An Evidence Combining Approach to Shape-from-shading

Eraldo Ribeiro, Fabio Sartori and Edwin R. Hancock,
Department of Computer Science, University of York, York Y01 5DD, UK.

Abstract

In this paper we describe a statistical framework for shape-from-shading. First, we commence by making least squares estimates of local Hessian matrices using samples of surface normals. These estimated Hessians are used to perform parallel transport of the surface normals across the surface. By transporting neighbouring surface normals in this way we are able collect samples of votes for the surface normal direction at each location. We select between these alternatives on the basis of compliance with the image irradiance equation.

1. Introduction

Shape-from-shading is a problem that has been studied for over 25 years in the vision literature [1, 5, 7, 8]. Stated succinctly, the problem is to recover local surface orientation information, and hence reconstruct the surface height function, from information provided by the surface brightness. Since the problem is an ill-posed one, in order to be rendered tractable recourse must be made to strong simplifying assumptions and constraints. Hence, the process is usually specialised to matte reflectance from a surface of constant albedo, illuminated by a single collimated light source of known direction. To overcome the problem that the two parameters of surface slope can not be recovered from a single brightness measurement, the process is augmented by constraints on surface normal direction at occluding contours or singular points, and also by constraints on surface smoothness.

The observation underpinning this paper is that although considerable effort has gone into the development of improved shape-from-shading methods, there are two areas which leave scope for further development. The first of these is the use of statistical methods in the recovery of surface normal information. The second is that relatively little effort has been expended in the use of ideas from differential geometry for surface modeling.

Our aim in this paper is to develop a sample-based algorithm for shape-from-shading which exploits curvature consistency information. As suggested by Worthington and Hancock [8], we commence with the surface normals positioned on their local irradiance cone and are aligned in the direction of the local image gradient. ¿From the initial surface normals, we make local estimates of the Hessian matrix. This allows us to transport neighbouring normals across the surface in a manner which is consistent with the local surface topography. The resulting sample of surface normals represent predictions of the local surface orientation which are consistent with the local surface curvature. Moreover, each transported vector can be used to make a prediction of the local image image brightness. We select from these alternatives the one which gives the lowest brightness error.

Hence, we facilitate a direct coupling between consistent surface normal estimation and reconstruction of the image brightness. Moreover, our method overcomes the problem of estimating surface normal directions in a natural way. This offers two advantages over existing methods for shape-from-shading. First, because it is evidence-based, unlike the Horn and Brooks method, it is not model dominated and does not oversmooth the recovered field of surface normal directions. The data-closeness and surface-smoothness errors are not simply compounded in an additive way as is the case in the regularisation method. Second, and unlike the Worthington and Hancock method, it relaxes the image irradiance equation and hence allows for brightness errors to be corrected. Another interesting property of the method, is that we parameterise the local surface structure using the Hessian matrix, rather than quadric patch parameters. Hence we exploit the intrinsic geometry of the Gauss map rather than its extrinsic geometry.

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 co-ordinates (x, y), and the direction and strength of illumination.

The *reflectance map*, $R(p,q)$ characterizes 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 [7] 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 [3] 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. Normalizing both the image intensity, $E(x,y)$, and the reflectance map, the constant of proportionality becomes unity, and the image irradiance equation is simply $E(x,y) = R(p,q)$.

Although the image irradiance equation succinctly describes the mapping between the x,y co-ordinate space of the image and the p,q gradient-space of the surface, it provides insufficient constraints for the unique recovery of the needle-map. To overcome this problem, a further constraint must be applied. Usually, the needle-map is assumed to vary smoothly.

3. Differential Surface Structure

In this paper we are interested in the local differential structure of surfaces represented in terms of a field of surface geometry. In the differential geometry this representation is known as the Gauss map. The differential structure of the surface is captured by the second fundamental form or Hessian matrix

$$\mathcal{H} = \begin{pmatrix} \alpha & \beta \\ \beta & \gamma \end{pmatrix} \tag{1}$$

where $\alpha = \left(\frac{\partial \mathbf{Q}}{\partial x}\right)_x$, $\gamma = \left(\frac{\partial \mathbf{Q}}{\partial x}\right)_y$ and $\beta = \left(\frac{\partial \mathbf{Q}}{\partial y}\right)_x = \left(\frac{\partial \mathbf{Q}}{\partial x}\right)_y$. We make least squares estimates of the elements of the Hessian using samples of surface normals. Let $\mathbf{Q}_o$ represent the surface normal at the position (x_o,y_o) and let $\mathbf{Q}_m$ be a neighbouring surface normal with position (x_m,y_m). Suppose that $\Delta_m = Q_m - Q_o$ is the difference in surface normal direction at the locations o and m, and that $\Delta x_m = x_m - x_o$ and $\Delta y_m = y_m - y_o$ are the co-ordinate displacements between the locations. Let $\mathbf{N} = ((\Delta Q_1)_x, (\Delta Q_1)_y, (\Delta Q_2)_x, ...)$ and

$$\mathbf{X} = \begin{pmatrix} \Delta x_1 & \Delta y_1 & 0 \\ 0 & \Delta x_1 & \Delta y_1 \\ \Delta x_2 & \Delta y_2 & 0 \\ & \vdots & \end{pmatrix}$$

By performing a first-order Taylor expansion, on the directional derivatives of the surface normals, it is a straightforward task to show that the least-squares estimator of the elements of the Hessian is

$$\Phi_o = (\alpha_o, \beta_o, \gamma_o)^T = (\mathbf{X^T X})^{-1}\mathbf{X^T N} \tag{2}$$

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 [2] and Worthington and Hancock [8], 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 [2] have addressed the problem using local Darboux frame smoothing. Worthington and Hancock [8], 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.

Our idea is as follows. At each location on the surface we make an estimate of the vector of curvature parameters. Suppose that we are positioned at the point $\vec{X}_o = (x_o,y_o)^T$ where the vector of estimated curvature parameters is Φ_o and that the resulting estimate of the Hessian matrix is $\mathcal{H}_o$. Further suppose that $\mathbf{Q}_m$ is the surface normal at the point $\vec{X}_m = (x_m,y_m)^T$ in the neighbourhood of $\vec{X}_o$. We use the local curvature parameters Φ_o to transport the vector $\mathbf{Q}_m$ to the location $\vec{X}_o$. The first-order approximation to the transported vector is

$$\mathbf{Q}_m^o = \mathbf{Q}_m + \mathcal{H}_o(\vec{X}_m - \vec{X}_o) \tag{3}$$

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. The geometry of the parallel transport procedure is illustrated in Figure 1.

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 characterization of the sample of available surface normals. We assume that the observed brightness E_o at the point $\vec{X}_o$ follows a Gaussian distribution. As a result the probability density function for the transported surface normals is

$$p(E_o|\mathbf{Q}_m, \Phi_o) = \frac{1}{\sqrt{2\pi}\sigma} \exp\left[-\frac{(E_o - \mathbf{Q}_m^o.\mathbf{s})^2}{2\sigma^2}\right] \quad (4)$$

where σ^2 is the noise-variance of the brightness errors. With this distribution to hand, we can use the image irradiance equation to compute the expected value of the image brightness at the location $\vec{X}_o$ for the sample of transported surface normals. The expected brightness is given by

$$\hat{E}_o = \sum_{m \in R_o} p(E_o|\mathbf{Q}_m, \Phi_o)\mathbf{Q}_m^o.\mathbf{s} \quad (5)$$

To update the surface normal direction, we select from the sample the one which results in a brightness value which is closest to $\hat{E}_o$. This surface normal is the one for which

$$\hat{\mathbf{Q}}_o = \arg\min_{m \in R_o} \left[\hat{E}_o - \mathbf{Q}_m^o.\mathbf{s}\right]^2 \quad (6)$$

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 $\mathbf{Q}_o$.
- 2: At each image location $\vec{X}_o$ obtain a sample of surface normals $S_o = \{\mathbf{Q}_m^o | m \in R_o\}$ 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 S_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 [8]. 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 initialized in this way, then they satisfy the image irradiance equation.

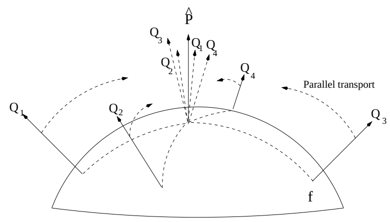

Figure 1. Parallel transport used for predicting the surface normal vector using local curvature estimation.

5. Experiments

In this section, we present some experimental evaluation of the new method.

We commence by exploring some of the iterative properties of the method. Here use experiment with an image of a toy duck from the Columbia COIL data-base. Figure 2 shows the effect of re-illuminating the final needle-map with different light source directions. This highlights the curvature detail on the surface, which appears to be well reconstructed.

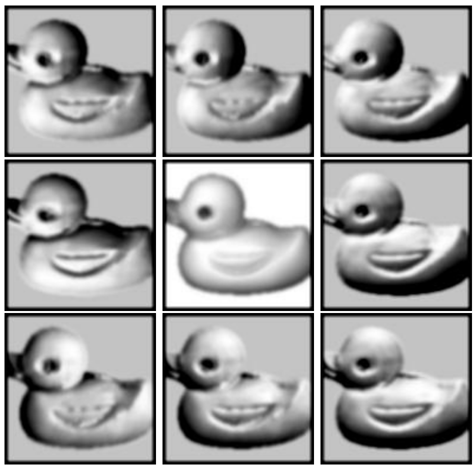

Figure 2. Reconstructed image with different illumination directions

In Figure 3 we show results for images of marble statues. In the left-hand column we show the original image. In the right-hand column we show value of $\hat{E}_o$. In each case the reconstructed brightness images reproduce the curvature structure well, at all but the points of highest curvature.

Finally, in Figure 4, we show the re-illumination of the statue Venus. This captures the surface detail well. In particular, the folds in the draping around the legs is well reproduced.

Figure 3. Results for classical statuary.

6. Conclusions

In this paper we have described a new method for shape-form-shading which relies on vector transport to accumulate evidence for surface normal directions which are consistent with the observed image brightness. The method uses a two-step iterative algorithm. First, estimates of the Hessian matrix are made using the available surface normals. These Hessian matrices are used to perform vector transport on the surrounding surface normals to accumulate a sample of orientation hypotheses. These putative directions are used to compute an expected value for the image brightness. In the second step of the algorithm, the surface normal direction is updated. The direction is taken to be that of the transported vector which yields the brightness which is closest to the expected value. The method is evaluated on a variety of real-world images where it provides promising results.

Figure 4. Re-illumination of Venus

References

[1] Bichsel, M. and Pentland, A.P. (1992) A Simple Algorithm for Shape from Shading, *Proc. IEEE Conf. on Computer Vision and Pattern Recognition,* pp. 459-465.

[2] 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.

[3] Horn, B.K.P. and Brooks, M.J. (eds.), *Shape from Shading,* MIT Press, Cambridge, MA, 1989.

[4] Ikeuchi, K. and Horn, B.K.P. (1981) Numerical Shape from Shading and Occluding Boundaries, *Artificial Intelligence,* Vol. 17, No. 3, pp. 141-184.

[5] Kimmel, R. and Bruckstein, A.M. (1995) Tracking Level-sets by Level-sets: A Method for Solving the Shape from Shading Problem, *CVIU,* Vol. 62, No. 1, pp. 47-58.

[6] Kimmel R., Siddiqi K., Kimia B.B., and Bruckstein A.M., "Shape from shading - Level set propagation and viscosity solutions", Int J. Computer Vision Vol 16, No 2, pp. 107-133, 1995.

[7] 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.

[8] Worthington, P.L. and Hancock, E.R. (1999) New Constraints on Data-closeness and Needle-map consistency for SFS, *IEEETransactions on Pattern Analysis,* Vol. 21, pp. 1250-1267.