Curvature Consistency 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 synthetic data.

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 [9]. 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 post-processing 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 [10]. 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 [7]. 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 combing one. As suggested by Worthington and Hancock [10], 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, 6, 2, 8] 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 [5] 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

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

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:

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

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 [3]. On the tangent-plane to the surface, the principal curvature direction is given by the 2-component vector

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

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 [10], 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 [10], 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 $\boldsymbol{X}_o = (x_o, y_o)^T$ 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 $\boldsymbol{X}_m = (x_m, y_m)^T$ in the neighbourhood of $\boldsymbol{X}_o$. We use the local estimate of the Hessian matrix to transport the vector $\mathbf{Q}_m$ to the location $\boldsymbol{X}_o$. The first-order approximation to the transported vector is

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

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

$$\kappa_{o,m} = (\boldsymbol{T}_{o,m}.\boldsymbol{e}_o^{max})^2(K_o^{max} - K_o^{min}) + K_o^{min} \tag{5}$$

where $\boldsymbol{T}_{o.m} = \frac{\boldsymbol{X}_m - \boldsymbol{X}_o}{|\boldsymbol{X}_m - \boldsymbol{X}_o|}$ 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 $\boldsymbol{X}_o$ follows a Gaussian distribution. As a result the probability density function for the transported surface normals is

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

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

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

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

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

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

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

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.

5   Experiments

We present some experimental evaluation of the new method by showing results on synthetic data with known ground truth and on real world.

Turning our attention first to synthetic data, we have generated height data from which surface normal directions have been computed. The surface normals have been used to generate shading information using a Lambertian reflectance model. We have applied our shape-from-shading algorithm to the images generated in this way. The objects studied are shown in Figure 1. These are a sphere surrounded by a torus and two domes intersected by a parabolic dome. The figure shows the synthetic images of the surfaces together with their ground-truth needle-maps. In Figures 2 and 3 we show some examples of Lambertian
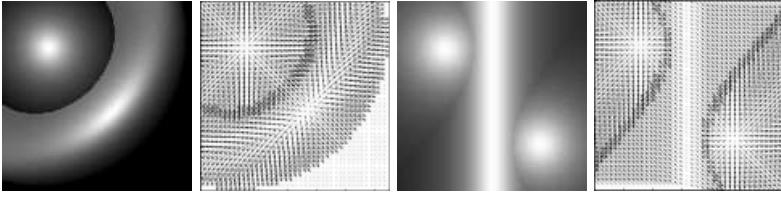
Fig. 1. Original synthetic images and respective needle-maps.

re-illuminations computed from the recovered surface normals delivered by our new method and the method of Worthington and Hancock [10]. The main feature to note is that re-illuminations generated by our new method preserve more of the fine surface structure. Next, we present a series of plots which compare the results obtained using shape-from-shading with the corresponding ground truth for the synthetic images. In Figure 4, from left to right, we show the needle-maps obtained using the Worthington and Hancock method, the method described in this paper and the field of vector differences between these two needle-maps. It is clear that the main differences between the two methods occur where the surface is most curved. Figures 5 and 6 show scatter plots of estimated versus ground truth shape-index $\phi_o = \frac{2}{\pi} \arctan \frac{K_o^{max} + K_o^{min}}{K_o^{max} - K_o^{min}}$. This is an angular measure of local surface topography, which has proved effective in the visualisation of surfaces [7]. The first column of scatter plots is for the initial needle map, and this is common to both the vector transport method and the Worthington and Hancock method. The second column is for the middle iteration, and the final column is for the last iteration. The top row is for our new method, and the bottom row is for the Worthington and Hancock method. Figure 5 shows the results for the torus+sphere and Figure 6 shows the results for the oval domes. The shape-index distributions in the initial scatter plots are dispersed about the diagonal. However, in the final iteration the dispersion is considerably reduced. It is clear that the Worthington and Hancock method introduces some systematic biases as it uses the shape-index to control smoothing. Our method does not appear to suffer from this problem and has dispersed but unstructured scatter plot. It is also important to note that our method typically converges in 10 iterations, while WH method takes some 50 iterations.

Finally we show the results for real world images. The result in Figure 7 is for the "Three Graces" relief. The original image is in the left-most panel and a sequence of re-illuminations generated from the surface normals delivered using our new method are shown in the remaining panels. The surface detail is well reconstructed and the re-illumination captures well the changes in light source direction. Figure 8 repeats this sequence for an image of the head of Canova. The re-illuminations reveal important surface details including the fine structure of the hair, the indentations of the left cheek and the detail of the nose. The results show that our method is able to recover quite fine surface detail, including high curvature structure.
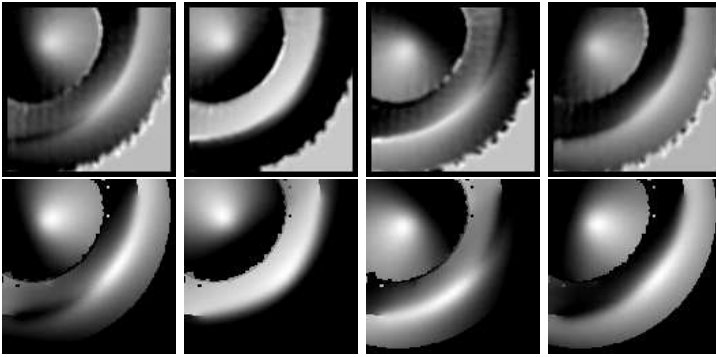
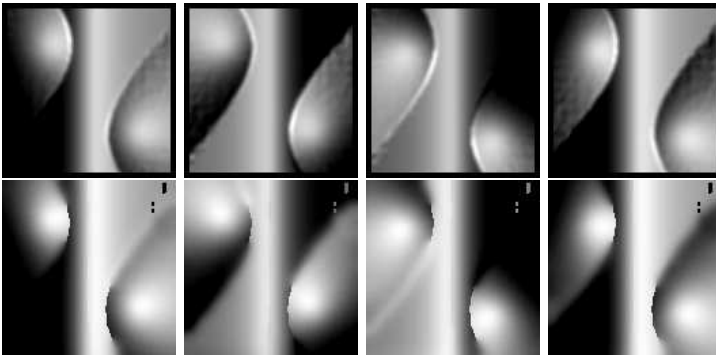**Fig. 2.** Needle-maps re-illuminations from the new method (first row) and Worthington and Hancock method (second row).

Fig. 3. Needle-maps re-illuminations from the new method (first row) and Worthington and Hancock method (second row).
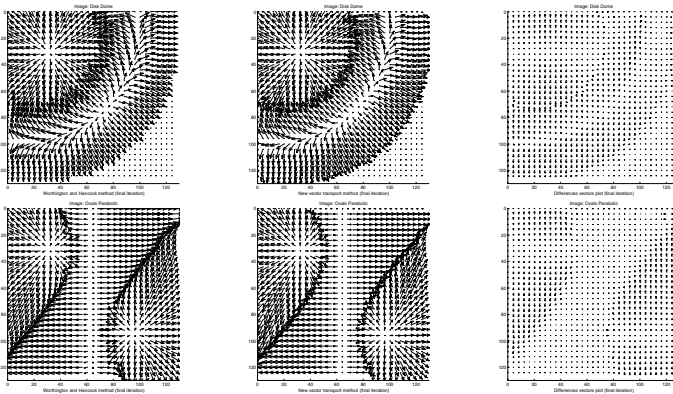

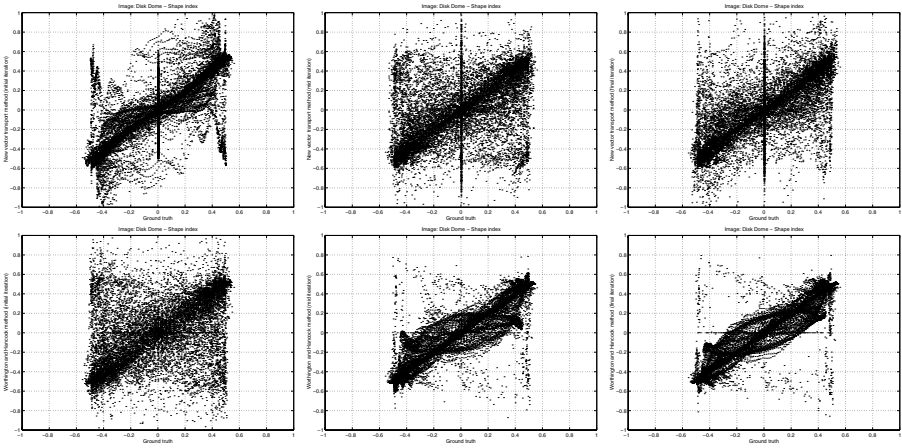**Fig. 4.** Needle-maps delivered by the two different methods.

Fig. 5. Scatterplots of the shape index for the sphere and torus image. New method (first row) and Worthington and Hancock method (second row).
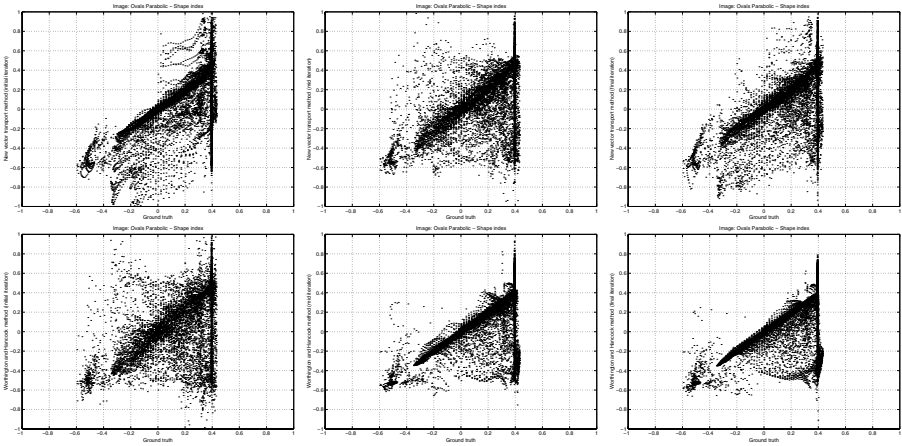

Fig. 6. Scatterplots of the shape index for the domes image. New method (first row) and Worthington and Hancock method (second row).

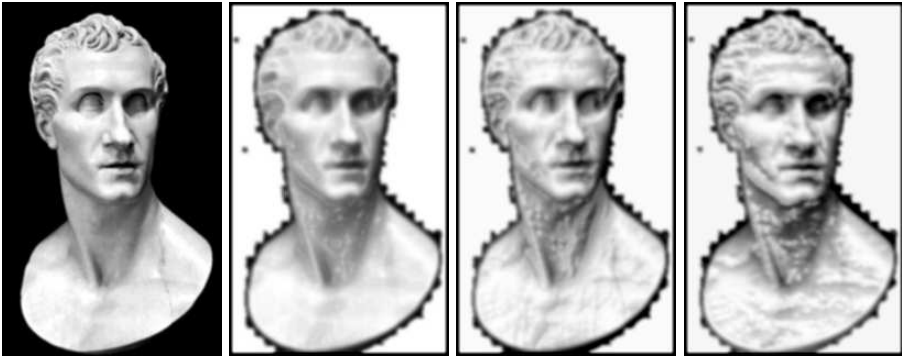**Fig. 7.** Re-illuminations from the new method.

Fig. 8. Re-illuminations from new method.

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 synthetic and 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.
2. Bruckstein, A.M. (1988) On Shape from Shading, *CVGIP,* Vol. 44, pp. 139-154.
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.
5. Horn, B.K.P. and Brooks, M.J. (eds.), *Shape from Shading,* MIT Press, Cambridge, MA, 1989.
6. Horn, B.K.P. (1990) Height and Gradient from Shading, *IJCV,* Vol. 5, No. 1, pp. 37-75.
7. Koenderink, J.J. (1990) *Solid Shape,* MIT Press, Cambridge MA.
8. 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.
9. Sander, P. and Zucker, S. (1990) Inferring surface trace and differential structure from 3-d images, *PAMI,* 12(9):833–854.
10. 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.