E_o - \mathbf{Q}_m^o \cdot \mathbf{s})^2}{\sigma_E^2} + \frac{\kappa_{o,m}^2}{\sigma_k^2} \right) \right]} \quad (9)$$

and the predicted brightness is $\hat{E}_o = \hat{\mathbf{Q}}_o \cdot \mathbf{s}$

This procedure is repeated at each location in the field of surface normals. We iterate the method as follows:

- 1: At each location compute a local estimate of the Hessian matrix \mathcal{H}_o from the currently available surface normals \mathbf{Q}_o .

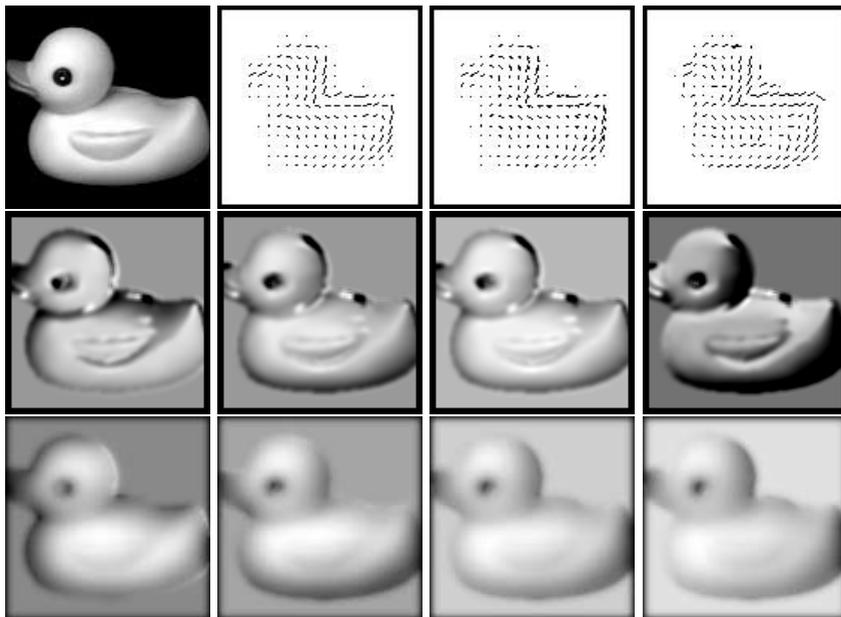


Fig. 1. Toy duck: original image, initial needle-map, final needle-maps and reconstructed images from different illumination directions

- 2: At each image location \mathbf{X}_o obtain a sample of surface normals $N_o = \{\mathbf{Q}_m^o | m \in R_o\}$ by applying parallel transport to the set of neighbouring surface normals whose locations are indexed by the set R_o .
- 3: From the set of surface normals N_o compute the expected brightness value \hat{E}_o and the updated surface normal direction $\hat{\mathbf{Q}}_o$. Note that the measured intensity E_o is kept fixed throughout the iteration process and is not updated.
- 4: With the updated surface normal direction to hand, return to step 1, and recompute the local curvature parameters.

To initialise the surface normal directions, we adopt the method suggested by Worthington and Hancock [11]. This involves placing the surface normals on the irradiance cone whose axis is the light-source direction \mathbf{s} and whose apex angle is $\cos^{-1} E_o$. The position of the surface normal on the cone is such that its projection onto the image plane points in the direction of the local image gradient, computed using the Canny edge detector. When the surface normals are initialised in this way, then they satisfy the image irradiance equation.

5 Experiments

We have experimented with our new shape-form-shading method on a variety of real world imagery and have compared it with the Horn and Brooks method [7].

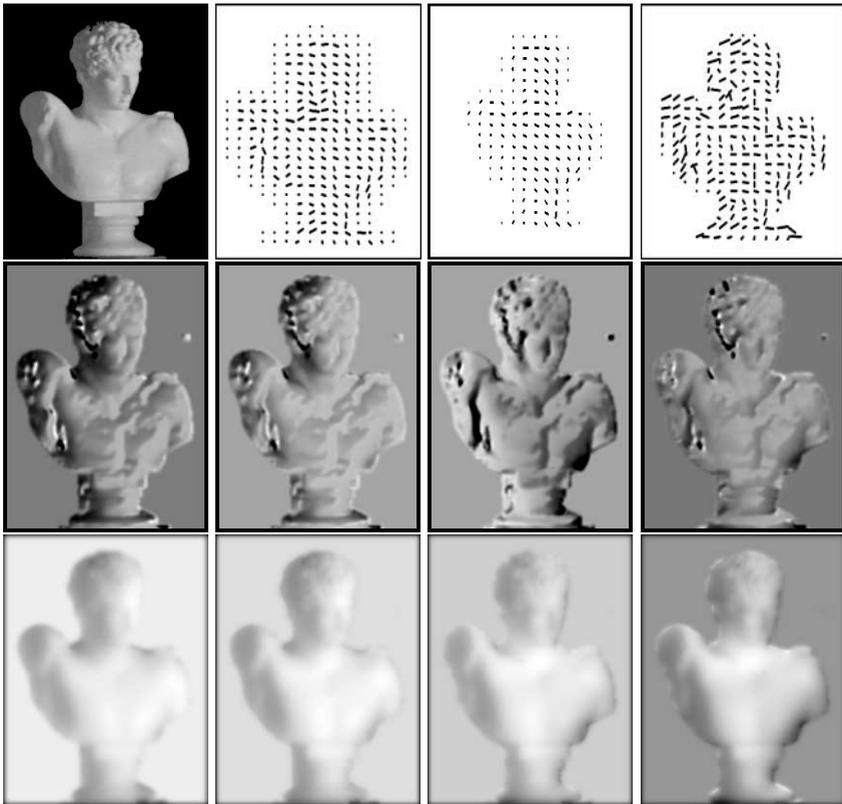


Fig. 2. Hermes: original image, initial needle-map, final needle-maps and reconstructed images from different illumination directions

The images are taken from a variety of sources including the Columbia COIL database and web epositories of images of statues.

In Figure 1 we show a sequence of results for a toy duck image from the Columbia COIL data-base. In the top row, from left to right, we show the original image, the needle-map obtained using the Horn and Brooks method [5] and the needle-map obtained using our new method. The needle-map obtained using our new method is more detailed than that obtained using the Horn and Brooks method. In the subsequent two rows of the figure, we show the results of reilluminating the needle-map from various directions. The reilluminations reveal that the method recovers fine surface detail, especially around the wings and the head. The bottom row of the figure shows the result of reillumination using the Horn and Brooks method. This is blurred and does not convey the surface detail delivered by our new method.

Figures 2, 3 and 4 and repeat this sequence for different images. In Figure 2 we study an image of a statue of Hermes, in Figure 3 an image of a terracotta

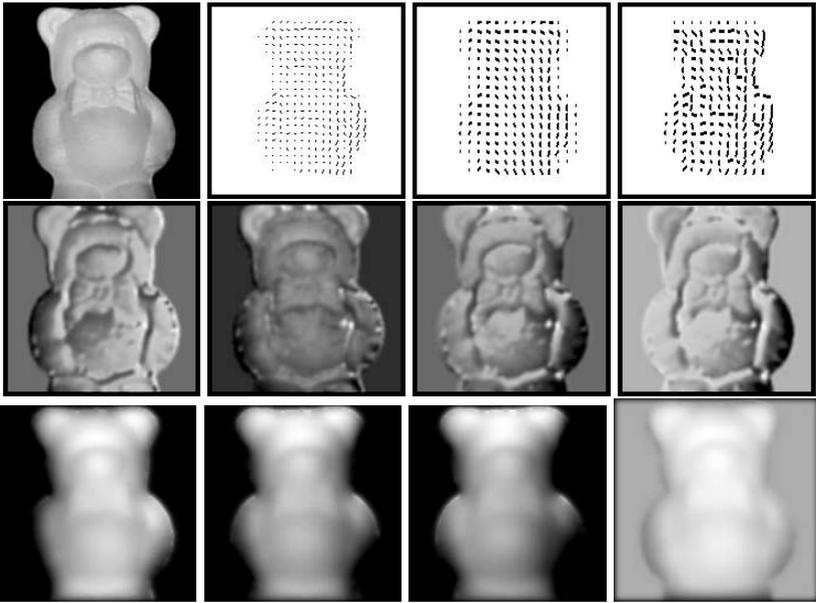


Fig. 3. Bear: Original image, initial needle-map, final needle-maps and reconstructed images from different illumination directions

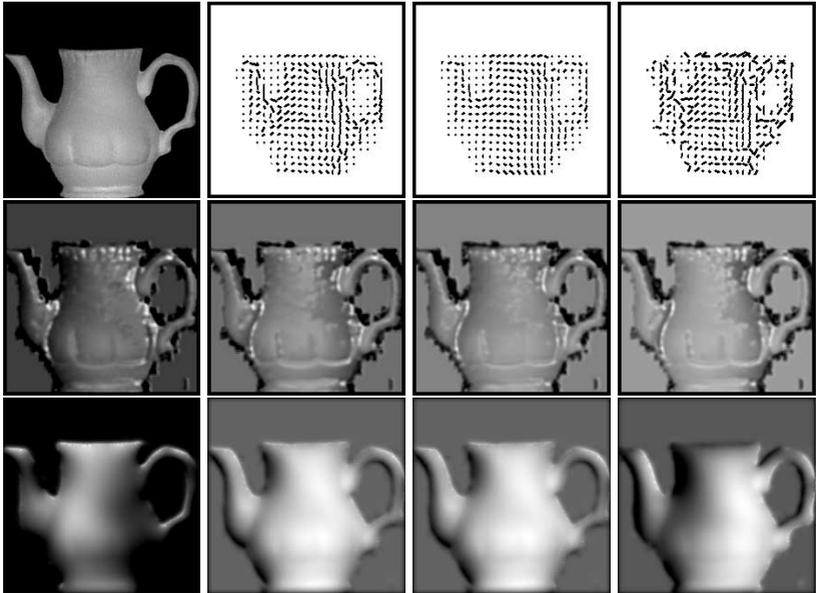


Fig. 4. Pot: Original image, initial needle-map, final needle-maps and reconstructed images from different illumination directions

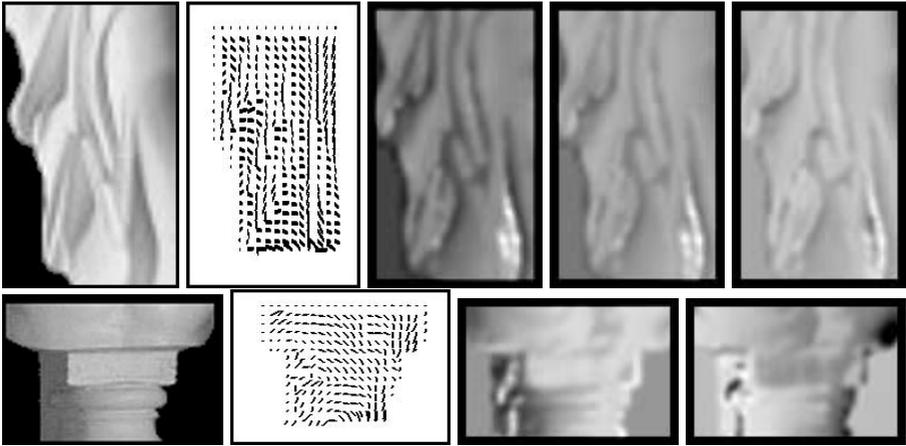


Fig. 5. Details from Venus and Hermes: needle-map and reillumination examples

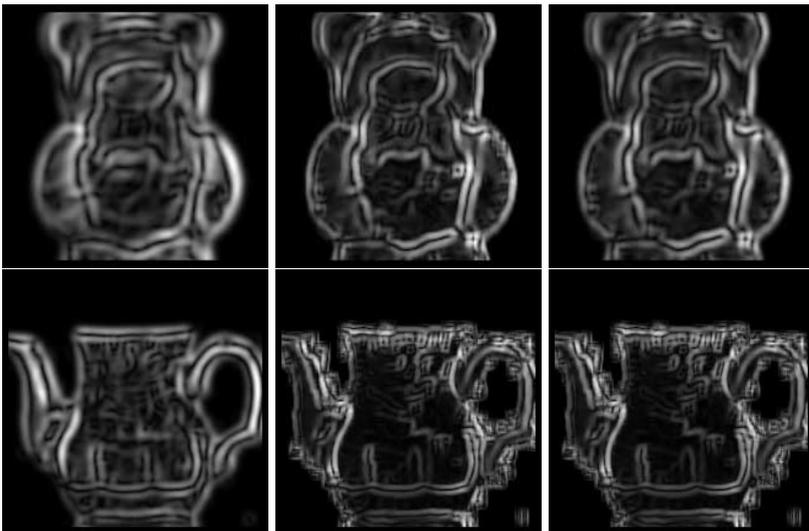


Fig. 6. Evolution of the curvedness during the minimisation process

bear, and in Figure 4 an image of a terracotta tea pot. The detail in the re-illuminations of the needle-maps obtained with our new method are much clearer than those obtained with the Horn and Brooks method.

In Figure 5 we show the results for images of highly structured surfaces. The top row shows the results for a detail in the folds of the drapery for the Venus de Milo statue. The bottom row shows the results for a detail around the plinth of the Hermes statue shown earlier. The images in each row of the figure are from left to right, the original image, the needle-map, and some example

reilluminations. The results show that our method is able to recover quite fine surface detail, including high curvature structure.

Finally, we focus on the iterative qualities of the algorithm. In Figure 6, we show the curvedness $K = \sqrt{(K_o^{max})^2 + (K_o^{min})^2}$ as a function of iteration number. It is clear that the method has the effect of sharpening the curvature detail as it iterates.

6 Conclusions

In this paper we have described a curvature consistency method for shape-from-shading. The idea underpinning this work is to compute a weighted average of linearly transported surface normals. The transport is realized using a local estimate of the Hessian matrix and the weights are computed using the normal curvature of the transport path. The method proves effective on a variety of real world images.

References

1. Belhumeur, P. N. and Kriegman, D. J. (1996) What is the Set of Images of an Object Under All Possible Lighting Conditions? *Proc. IEEE Conference on Computer Vision and Pattern Recognition*, pp. 270-277. [217](#)
2. Bruckstein, A. M. (1988) On Shape from Shading, *CVGIP*, Vol. 44, pp. 139-154. [217](#)
3. Do Carmo M. P. (1976), *Differential Geometry of Curves and Surfaces*, *Prentice Hall*.
4. Ferrie, F. P. and Lagarde, J. (1990) Curvature Consistency Improves Local Shading Analysis, *Proc. IEEE International Conference on Pattern Recognition*, Vol. I, pp. 70-76. [216](#), [218](#)
5. Horn, B. K. P. and Brooks, M. J. (1986) The Variational Approach to Shape from Shading, *CVGIP*, Vol. 33, No. 2, pp. 174-208. [221](#)
6. Horn, B. K. P. and Brooks, M. J. (eds.), *Shape from Shading*, MIT Press, Cambridge, MA, 1989. [217](#)
7. Horn, B. K. P. (1990) Height and Gradient from Shading, *IJCV*, Vol. 5, No. 1, pp. 37-75. [217](#), [220](#)
8. Koenderink, J. J. (1990) *Solid Shape*, MIT Press, Cambridge MA. [216](#)
9. Oliensis, J. and Dupuis, P. (1994) An Optimal Control Formulation and Related Numerical Methods for a Problem in Shape Reconstruction, *Ann. of App. Prob.* Vol. 4, No. 2, pp. 287-346. [217](#)
10. Sander, P. and Zucker, S. (1990) Inferring surface trace and differential structure from 3-d images, *PAMI*, 12(9):833-854. [216](#)
11. Worthington, P. L. and Hancock, E. R. (1999) New Constraints on Data-closeness and Needle-map consistency for SFS, *IEEE Transactions on Pattern Analysis*, Vol. 21, pp. 1250-1267. [216](#), [217](#), [218](#), [220](#)