Vector Transport for Shape-from-shading

Eraldo Ribeiro Edwin Hancock
University of York,York, UK.
erh@cs.york.ac.uk

Abstract

In this paper we describe a new shape-from-shading method. We show how the parallel transport of surface normals can be used to impose curvature consistency and to iteratively update surface normal directions so as to improve the brightness error. We commence by showing how to make local estimates of the Hessian matrix from surface normal information. With the local Hessian matrix to hand, we develop an "EM-like" algorithm for updating the surface normal directions. At each image location, parallel transport is applied to the neighbouring surface normals to generate a sample of local surface orientation predictions. From this sample, a local weighted estimate of the image brightness is made. The transported surface normal which gives the brightness prediction which is closest to this value is selected as the revised estimate of surface orientation. The revised surface normals obtained in this way may in turn be used to re-estimate Hessian matrix, and the process iterated until stability is reached.

1 Introduction

Shape-from-shading is a problem that has been studied for over 25 years in the vision literature [1, 7, 10, 13]. 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.

There have been several distinct approaches to the shape-from-shading problem. The classic approach developed by Ikeuchi and Horn [6] and, by Horn and Brooks [3], among others, is an energy minimisation one based on regularisation theory. Here the dual constraints of compliance with the image irradiance equation and local surface smoothness are captured by an error function. This has distinct terms corresponding to data-closeness, i.e. compliance with the image irradiance equation, and for surface smoothness, i.e. the constraint that the local variation in the surface normal directions should be small. The shortcomings with this method are threefold. First, it is sensitive to the initial surface normal directions, second the data-closeness and surface smoothness must be carefully balanced, and, finally, the solution found is invariably dominated by the smoothness model and as a result fine surface detail is lost. The second approach to the problem of shape-from-shading has been to adopt the apparatus of level-set theory to solve the underlying differential equation [7, 8, 11]. This offers two advantages. First, the recovered solution is provably correct, and second, the surface height function is recovered at the same time as the field of surface normals. The third approach was recently developed by Worthington and Hancock [13]. This method adopts the view that the image irradiance equation should be treated as a hard constraint and that curvature consistency constraints should be used in preference to local smoothing. They develop a shape-form-shading algorithm in which the surface normals are constrained to fall on the irradiance cone whose apex angle is determined by the local image brightness. The surface normals are initialised to point in the direction of the local Canny image gradient. These directions on the cone are updated by smoothing the surface normal directions is a manner which is sensitive to local surface topography.

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 [13, 12], 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 adopt a simple model of the distribution of brightness estimates based on the assumption that the original intensity image is subject to Gaussian measurement errors. Using this distribution, we compute the mean predicted brightness value for the sample of transported surface normals. We select a revised local surface normal direction by identifying the transported vector which gives the brightness that is closest to the mean-value.

This process may be iterated until stability is reached. From the revised surface normal directions, we make new estimates of the local Hessian matrices. These matrices in-turn are used for neighbouring surface normal transportation, and the samples of surface normals so-obtained are used to estimate mean brightness. Viewed in this way our algorithm has many features reminiscent of the EM algorithm. The surface normals may be regarded as hidden or missing data that must be recovered from the observed image brightness. In the expectation-step, we compute the mean image brightness. The maximisation step is concerned with finding the revised surface normal directions that minimise the weighted brightness error. From the perspective of differential geometry, one of the attractive features of our algorithm is that it provides a statistical framework for combining evidence for shading patterns from the Gauss map.

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.

Hence, the novelty of our contribution is twofold. First, we develop an evidence combining algorithm for shape-from-shading. There has been little previously documented attempts to do this in the literature. Second, is our idea of using parallel transport to ensure consistency with differential geometry. Here there are two pieces of related work. First, Lagarde and Ferri [2] have shown how the Darboux smoothing idea of Sander and Zucker can be applied to smooth extracted needle-maps as a post-processing step. Second, Worthington and Hancock [13] have shown how the variance of the Koenderinck and Van Doorn shape-index can be used to control the robust smoothing of surface normal directions.

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 [10, 5, 1] 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

$$R(p, q) = \mathbf{Q} \cdot \mathbf{s} \tag{1}$$

The image irradiance equation [4] 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

$$E(x, y) = R(p, q) \tag{2}$$

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.

The process of smooth surface recovery is posed as a variational problem in which a global error-functional is minimized through the iterative adjustment of the needle map. Surface normals are updated with a step-size dictated by Euler's equation. Here we consider the formulation of Brooks and Horn [**?**] which is couched in terms of unit surface normals. Their error functional is defined to be

$$I = \int\int \underbrace{\left(E(x, y) - \mathbf{Q} \cdot \mathbf{s} \right)^2}_{BrightnessError} + \lambda \underbrace{\left(\left\| \frac{\partial \mathbf{Q}}{\partial x} \right\|^2 + \left\| \frac{\partial \mathbf{Q}}{\partial y} \right\|^2 \right)}_{RegularizingTerm} dx\, dy \tag{3}$$

The terms $\frac{\partial \mathbf{Q}}{\partial x}$ and $\frac{\partial \mathbf{Q}}{\partial y}$ above are the directional derivatives of the needle-map in the x and y directions respectively. The magnitudes of these quantities are used to measure the smoothness of the surface, with a large value indicating a highly-curved region. However, it should be noted that a planar surface has $\frac{\partial \mathbf{Q}}{\partial x} = \frac{\partial \mathbf{Q}}{\partial y} = \mathbf{0}$ in this case.

The first term of Equation 3 is the brightness error, which encourages data-closeness of the measured image intensity and the reflectance function. The *regularizing term* imposes the smoothness constraint on the recovered surface normals, penalising large local changes in surface orientation. The constant λ is a Lagrange multiplier. For numerical stability, λ must often be large, resulting in the smoothness term dominating.

Minimization of the functional defined in Equation 3 is accomplished by applying the calculus of variations and solving the resulting Euler equation. The solution is

$$\mathbf{Q}_o^{(k+1)} = \bar{\mathbf{Q}}_o^{(k)} + \frac{\epsilon^2}{2\lambda}\left(E_o - \mathbf{Q}_o^{(k)} \cdot \mathbf{s}\right)\mathbf{s} \tag{4}$$

where ϵ is the spacing of pixel-sites on the lattice and $\bar{\mathbf{Q}}_o$ is the local mean of the surface normals around the neighbourhood R_o of the pixel at position o

$$\bar{\mathbf{Q}}_o = \frac{1}{|R_o|}\sum_{m \in R_o} \mathbf{Q}_o \tag{5}$$

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

$$\mathcal{H} = \begin{pmatrix} \frac{\partial^2 z}{\partial x^2} & \frac{\partial^2 z}{\partial x \partial y} \\ \frac{\partial^2 z}{\partial x \partial y} & \frac{\partial^2 z}{\partial y^2} \end{pmatrix} \tag{6}$$

The eigen-structure of the Hessian matrix can be used to gauge the curvature of the surface. The two eigen-values of $\mathcal{H}$ are the maximum and minimum curvatures. The orthogonal eigen-vectors of $\mathcal{H}$ are known as the principal curvature directions. The mean-curvature of the surface is found by averaging the maximum and minimum curvatures. Finally, the Gaussian curvature is equal to the product of the two eigenvalues.

In the case when surface normal information is being used to characterise the surface, then the Hessian matrix takes on the following form

$$\mathcal{H} = \begin{pmatrix} \alpha & \beta \\ \beta & \gamma \end{pmatrix} \tag{7}$$

The diagonal elements of the Hessian are related to the rate-of change of the surface normal components via the equations

$$\alpha = \left(\frac{\partial \mathbf{Q}}{\partial x}\right)_x \tag{8}$$

$$\gamma = \left(\frac{\partial \mathbf{Q}}{\partial x}\right)_y \tag{9}$$

where the subscripts x and y on the large brackets indicate that the x or y components of the vector-derivative are being taken.

Treatment of the off-diagonal elements is more subtle. However, if we assume that the surface can be represented by a twice differentiable function $z = f(x, y)$, then we can write

$$\beta = \left(\frac{\partial \mathbf{Q}}{\partial y}\right)_x = \left(\frac{\partial \mathbf{Q}}{\partial x}\right)_y \tag{10}$$

In the next Section we will describe how the elements of the Hessian, i.e. α, β and γ, can be estimated from raw surface normal data using the method of least-squares.

4 Estimating the Hessian

In this section we describe how to make a statistical estimate of the Hessian matrix from a sample of surface normals delivered by shape-from-shading. Specifically, we use the method of least squares to estimate the elements of $\mathcal{H}$.

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). If the normals are close to each other, then we can approximate the change in the components of the surface normal using a first-order Taylor expansion. Accordingly,

$$(\Delta Q_m)_x = \left(\frac{\partial \mathbf{Q}}{\partial x}\right)_x \Delta x_m + \left(\frac{\partial \mathbf{Q}}{\partial y}\right)_x \Delta y_m \tag{11}$$

$$(\Delta Q_m)_y = \left(\frac{\partial \mathbf{Q}}{\partial x}\right)_y \Delta x_m + \left(\frac{\partial \mathbf{Q}}{\partial y}\right)_y \Delta y_m \tag{12}$$

where the measured change in the components of the surface normal is given by $\mathbf{Q}_m - \bar{\mathbf{Q}}_o = ((\Delta Q_m)_x, (\Delta Q_m)_x)^T$. The displacements in point co-ordinates are $\Delta x_m = x_m - x_o$ and $\Delta y_m = y_m - y_o$. We can rewrite the first-order Taylor expansion in terms of elements of the Hessian matrix, i.e.

$$(\Delta Q_m)_x = \alpha_o \Delta x_m + \beta_o \Delta y_m \tag{13}$$

$$(\Delta Q_m)_y = \beta_o \Delta x_m + \gamma_o \Delta y_m \tag{14}$$

where α_o, β_o and γ_o are the elements of the Hessian matrix at the pixel indexed o. These equations govern the parallel transport of the vector across the curved geometry of the surface. So, to first-order, the change in the normal is linear in the elements of the Hessian matrix. Unfortunately, for the single neighbouring normal these equations are under-constrained and we can not recover the Hessian. However, if we have a sample of N neighboring surface normals, then there are $2N$ homogenous linear equations in the elements of $\mathcal{H}$ and the problem of recovering differential structure is no-longer under-constrained. Under these circumstances, we can estimate the elements of the Hessian matrix using the method of least-squares.

To proceed, we make the homogeneous nature of the equations more explicit by writing

$$\begin{array}{lccccc} (\Delta Q_m)_x = & \Delta x_m \cdot \alpha_o & + & \Delta y_m \cdot \beta_o & + & 0 \cdot \gamma_o \\ (\Delta Q_m)_y = & 0 \cdot \alpha_o & + & \Delta x_m \cdot \beta_o & + & \Delta y_m \cdot \gamma_o \end{array} \tag{15}$$

In order to simplify notation, we can write the full system of 2N equations in matrix form as

$$\mathbf{N} = \mathbf{X}\mathbf{\Phi_o} \tag{16}$$

where $\mathbf{N}$ is an aggregated column-vector of normal components

$$\mathbf{N} = \begin{pmatrix} (\Delta Q_1)_x \\ (\Delta Q_1)_y \\ (\Delta Q_2)_x \\ \vdots \end{pmatrix},$$

The design matrix $\mathbf{X}$ is a matrix of co-ordinate displacements

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

and Φ_o is the parameter vector

$$\Phi_o = \begin{pmatrix} \alpha_o \\ \beta_o \\ \gamma_o \end{pmatrix}$$

When the system of equations is over-specified in this way, then we can extract the set of parameters that minimises the vector of error-residuals $\mathbf{N} - \mathbf{X}\Phi_o$. We pose this parameter recovery process as a least-squares estimation problem. In other words we seek the vector of estimated parameters $\hat{\mathbf{\Phi}}_\mathbf{o} = (\hat{\alpha}_\mathbf{o}, \hat{\beta}_\mathbf{o}, \hat{\gamma}_\mathbf{o})^\mathbf{T}$ which satisfy the condition

$$\hat{\mathbf{\Phi}}_\mathbf{o} = \arg\min_{\mathbf{\Phi}} (\mathbf{N} - \mathbf{X\Phi})^\mathbf{T}(\mathbf{N} - \mathbf{X\Phi}) \qquad (17)$$

The solution-vector is found by computing the pseudo-inverse of the design matrix $\mathbf{X}$ thus

$$\hat{\Phi}_o = (\mathbf{X^T X})^{-1}\mathbf{X^T N} \qquad (18)$$

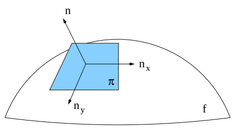

Figure 1. Surface normal vector and its components.

The surface geometry is illustrated in Figure 1.

5 Parellel Transport

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 [2] and Worthington and Hancock [13], 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 [13], 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 approaximation to the transported vector is

$$\mathbf{Q}_m^o = \mathbf{Q}_m + \mathcal{H}_o(\vec{X}_m - \vec{X}_o) \qquad (19)$$

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 2.

6 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. 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

$$p(E_o|\mathbf{Q}_m, \Phi_o) = \frac{1}{\sqrt{2\pi}\sigma} \exp\left[-\frac{(E_o - \mathbf{Q}_m^o.\mathbf{s})^2}{2\sigma^2}\right] \qquad (20)$$

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

$$\hat{E}_o = \sum_{m \in R_o} p(E_o|\mathbf{Q}_m, \Phi_o)\mathbf{Q}_m^o.\mathbf{s} \qquad (21)$$

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

$$\hat{\mathbf{Q}}_o = \arg\min_{m \in R_o} \left[\hat{E}_o - \mathbf{Q}_m^o.\mathbf{s}\right]^2 \qquad (22)$$

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 [13]. 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.

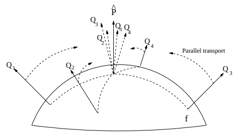

Figure 2. Parallel transport used for predicting the surface normal vector using local curvature estimation.

7 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. In Figure 3 we show the field of surface normal directions with iteration number. The main feature to note here is that the surface details become more marked with iteration number. In Figure 4, we plot the difference between the measured and predicted image brightness as a function of iteration number. Initially, the error is greatest on the curved regions of the surface (near the head, beak, neck, wings and tail). After 10 iterations, the only region where there is a significant error is around the eye, where there is an albedo difference. There are some high curvature points around the neck and the wing where there is also some residual brightness error. The average brightness error is plotted as a function of iteration number in Figure 5. Figure 6 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.

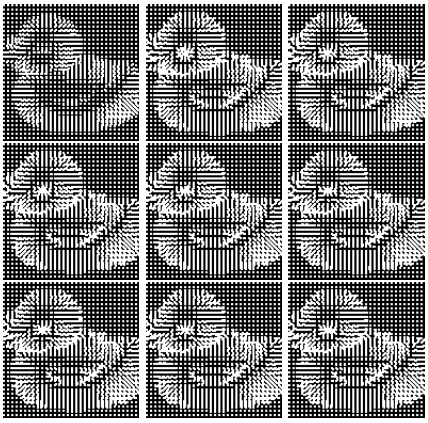

Figure 3. Needle maps for the small duck image

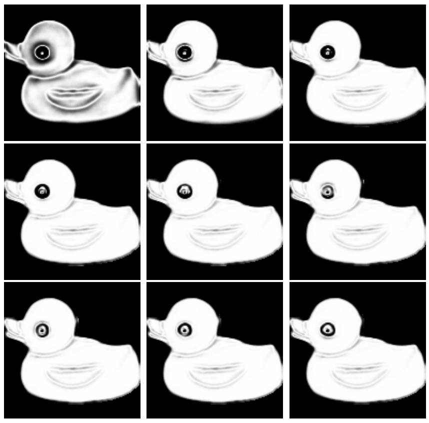

Figure 4. Increasing probability of image intensity agreement

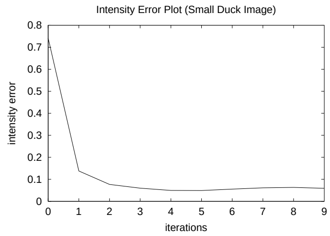

Figure 5. Error plot for the intensity value recovering for the small duck image

In Figures 7, 8 and 9 we show results for images of marble statues. In each image the top left-hand panel is the original image, the top right-hand panel is the final value of $\hat{E}_o$, and the bottom two images show the initial (left) and final (right) needle-maps. In each case the final needle-maps and brightness images reproduce the curvature structure well, at all but the points of highest curvature.

Finally, in Figure 10, 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.

8 Conclusions

In this paper we have described a new method for shapeform-shading which relies on vector transport to accumulate evidnece 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

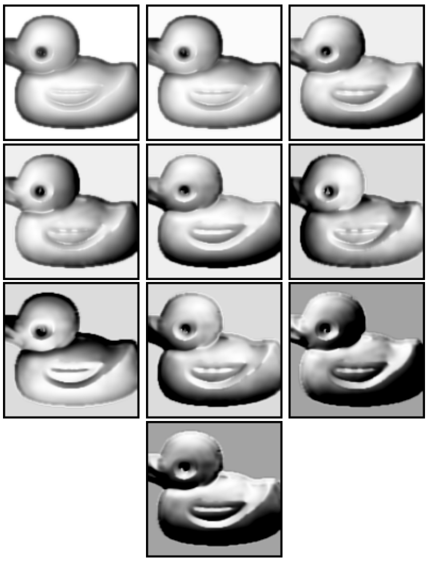

Figure 6. Reconstructed image with different illumination directions

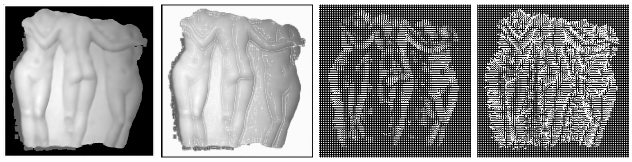

Figure 7. The three graces

to the expected value. The method is evaluated on a variety of real-world images where it provides promissing results.

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] 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. (1986) The Variational Approach to Shape from Shading, *CVGIP,* Vol. 33, No. 2, pp. 174-208.

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

[5] Horn, B.K.P. (1990) Height and Gradient from Shading, *IJCV,* Vol. 5, No. 1, pp. 37-75.

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

[7] 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.

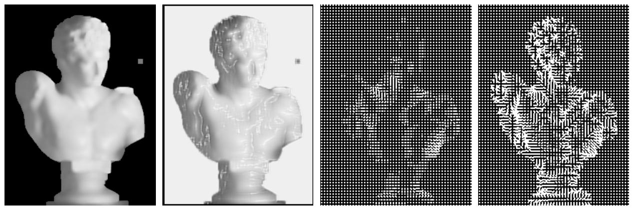

Figure 8. Hermes

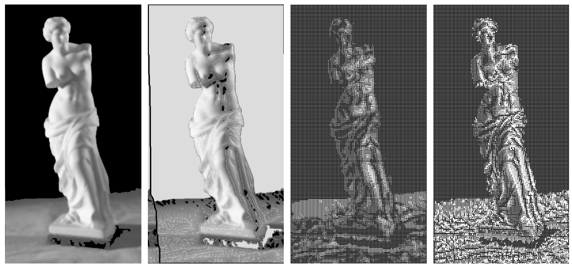

Figure 9. Venus

[8] 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.

[9] Koenderink, J.J. (1990) *Solid Shape,* MIT Press, Cambridge MA.

[10] 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.

[11] Rouy E. and Tourin A., "A viscosity solutions approach to shape-from-shading", SIAM J. Numerical Analysis, vol 29, No 3, pp. 867-884, 1992.

[12] Worthington, P.L. and Hancock, E.R. (1998) Needle Map Recovery using Robust Regularizers, *Image and Vision Computing,* Vol. 17, No. 8, pp. 545-559.

[13] 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.

Figure 10. Re-illumination of Venus