



Deposited via The University of York.

White Rose Research Online URL for this paper:

<https://eprints.whiterose.ac.uk/id/eprint/1984/>

Article:

Robles-Kelly, Antonio and Hancock, Edwin R. (2007) Shape-from-shading using the heat equation. IEEE Transactions on Image Processing. pp. 7-21. ISSN: 1057-7149

<https://doi.org/10.1109/TIP.2006.884945>

Reuse

Items deposited in White Rose Research Online are protected by copyright, with all rights reserved unless indicated otherwise. They may be downloaded and/or printed for private study, or other acts as permitted by national copyright laws. The publisher or other rights holders may allow further reproduction and re-use of the full text version. This is indicated by the licence information on the White Rose Research Online record for the item.

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.

Shape-From-Shading Using the Heat Equation

Antonio Robles-Kelly, *Member, IEEE*, and Edwin R. Hancock

Abstract—This paper offers two new directions to shape-from-shading, namely the use of the heat equation to smooth the field of surface normals and the recovery of surface height using a low-dimensional embedding. Turning our attention to the first of these contributions, we pose the problem of surface normal recovery as that of solving the steady state heat equation subject to the hard constraint that Lambert's law is satisfied. We perform our analysis on a plane perpendicular to the light source direction, where the z component of the surface normal is equal to the normalized image brightness. The $x - y$ or azimuthal component of the surface normal is found by computing the gradient of a scalar field that evolves with time subject to the heat equation. We solve the heat equation for the scalar potential and, hence, recover the azimuthal component of the surface normal from the average image brightness, making use of a simple finite difference method. The second contribution is to pose the problem of recovering the surface height function as that of embedding the field of surface normals on a manifold so as to preserve the pattern of surface height differences and the lattice footprint of the surface normals. We experiment with the resulting method on a variety of real-world image data, where it produces qualitatively good reconstructed surfaces.

I. INTRODUCTION

SHAPE-from-shading is a problem that has been studied for over 25 years in vision literature [1]–[7]. Stated succinctly, the problem is to recover local surface orientation information and, hence, reconstruct the surface height function, from the information provided by the surface brightness. Since the problem is an under-constrained one, in order to be rendered tractable, recourse must be made to strong simplifying assumptions and constraints. Hence, the process is usually specialized to matte reflectance from a surface of constant albedo, illuminated by a single point 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, singular points or constraints on surface smoothness.

In this paper, we make two contributions to the shape-from-shading problem. First, we show how the recovery of the field of surface normals can be posed as the solution of the linear heat equation. Second, we present a new method for surface height

Manuscript received February 1, 2005; revised May 1, 2006. A. Robles-Kelly was supported in part by the Australian Governments Backing Australia's Ability Initiative and in part by the Australian Research Council. The associate editor coordinating the review of this manuscript and approving it for publication was Dr. Benoit Macq.

A. Robles-Kelly is with the National ICT Australia, Australian National University, Canberra ACT 0200, Australia (e-mail: antonio.robles-kelly@nicta.com.au; antonio.robles-kelly@anu.edu.au).

E. R. Hancock is with the Department of Computer Science, University of York, York YO1 5DD, U.K. (e-mail: auerh@cs.york.ac.uk).

Color versions of Figs. 2, 3, 9–14, and 16 are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TIP.2006.884945

recovery from the field of surface normals. In this section of the paper, we commence by reviewing the literature in these two areas to motivate our contribution and highlight the novelty of the proposed methods.

A. Related Literature

There are a number of perspectives from which the shape-from-shading literature can be viewed. For instance, Zhang *et al.* [8] divide existing shape-from-shading methods into those that are local and those that are global. Local methods involve quilting local surface patches to recover the overall surface shape. Although fast, the methods require prior information concerning height, which may be provided by, for instance, the elevation of singular points. As a result, local methods are often sensitive to noise. Global methods, on the other hand, recover the height through the propagation of smoothness constraints or by minimisation of an energy function. Although more robust to noise, global methods can have a tendency to over-smooth the recovered surface. An exhaustive review of the topic, which includes a detailed comparative study can be found in the recent paper by Zhan *et al.* [8].

The classic approach developed by Ikeuchi and Horn [9] and by Horn and Brooks [10], 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. Third, and finally, the solution found is invariably dominated by the smoothness model and, as a result, fine surface detail is lost. There have been attempts to overcome these problems. For instance, Zheng and Chellappa [11] have imposed a gradient consistency constraint that penalizes differences between the image intensity gradient and the surface gradient for the recovered surface. Worthington and Hancock [12] impose the Lambertian radiance constraint in a hard manner by demanding that the recovered surface normals lie on cones whose axis is in the light source direction and whose apex angle is the inverse cosine of the normalized image radiance. The smoothing process may also be effected using deformable implicit surfaces [13].

The shape-from-shading problem can also be viewed as that of solving a nonlinear partial differential equation. The early work mentioned above posed the recovery of the solution of the irradiance equation in a variational setting and developed iterative numerical schemes [10]. Later work by Dorou and Maitre [14] unearthed problems with this method, and demonstrated

both that the method was not linearly convergent and that multiple ambiguous solutions were possible. Dupuis and Oliensis [6] provided a means of overcoming these problems and recovering probably correct solutions with respect to the image irradiance equation. Under conditions in which the light-source and viewer directions are identical, Rouy and Tourin [15] showed how the Hamilton–Jacobi equations could be used to recover solutions that are probably convergent. Recently, Prados and Faugeras [16] have extended this work to the case where the light source and viewer directions are no longer coincident. This leads to continuous solutions to the image irradiance equation, that, in general, lead to a shape-from-shading scheme that can work in more realistic situations. In related work, Kimmel and Bruckstein have used the apparatus of level sets methods to recover solutions to the Eikonal equation [17]. The resulting framework can also be applied to image regularisation [18]. Finally, we note that the problem of recovering implicit surfaces to represent objects can also be posed as solving a system of partial differential equations (PDEs) [13].

Another important issue concerning the shape-from-shading process is that which pertains the reconstruction of the surface from the estimated surface normal directions. The so-called surface integration process involves selecting a path through the surface normal locations. This may be done using either a curvature minimizing path or by advancing a wavefront from the occluding boundary or singular points. If the surface normals do not satisfy the integrability constraint (i.e., the Hessian matrix is symmetric) then the recovered height can depend on the path chosen through the surface normals.

The analysis of the literature on the topic of surface height recovery is not a straightforward task. The reason for this is that surface recovery is frequently viewed as an integral part of shape-from shading or shape-from-texture processes. Horn and Brooks [19] realize surface height recovery as a postprocessing step. The process proceeds from the occluding boundary and involves incrementing the surface height by an amount determined by the distance traversed and the slope angle of the local tangent plane. In some of the earliest work, Wu and Li [20] average the surface normal directions to obtain a height estimate. A more elegant solution is proposed by Frankot and Chellappa [21] who project the surface normals into the Fourier domain to impose integrability constraints and surface height is recovered using an inverse Fourier transform. Wei and Klette [22] have enhanced this approach by showing how more complex regularisation constraints can be formulated in the Fourier domain. Dupuis and Oliensis [6] have developed a method which draws on differential geometry and involves propagation in the direction of the steepest gradient from singular points. A fast variant of this algorithm is described by Bichsel and Pentland [3] who compute the relative height of the surface with respect to the highest intensity point.

B. Contribution and Paper Outline

To set our original contributions in the context of this literature, we make the following observations. First, most of the existing work hinges around the solution of a nonlinear partial differential equation and this can prove cumbersome. Our first contribution in this paper, is to show how the shape-from-shading

problem can be posed as the solution of the linear heat-equation. We pose the problem as that of solving the heat equation subject to the constraint that the recovered surface normals satisfy Lambert’s law. We work on a plane perpendicular to the light source direction. To this effect, we perform a rotation of the coordinate system. In the transformed coordinate system, the z component of the surface normal is constrained to be equal to the normalized image brightness. To compute the x and y , i.e., azimuthal, components of the surface normal, we introduce a scalar potential. We assume that the time-dependance of the scalar potential during the surface normal smoothing process is governed by the heat equation and show that this scalar potential is given by the average of the image intensity. This allows us to use a standard finite element method to compute the components of the surface normals. Hence, smoothed surface normals that satisfy the image irradiance equation are recovered as solutions of the linear heat equation. This offers a number of advantages over existing differential equation methods for shape-from-shading.

The second observation that we draw from the literature is that the problem of surface height recovery is usually posed as one of path-based integration. This process can prove sensitive to noise, and does not model the relationship between the observed set of surface normals and the surface that generated them. Our second contribution in this paper is to make this link explicit. To do this, we pose the problem of recovering the surface height function as that of embedding the field of surface normals on a manifold. Our approach is as follows. From the field of surface normals, we compute the surface height increments corresponding to each location on the pixel lattice. The height increments, in turn, can be used to estimate the surface height difference between each pair of pixel-locations. We pose the problem of surface height recovery as that of embedding the surface normals into a manifold in a 3-D space so as to preserve both the pattern of surface height differences and the lattice footprint of the field of surface normals.

The outline of the paper is as follows. In Section II, we briefly review the physics of Lambertian reflectance. Section III presents the first novel contribution of the paper and demonstrates how the recovery of the surface normal directions, which satisfies Lambert’s law, can be posed as the solution of the linear heat equation. Section IV describes the second novel contribution, i.e., the recovery of the surface height function using a unidimensional embedding method. Section V presents experiments on data with known ground truth. Here, we separately compare the surface normal recovery method and the surface height recovery method with alternatives described elsewhere in the literature. Finally, Section VI offers some conclusions and suggests directions for future research.

II. LAMBERTIAN REFLECTANCE

We confine our attention to surfaces which exhibit Lambertian reflectance. For Lambertian reflectance from a matte surface, the image irradiance equation constrains the surface normal directions to fall on a cone whose axis is aligned with the light source direction and whose apex angle is the inverse cosine of the normalized image brightness [12]. To be more formal, let \vec{L} be a unit vector in the light source direction, $I(p_k)$ be the normalized brightness at the pixel-site p_k on the

image plane Π , and $\vec{N}(s_k) = [N_x(s_k), N_y(s_k), N_z(s_k)]^T$ be the corresponding surface normal at a point s_k on the surface under study S . From Lambert's law, we have

$$I(p_k) = \vec{N}(s_k) \cdot \vec{L}. \quad (1)$$

Thus, the surface normal is constrained to fall on a cone, whose axis is the direction on the light source \vec{L} and whose opening angle is $\gamma_{s_k} = \arccos(I(p_k))$.

Our aim is to recover the on-cone surface normal which satisfies the image irradiance equation as a hard constraint by computing a vector $\vec{N}'(s_k) = [N'_x(s_k), N'_y(s_k), N'_z(s_k)]^T$ that satisfies the conditions

$$I(p_k) = \vec{L} \cdot \vec{N}'(s_k) \text{ and } \vec{N}(s_k) = \mathbf{R}\vec{N}'(s_k) \quad (2)$$

where \mathbf{R} is a rotation matrix.

To simplify our analysis, we choose to work in a transformed coordinate system, in which the image plane is rotated to be perpendicular to the light sources direction. This rotation can be realized as follows. As a consequence of Lambert's law, the apex angle $\gamma_{s_k} = \arccos(I(p_k))$ does not depend on the viewer direction \vec{V} . Hence, we can consider the vector

$$\begin{aligned} \vec{N}'(s_k) &= [N'_x(s_k), N'_y(s_k), N'_z(s_k)]^T \\ &= [\sin(\theta'_{s_k}) \cos(\alpha'_{s_k}), \\ &\quad \sin(\theta'_{s_k}) \sin(\alpha'_{s_k}), \cos(\theta'_{s_k})]^T \end{aligned} \quad (3)$$

to lay on a plane $\hat{\Pi}$ perpendicular to the light source direction \vec{L} whose x and y axes correspond to the image row and column directions rotated about the light source direction $\vec{L} = [L_x, L_y, L_z]^T$. The matrix \mathbf{R} rotates the vector $\vec{N}'(s_k)$ by the angle difference between the light source direction \vec{L} , and the viewer direction \vec{V} . The rotation axis is the light source vector $\vec{L} = [L_x, L_y, L_z]^T$. The angle of rotation is $\theta_L = \arccos(L_z)$. Hence, the rotation matrix is as follows:

$$\mathbf{R} = \begin{bmatrix} L_z + L_x^2 c & -L_z s + L_x L_y c & L_y s + L_x L_z c \\ L_z s + L_x L_y c & L_z + L_y^2 c & -L_x s + L_y L_z c \\ -L_y s + L_x L_z c & L_x s + L_y L_z c & L_z(1 + L_z c) \end{bmatrix} \quad (4)$$

where $c = 1 - L_z$ and $s = \sin(\arccos(L_z))$. The geometry of this procedure is shown in Fig. 1.

III. SURFACE NORMAL COMPUTATION VIA VECTORIAL FLUX

Lambert's law provides only a constraint on the zenith angle between the light source direction and the surface normal direction and, as a result, the azimuth angle of the surface normal remains undetermined. There are number of ways in which the azimuth angle can be recovered. For instance, the surface normal may be aligned on the irradiance cone so that it points in the direction of the image gradient. The direction may also be determined using local smoothing. In this section, we describe a method that can be used to recover the azimuth angle using a smoothing method based on a heat flow analogy.

We work on a plane $\hat{\Pi}$ perpendicular to the light source direction. Let $\hat{p}_k = (u, v, w)$ be the coordinates of the point with pixel index k on the plane $\hat{\Pi}$. Since the surface normals is constrained to fall on the irradiance cone and we are working on




Fig. 1. Geometry of the SFS process.

the plane perpendicular to the light source direction, the z component of the surface normal is equal to the normalized image brightness $I(p_k)$.

In this paper, our idea is to recover the azimuthal component using a heat-equation smoothing process. To do this we introduce a scalar field on the rotated image plane $\hat{\Pi}$. The azimuthal component of the surface normal is the gradient of this scalar field. The scalar field is assumed to be subject to the heat equation, and the associated heat flow or diffusion process is responsible for smoothing the azimuthal component of the surface normal. Although the azimuthal component of the surface normal is determined by the gradient of the scalar field, in practice we are only interested in the polar angle of the gradient vector since this determines the azimuthal angle of the surface normal. The heat equation provides a means of smoothing the azimuthal directions of the surface normals on the plane $\hat{\Pi}$.

To be more formal, let $Q(x, y, t)$ be a time-dependent field on the plane $\hat{\Pi}$, where $(x, y)^T$ is the coordinate vector of the point \hat{p}_k and t is the time epoch. According to Fourier's law, the heat flow vector \vec{q} is related to the gradient of the scalar potential field $Q(x, y, t)$ through the expression

$$\vec{q} = -\kappa \nabla Q(x, y, t). \quad (5)$$

The magnitude of the time derivative of the heat flow \dot{q} is related to the time derivative of the scalar potential via the heat diffusion equation

$$\nabla \cdot \kappa \nabla Q(x, y, t) + \dot{q} = \frac{1}{\alpha} \frac{\partial Q(x, y, t)}{\partial t} \quad (6)$$

where κ and α are proportionality constants. With these ingredients, the azimuthal angle of the surface normal on the plane $\hat{\Pi}$ can then be found by computing the angle between the components of the gradient for the scalar field $Q(x, y, t)$, which is given by

$$\nabla \cdot Q(x, y, t) = \left[\frac{\partial Q(x, y, t)}{\partial x}, \frac{\partial Q(x, y, t)}{\partial y} \right]^T.$$

Here, we consider the heat release in the direction perpendicular to the plane $\hat{\Pi}$, i.e., \dot{q} , to be constant and equal to the




Fig. 2. Top row, from left to right: Ground truth height data, rendering of the test shape using Lambert's law with added noise, field of surface normals at the first and final iterations of the algorithm. Bottom row, from left to right: Heat function $Q(x, y, t)$ for $t = 1, t = 3, t = 6,$ and $t = 9$.




Fig. 3. Shape-from-shading results on real-world data. Top row: Input images whose specularities have been removed. Bottom row: Ground truth depth data.

brightness at the point $p_k = (x, y)$. Hence, and by expanding the Laplacian, we can rewrite the heat diffusion equation as follows:

$$\kappa \left(\frac{\partial^2 Q(x, y, t)}{\partial x^2} + \frac{\partial^2 Q(x, y, t)}{\partial y^2} \right) + I(x, y) = \frac{1}{\alpha} \frac{\partial Q(x, y, t)}{\partial t}. \quad (7)$$

To take our analysis further, we rely on the physical meaning of the right-hand side of (7). The quantity $(1/\alpha)(\partial Q(x, y, t)/\partial t)$ is the rate of energy increase for an element of surface dS enclosed by a region R . Hence, we can normalize the heat-flow to unity by setting $(1/\alpha)(\partial Q(x, y, t)/\partial t) = 1$. As a result, and after rearranging terms, we get

$$\kappa \left(\frac{\partial^2 Q(x, y, t)}{\partial x^2} + \frac{\partial^2 Q(x, y, t)}{\partial y^2} \right) - (1 - I(x, y)) = 0 \quad (8)$$

and, therefore, we have

$$\int_R \nabla \cdot \nabla Q(x, y, t) dR = \int_R (1 - I(x, y)) dR \quad (9)$$

where we have set $\kappa = 1$.

At this point, since we are working on a uniformly sampled pixel lattice, we can consider uniform regions of area ρ across the $x - y$ plane and rewrite the equation above as follows:

$$\int_R \nabla \cdot \nabla Q(x, y, t) dR = \rho(1 - I(x, y)). \quad (10)$$

With the equation above at hand, we can make use of a simple finite element approach to recover the steady state of the func-

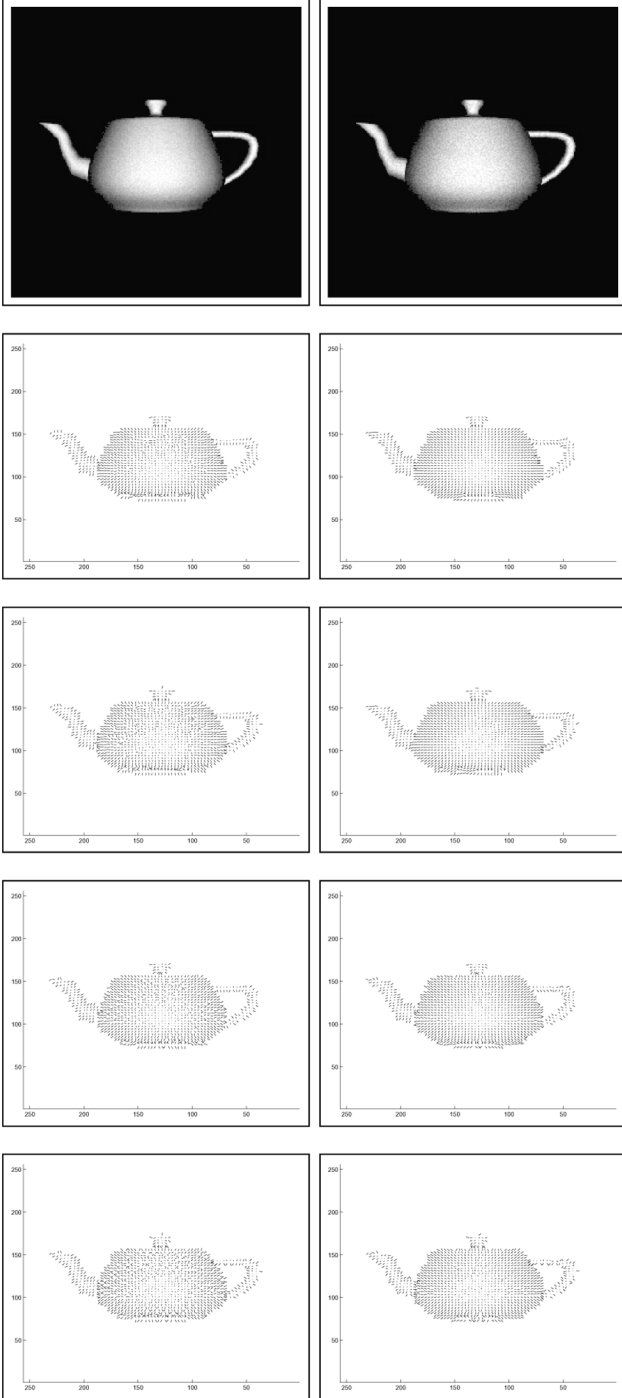


Fig. 4. First row: Input images with noise added. Second and third rows: Field of surface normals delivered by the algorithm of Worthington and Hancock after ten and 100 iterations. Fourth and fifth rows: Field of surface normals at the initial and final iterations of our method.

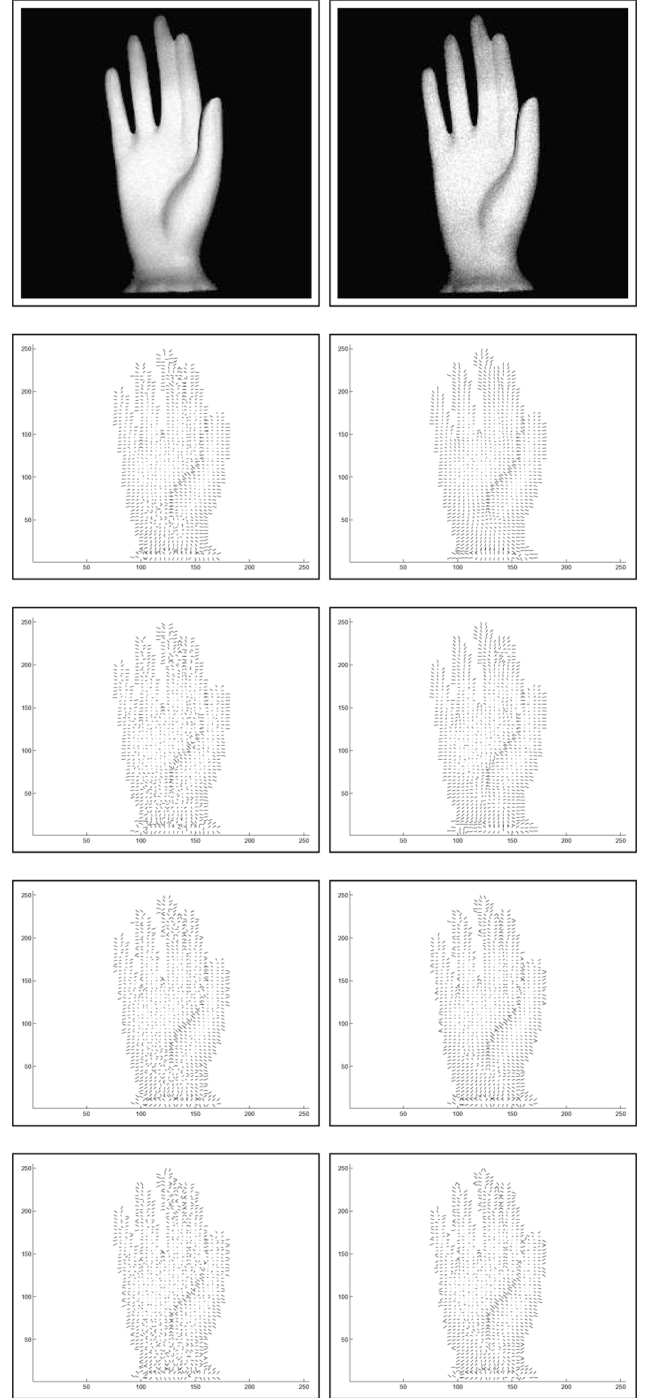


Fig. 5. First row: Input images with noise added. Second and third rows: Field of surface normals delivered by the algorithm of Worthington and Hancock after ten and 100 iterations. Fourth and fifth rows: Field of surface normals at the initial and final iterations of our method.

tion $Q(x, y, t)$. We commence by setting $Q(x, y, t = 0) \equiv 1$. At time step t_n , the update equation for the point indexed k is

$$Q(x_k, y_k, t_n) \approx \frac{\rho}{|\Omega_{\hat{p}_k}|} \sum_{p_l \in \Omega_{\hat{p}_k}} (1 - I(p_k)) \hat{Q}(x_k, y_k, t_{n-1}) \quad (11)$$

where $\Omega_{\hat{p}_k}$ is the set of points adjacent to \hat{p}_k . The quantity \hat{Q} is the averaged and normalized scalar potential, computed using

the formula

$$\hat{Q}(x_k, y_k, t_n) = \frac{\frac{\rho}{|\Omega_{\hat{p}_k}|} \sum_{p_l \in \Omega_{\hat{p}_k}} Q(x_l, y_l, t_n)}{\sum_{k \in \hat{\Pi}} \frac{\rho}{|\Omega_{\hat{p}_k}|} \sum_{p_l \in \Omega_{\hat{p}_k}} Q(x_l, y_l, t_n)}. \quad (12)$$

This process is repeated until the scalar field $Q(x_k, y_k, t_n)$ stabilizes. On convergence, the azimuth angle of the surface normal




Fig. 6. First row: Input images with noise added. Second and third rows: Field of surface normals delivered by the algorithm of Worthington and Hancock after ten and 100 iterations. Fourth and fifth rows: Field of surface normals at the initial and final iterations of our method.

at the location indexed k is given by

$$|\alpha'_{s_k}| = \left| \arctan \left(\frac{\Delta_x Q(x_k, y_k, t_n)}{\Delta_y Q(x_k, y_k, t_n)} \right) \right|. \quad (13)$$

In the above equation, $\Delta_x Q(x_k, y_k, t_n)$ and $\Delta_y Q(x_k, y_k, t_n)$ are the first difference approximations to the x and y compo-



Fig. 7. First row: Input images with noise added. Second and third rows: Field of surface normals delivered by the algorithm of Worthington and Hancock after ten and 100 iterations. Fourth and fifth rows: Field of surface normals at the initial and final iterations of our method.

nents of the gradient of Q , and are given by

$$\begin{aligned} \Delta_x Q(x_k, y_k, t_n) &= Q(x_k - 1, y_k, t_n) - Q(x_k + 1, y_k, t_n) \\ \Delta_y Q(x_k, y_k, t_n) &= Q(x_k, y_k - 1, t_n) - Q(x_k, y_k + 1, t_n). \end{aligned} \quad (14)$$

Since the value of $|\alpha'_{s_k}|$ constrains the azimuth of the rotated surface normal to be in the first quadrant, we make use of the

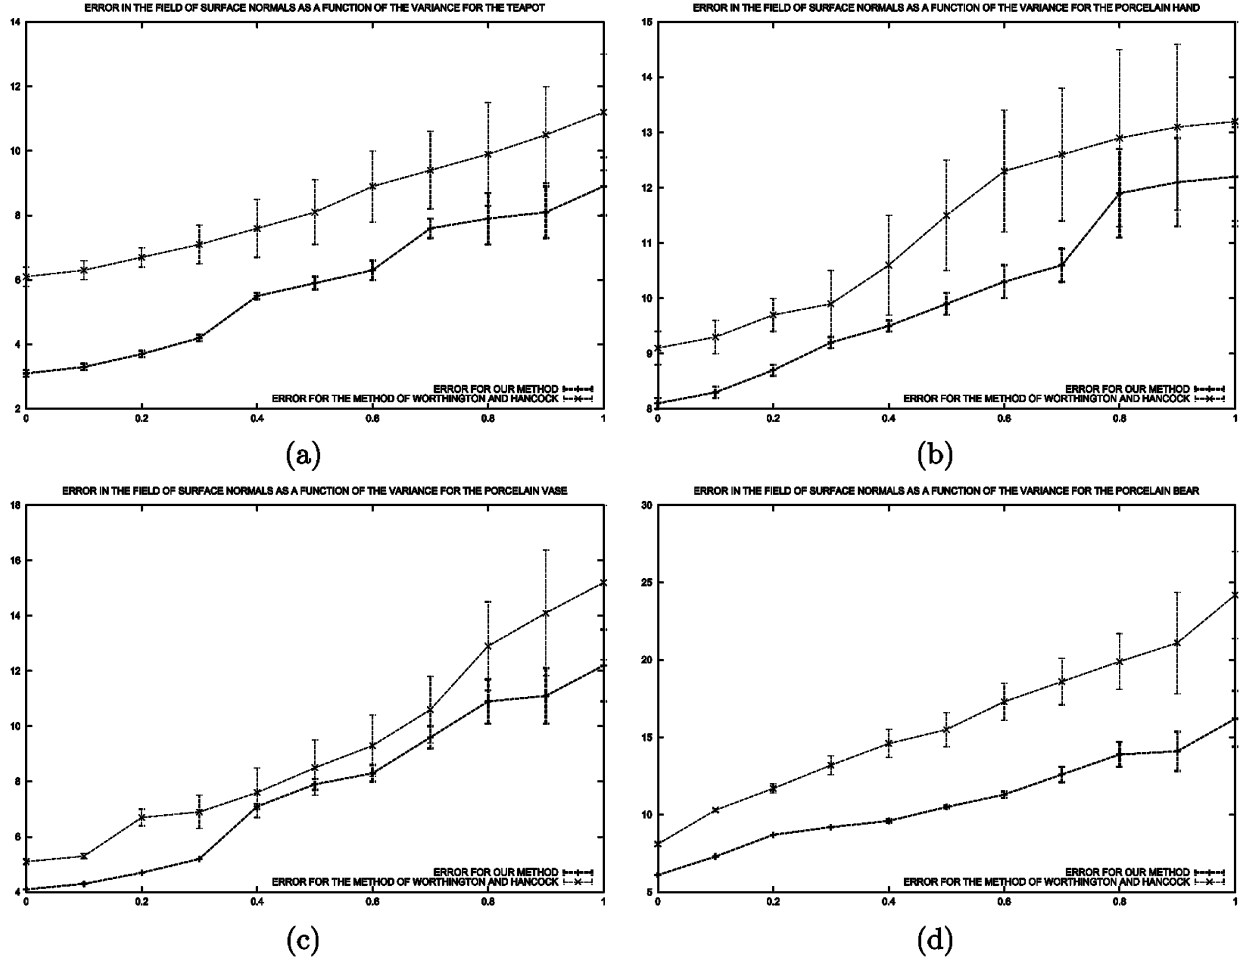


Fig. 8. Error as a function of variance for the four objects in our dataset.

sign of the flow in the x and y directions on the plane $\hat{\Pi}$ to determine the sign of surface normal components. Thus, in the rotated frame of reference the surface normal $\vec{N}'(s_k)$ is given by

$$\begin{aligned} \vec{N}'(s_k) &= \begin{bmatrix} N'_x(s_k) \\ N'_y(s_k) \\ N'_z(s_k) \end{bmatrix} \\ &= \begin{bmatrix} \text{sign}(\Delta_x Q(x_k, y_k, t_n)) \cos(|\alpha'_{s_k}|) \sin(\theta'_{s_k}) \\ \text{sign}(\Delta_y Q(x_k, y_k, t_n)) \sin(|\alpha'_{s_k}|) \sin(\theta'_{s_k}) \\ \cos(\theta'_{s_k}) \end{bmatrix} \end{aligned}$$

where $\theta'_{s_k} = \arccos(I(p_k))$. With the surface normal in the rotated frame of reference at hand, the surface normal on the image plane may be computed making use of the rotation matrix \mathbf{R} as given in (4).

IV. SURFACE HEIGHT RECOVERY

In this section, we address the problem of recovering the surface height function from the pattern of local height differences. To do this, we use the field of surface normals to estimate the height difference between each pair of pixel sites in the image. Once the height estimates have been computed, the surface may be then recovered by embedding them into the unidimensional Euclidean space perpendicular to the lattice footprint.

A. Height Difference Approximation

Let d_{s_k, s_l} be the estimate of the surface height difference between the pair of points s_k and s_l on the surface under study S whose x and y coordinates correspond to the row and column indexes of the pixels p_k and p_l on the image plane. An estimate of d_{s_k, s_l} may be recovered by making use of the increments along the pixel-sites falling on the path Γ_{p_k, p_l} that best describes the projection of the geodesic connecting the points on the surface onto the image or pixel lattice. To compute the quantity d_{s_k, s_l} , we traverse the path Γ_{p_k, p_l} and compute the height increments associated with the pixel-site transitions along the path. The quantity d_{s_k, s_l} is then given by the sum of these height increments. The approximation to the height increment associated with the transition from the pixel indexed $p_i \in \Gamma_{p_k, p_l}$ to the pixel indexed $p_j \in \Gamma_{p_k, p_l}$ can be computed by assuming that the two pixel sites are connected by a line whose slope is determined by the normal vectors $\vec{N}(s_i)$ and $\vec{N}(s_j)$. The height increment is given by

$$h_{s_i, s_j} = \frac{r_{p_i, p_j}}{2} \left\{ \frac{N_x(s_i)}{N_y(s_i)} + \frac{N_x(s_j)}{N_y(s_j)} \right\} \quad (15)$$

where $r_{p_i, p_j} = \|p_i - p_j\|$ is the Euclidean distance between the pixel centers on the image plane.




Fig. 9. Surface height recovery results. First row: Input images. Second and third rows: Recovered height and mean-squared error for our method. Fourth and fifth rows: Recovered height and mean-squared error for the method of Frankot and Chellappa. Sixth and seventh rows: Recovered height and mean-squared error for the method of Horn and Brooks.

B. Surface Integration via Unidimensional Embedding

The problem of recovering the surface from a set of pairwise height difference estimates may be viewed as one of unidimensional embedding subject to constraints provided by the footprint of the image lattice. We commence by rewriting the height difference estimate d_{s_k, s_l} between a pair of points on the surface S as

$$d_{s_k, s_l} = z_{s_k} - z_{s_l} + e_{s_k, s_l} \quad (16)$$

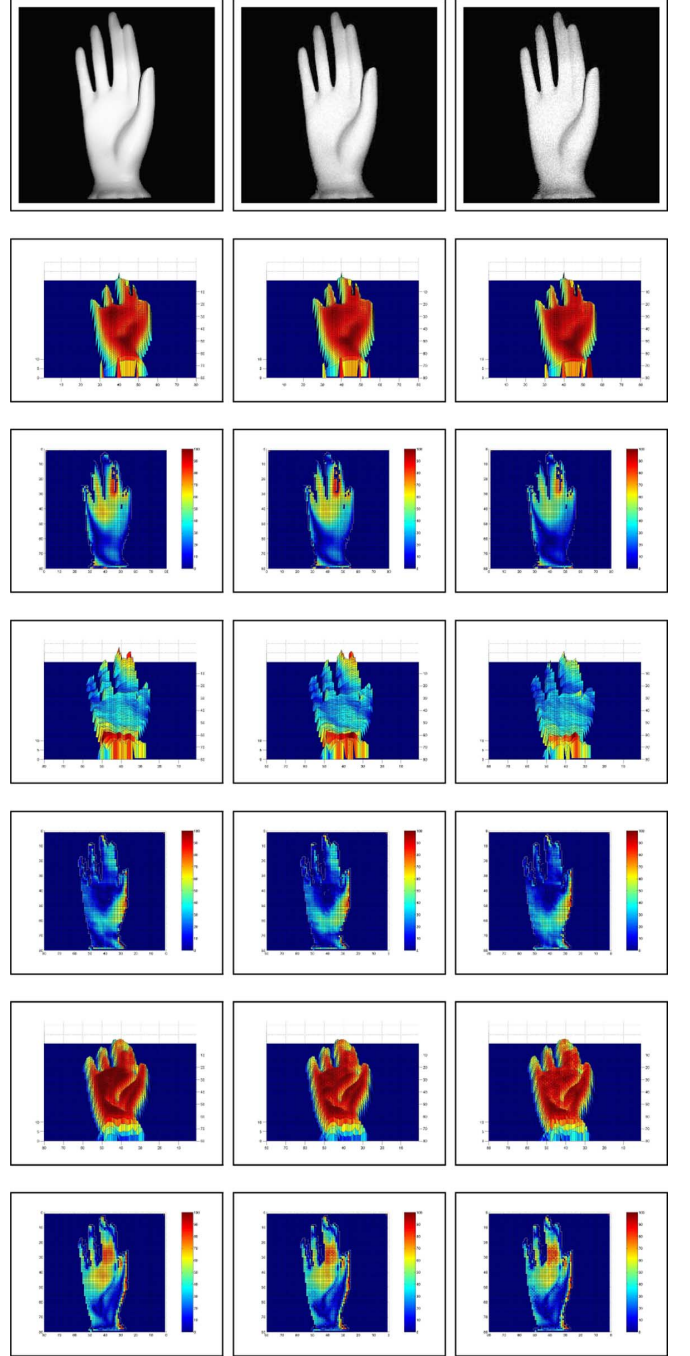


Fig. 10. Surface height recovery results. First row: input images. Second and third rows: Recovered height and mean-squared error for our method. Fourth and fifth rows: Recovered height and mean-squared error for the method of Frankot and Chellappa. Sixth and seventh rows: Recovered height and mean-squared error for the method of Horn and Brooks.

where z_{s_k} is the surface height at a point s_k corresponding to the pixel-site indexed p_k and e_{s_k, s_l} is the error of representation in the height difference.

Here, we assume that the errors of representation are independent random variables. Thus, if the image has M_{rows} rows and M_{cols} columns, then there are $M = |M_{\text{rows}} \times M_{\text{cols}}|$ random variables. Furthermore, the central limit theorem [23] states that, if $M \rightarrow \infty$, the distribution of normalized errors $\delta = e_{s_k, s_l}$ will

be Gaussian with mean μ and variance σ , i.e.,

$$P(\delta) = \frac{1}{\sigma\sqrt{2\pi}} \exp\left\{-\frac{(\delta - \mu)^2}{2\sigma^2}\right\}. \quad (17)$$

For the M samples, the error in the mean μ is

$$\epsilon = \frac{\sigma}{\sqrt{M}}. \quad (18)$$

As a result, as $M \rightarrow \infty$, the uncertainty of the estimator tends to zero.

Suppose that \mathbf{B} is the column-vector whose k th element is given by

$$\begin{aligned} \mathbf{B}_k &= \frac{1}{M} \sum_{l=1}^{l=M} d_{s_k, s_l} \\ &= \frac{1}{M} \sum_{l=1}^{l=M} (z_{s_k} - z_{s_l} + e_{s_k, s_l}) \\ &= z_{s_k} - \frac{1}{M} \sum_{l=1}^{l=M} z_{s_l} + \frac{1}{M} \sum_{l=1}^{l=M} e_{s_k, s_l}. \end{aligned} \quad (19)$$

The final term in the definition above is simply the mean height error at the location indexed k on the surface

$$\mu_k = \frac{1}{M} \sum_{l=1}^{l=M} e_{s_k, s_l}. \quad (20)$$

If we assume that the distribution of errors is stationary, then we can write

$$\mu_i = \frac{1}{M} \sum_{l=1}^{l=M} e_{s_i, s_l} \approx \mu_j = \frac{1}{M} \sum_{l=1}^{l=M} e_{s_j, s_l}. \quad (21)$$

As a result, and noting that the quantity $(1/M) \sum_{l=1}^{l=M} z_{s_l}$ in (19) does not depend on the index k , it follows that:

$$\begin{aligned} \mathbf{B}_i - \mathbf{B}_1 &= z_{s_i} + \frac{1}{M} \sum_{l=1}^{l=M} e_{s_i, s_l} - z_{s_1} - \frac{1}{M} \sum_{l=1}^{l=M} e_{s_1, s_l} \\ &= z_{s_i} - z_{s_1} + \mu_i - \mu_1 \approx z_{s_i} - z_{s_1}. \end{aligned} \quad (22)$$

Finally, if we set $z_{s_1} = 0$ and solve for z_{s_i} . The result is

$$z_{s_i} \approx \mathbf{B}_i - \mathbf{B}_1 = \frac{1}{M} \left(\sum_{l=1}^{l=M} d_{s_i, s_l} - \sum_{l=1}^{l=M} d_{s_1, s_l} \right). \quad (23)$$

It is important to stress that, whereas setting $z_{s_1} = 0$ may alter the position of the surface with respect to the coordinate system, the relative height configuration of each pair of points remains unchanged.

C. Approximation by Subsampling

The complexity of performing the embedding depends on the image dimension M . This may become burdensome for large images. To overcome this problem, for large images, we can

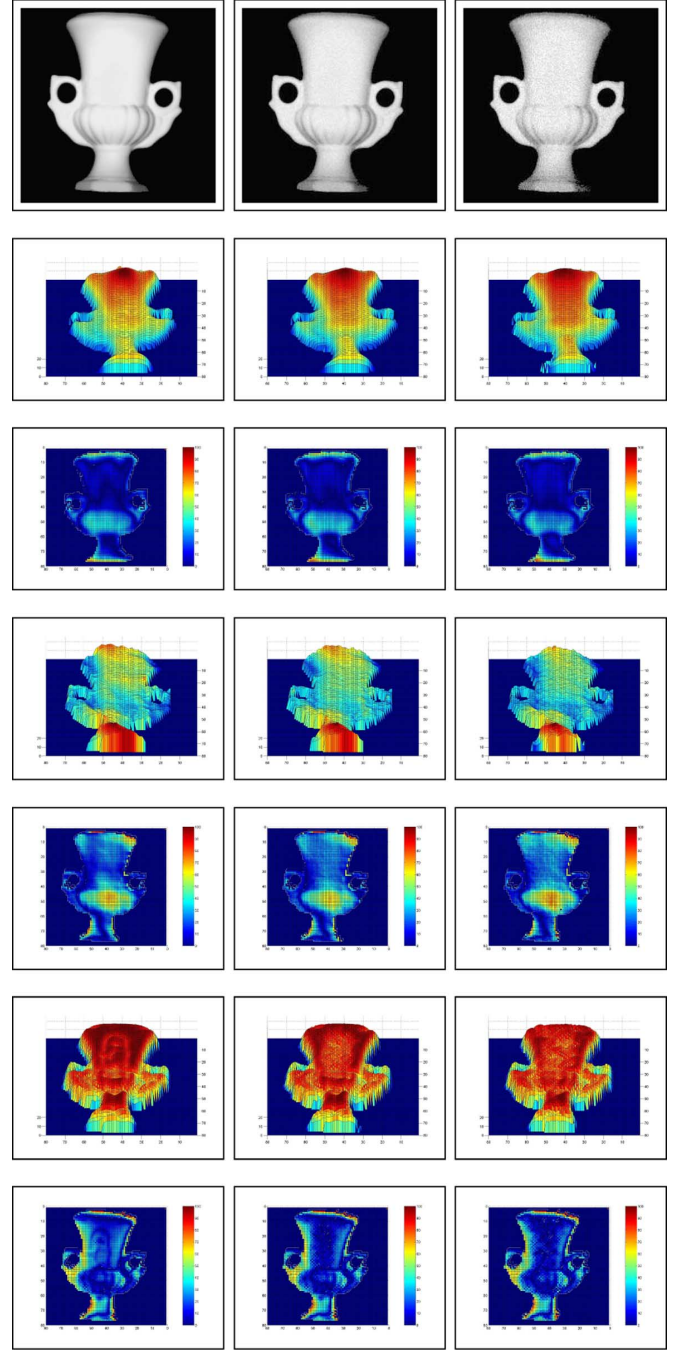


Fig. 11. Surface height recovery results. First row: input images. Second and third rows: Recovered height and mean-squared error for our method. Fourth and fifth rows: Recovered height and mean-squared error for the method of Frankot and Chellappa. Sixth and seventh rows: Recovered height and mean-squared error for the method of Horn and Brooks.

make use of the central limit theorem [23] to develop an accurate subsampling method. To do this, we commence by defining the k th coefficient of the vector $\hat{\mathbf{B}}$ as follows:

$$\begin{aligned} \hat{\mathbf{B}}_k &= \frac{1}{|\Phi|} \sum_{l \in |\Phi|} d_{s_k, s_l} \\ &= z_{s_k} - \frac{1}{|\Phi|} \sum_{l \in |\Phi|} z_{s_l} + \frac{1}{|\Phi|} \sum_{l \in |\Phi|} e_{s_k, s_l} \end{aligned} \quad (24)$$




Fig. 12. Surface height recovery results. First row: Input images. Second and third rows: Recovered height and mean-squared error for our method. Fourth and fifth rows: Recovered height and mean-squared error for the method of Frankot and Chellappa. Sixth and seventh rows: Recovered height and mean-squared error for the method of Horn and Brooks.

where Φ is the subset of sample pixel-sites used for computing the height difference estimates for each of the coefficients $\hat{\mathbf{B}}_k$. Since the length of the vectors $\hat{\mathbf{B}}$ and \mathbf{B} is the same, the resolution of the surface remains unchanged. If the probability distribution for z_{s_i} is Gaussian with variance $\bar{\sigma}$, then the surface height can be reconstructed using the formula

$$z_{s_i} \approx \hat{\mathbf{B}}_i - \hat{\mathbf{B}}_1 = \frac{1}{|\Phi|} \sum_{l \in \Phi} d_{p_i, p_l} - \frac{1}{|\Phi|} \sum_{l \in \Phi} d_{p_1, p_l}. \quad (25)$$

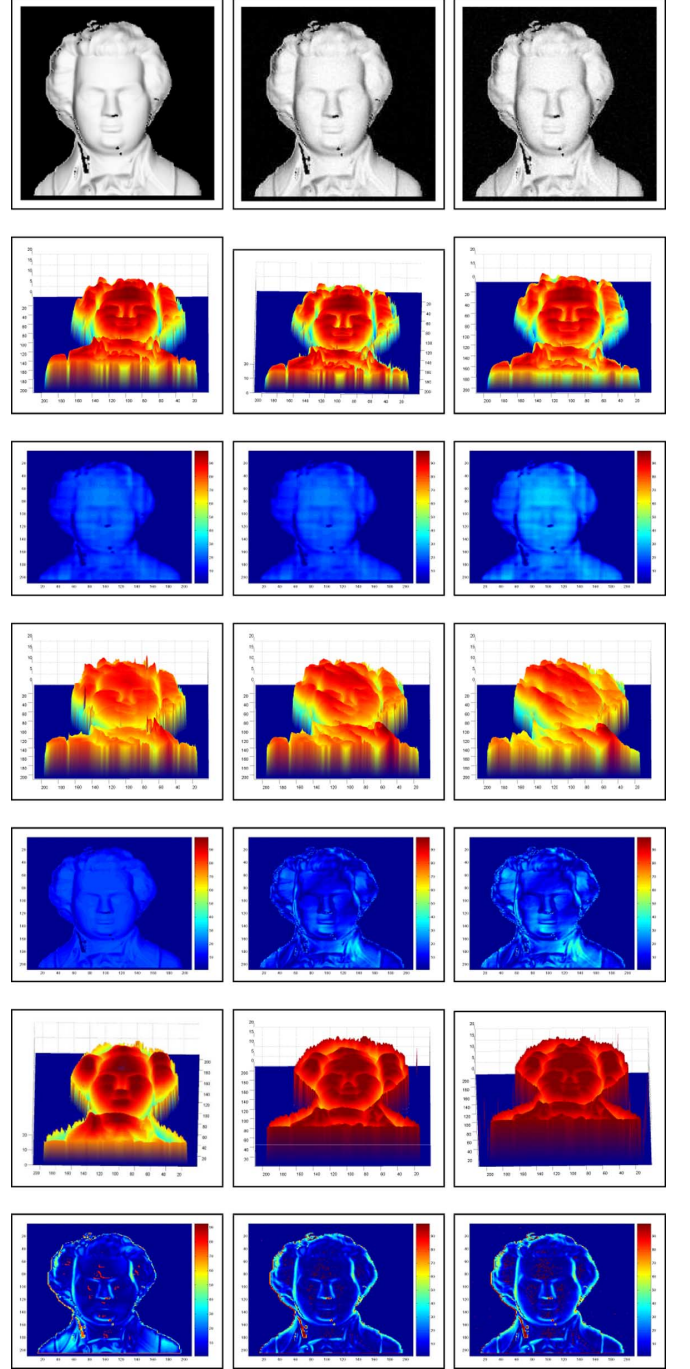


Fig. 13. Surface height recovery results. First row: input images. Second and third rows: Recovered height and mean-squared error for our method. Fourth and fifth rows: Recovered height and mean-squared error for the method of Frankot and Chellappa. Sixth and seventh rows: Recovered height and mean-squared error for the method of Horn and Brooks.

The precision of the estimate is

$$\bar{\epsilon} = \bar{\sigma} \left(1 - \frac{1}{\sqrt{|\Phi|}} \right). \quad (26)$$

In this way, the complexity may be greatly reduced and an interval of confidence can be established as a function of the number of samples $|\Phi|$.

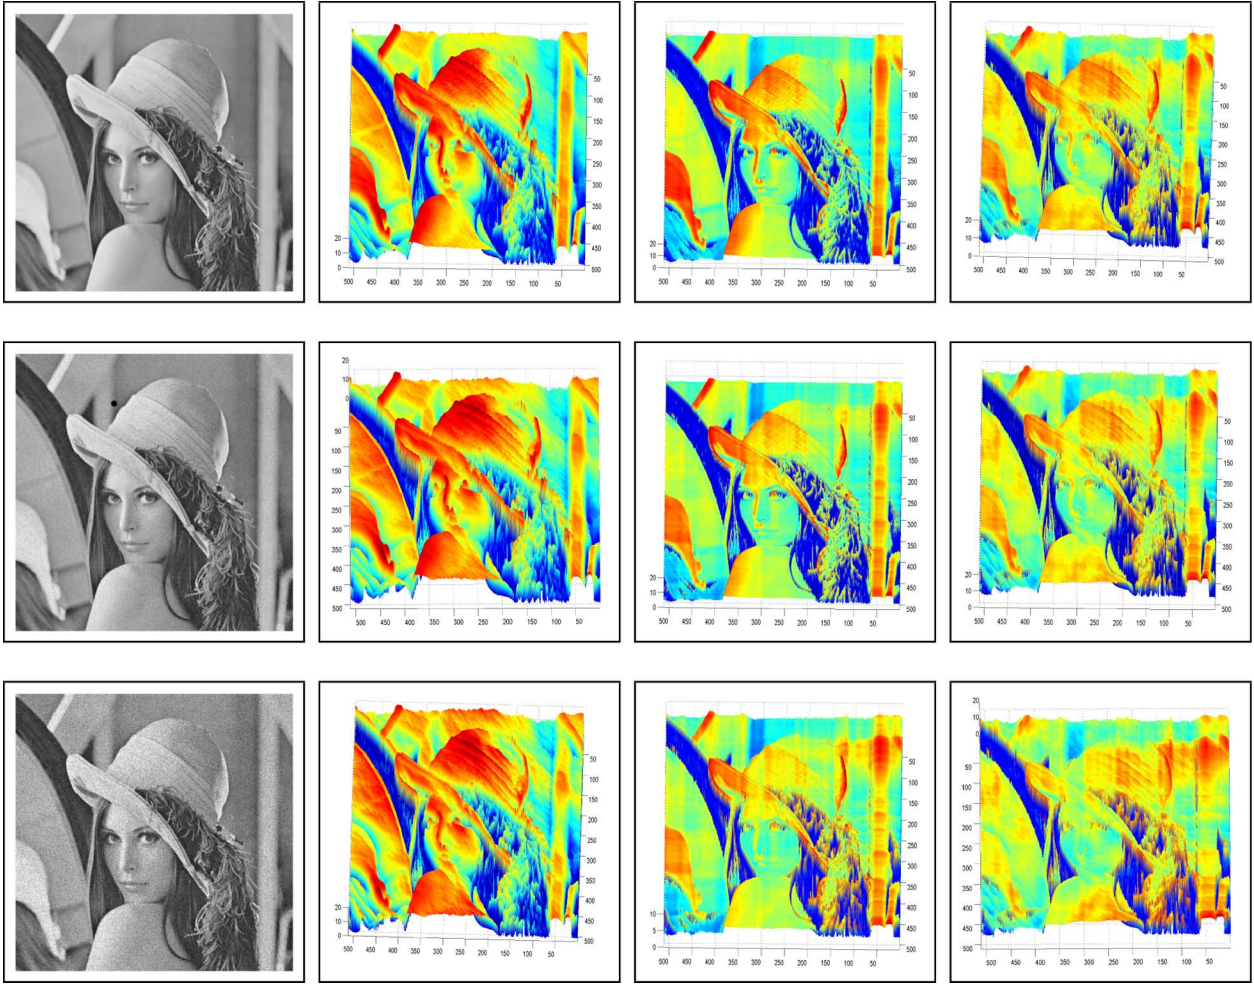


Fig. 14. Surface height recovery results. First column: Input images. Second column: Height recovered using our method. Third column: Recovered height for the method of Frankot and Chellappa. Fourth column: Recovered height for the method of Horn and Brooks.

V. EXPERIMENTS

In this section, we provide experiments with our new shape-from-shading method. The experimental study is divided into two parts. We commence by focussing on the performance of the heat-equation method for surface normal recovery. In the second part, we concentrate on the process of surface height recovery. The experimental study involves both, synthetic and real-world data.

A. Surface Normal Recovery

We commence our experimental study by illustrating the utility of the heat-equation method for recovering smoothed fields is surface normals. Here, we have experimented with images which exhibit significant levels of noise. To this end, we have added controlled levels of Gaussian noise to synthetic and real-world imagery of objects that exhibit Lambertian reflectance. In all our experiments, we follow Worthington and Hancock [12] and use the grayscale gradient as an initial estimate of the surface normal field directions.

In Fig. 2, we illustrate the behavior of the method on synthetic data. Here the left-most panel in the top row shows the synthetic surface used in our study. This consists of two parabolic domes superimposed on a parabolic ridge. The second panel

in the top row shows the result of rendering the surface using Lambert’s law. The third and fourth panels in the top row show the initial and final fields of surface normals delivered by the method. Initially, the field of surface normals is relatively disorganized, and appears noisy. The main effect to note here is that the heat-kernel smoothing method rotates the surface normal so that a consistent field emerges that reflects well the geometry of the underlying surface. In the bottom row of the figure, we show the heat function $Q(x, y, t)$ as a function of time t . Each plot displays the value of $Q(x, y, t)$ as a height value on the $x - y$ plane for fixed values of t , i.e., $t = 1, 3, 6, 9$. Here, convergence is reached after nine iterations. Therefore, in Fig. 2, we have chosen the values of t so as to be uniformly distributed from $t = 1$ to $t = 9$. The main elevation features of the surface emerge as peaks in the plots of $Q(x, y, t)$, and, as time evolves, they become sharpened.

We now turn our attention to real-world data. In Fig. 3, the top row of the figure shows images of four porcelain objects captured with an Olympus E10 digital camera. The objects are illuminated with a collimated tungsten light source with known direction. The original images contained specularities since the object surfaces are shiny. However, these have been removed using the method described in [24]. From left-to-right, the objects studied are a teapot, a hand, an urn and a model bear. In the




Fig. 15. Surface height error as a function of the variance for the objects in our dataset.

bottom line of the figure we show the ground truth data used in our study. This consists of depth images of the objects studied, captured using a Polhemus structured light range-sensor. The range data have been aligned with digital images using a simple registration algorithm.

In Figs. 4–7, we show the surface normal results obtained from the digital images. The top row of each figure shows noise corrupted versions of the input images. Here, we have added

Gaussian noise with zero mean and standard deviations 0.3 and 0.5 grey-levels to the raw image. In the second row of each figure we show the fields of surface normals obtained after ten and 100 iterations of the Worthington and Hancock algorithm. The fourth and fifth rows of the figures shows the surface normals at the initial and final (i.e., $n = 1, 10$) iterations of our heat-equation smoothing algorithm. The main features to note from the plots are that our new method gives fields of surface

normals that are both smoother and contain more fine surface detail than that of Worthington and Hancock. For instance, the surface structure of the hand, the ribs of the urn and the limbs of the bear all emerge more clearly when our method is used.

A more quantitative analysis of the results is shown in Fig. 8. Here, we show plots of the mean-squared error in the reconstructed surface normal directions as a function of the variance of the added Gaussian noise. For each object, we have computed the mean square error between the surface normal directions delivered by shape-from-shading method and those computed from the aligned ground truth range images. Fig. 8(a) is for the teapot, Fig. 8(b) is for the hand, Fig. 8(c) is for the urn, and Fig. 8(d) for the model bear. In each plot, the solid curve is the result of applying our method and the dotted curve that of applying the Worthington and Hancock method. The error bars show the standard error in the mean-squared errors. The features to note from the plots are as follows. First, the mean-squared error grows approximately linearly with the noise variance. Second, our new method gives errors which are consistently lower than those delivered by the Worthington and Hancock method.

B. Surface Height Recovery

Having examined the behavior of the heat-flow method for vector field smoothing, in this section we turn our attention to the effectiveness of our method for surface height recovery. We commence with experiments on synthetic data which are aimed at evaluating the systematics and noise sensitivity of the method, and then provide some experiments on real world data. We also present results on the “Mozart” and the “Lena” images. We have done this so as to provide the reader with results that can be used for the purposes of comparison with other work presented elsewhere in the literature. In all of our experiments, we have set $|\Phi| = 100$ so as to achieve a precision $\bar{\epsilon} = 0.9\bar{\sigma}$.

In Figs. 9–13, we investigate the effect of added image noise on the reconstructed height. In the top row of the figures, from left to right, we show the input images with zero added noise, and with added noise with zero mean and a standard deviation of 0.3 and 0.5 grey levels. In the second and third rows of each figure, we show the reconstructed height and the difference between the ground truth and reconstructed height using our method. In the fourth and fifth rows we show the corresponding results obtained using the Frankot and Chellappa [21] method, and the sixth and seventh rows those obtained using the Horn and Brooks [10] method. The main features to note from these plots are as follows. First, our method gives good reconstruction of the surface detail even under the highest level of noise. Second, the errors obtained with our method are confined to the boundaries of the objects, and the locations of sharp surface cusps. Third, the Horn and Brooks method recovers surfaces that show significant error over the entire object.

We show the results for the image of “Lena” in Fig. 14. Unfortunately, mean-squared error measures could not be provided since there is no ground truth available for this image. Nonetheless, we have investigated the effect of added image noise on the reconstructed height by adding noise with zero mean and known standard deviation. In the figure, the top row shows the results for the noise-free image. The second and third rows show the




Fig. 16. Surface height recovery results. First row: Input images. Second and third rows: Recovered height and mean-squared error for our method. Fourth and fifth rows: Recovered height and mean-squared error for the method of Frankot and Chellappa. Sixth and seventh rows: Recovered height and mean-squared error for the method of Horn and Brooks.

results for the noise-added imagery with standard deviation of 0.3 and 0.5, respectively. For each row, we show, from left to right, the input image, the recovered height for our algorithm and the depth maps for the algorithms of Frankot and Chellappa [21] and Horn and Brooks [10].

In Fig. 15, we repeat the noise sensitivity analysis for the reconstructed height data. The plots show the mean squared height error as a function of the noise variance for the objects studied in Figs. 9–13. In each plot, the dotted curve is the result obtained with our embedding method, the dashed curve that obtained using




Fig. 17. Surface height error as a function of the angle between the light source and the viewer directions for the teapot.

TABLE I
PROCESSING TIMES FOR THE SURFACE HEIGHT RECOVERY PROCESS

Object	Image size in pixels	Processing time in seconds (mean/var)		
		Our method	Frankot and Chellappa	Horn and Brooks
Porcelain vase	256 × 256	4.6/0.56	3.2/0.34	26/7.9
Porcelain bear	256 × 256	4.5/0.56	3.6/0.33	28/7.6
Porcelain hand	256 × 256	4.7/0.56	3.8/0.36	30/8.3
Teapot	480 × 480	7.4/0.67	4.5/0.456	68/11.2
Mozart	210 × 210	4.4/0.466	2.9/0.298	21/7.4

the Frankot and Chellappa’s method, and the solid curve is that obtained using the Horn and Brooks method. The height errors obtained using our method are consistently lower than those obtained with the two alternatives. It is interesting to note that the slope of the curves are very close, but that the error at zero noise is smallest for our method. This means that even when there is no added noise, the alternative methods result in significant surface reconstruction error. In other words they exhibit bias.

Next, we turn our attention to the effects of variation in light source direction. For this study we use the teapot. To this end, we have generated a set of six images illuminated with a single light source positioned in the direction $\vec{L} = [\sin(\theta_L), 0, \cos(\theta_L)]^T$. For our experiments, we have set $\vec{V} = [0, 0, 1]^T$ and varied the angle θ_L . In order to avoid self-shadowing effects, we have limited our attention to the values of the angle θ_L between 0° and 50° .

In the top row of Fig. 16, we show three of the images in our dataset. Here, we have ordered the example images, from left to right, in increasing θ_L , i.e., $\theta_L = 0^\circ, 20^\circ, 50^\circ$. The second and third rows in the figure show, again, the recovered height and, the difference between the recovered and the ground-truth heights. The results for the methods of Horn and Brooks and Frankot and Chellappa are shown in the bottom rows. The errors are largest at highly inclined locations on the surface. Whereas the effect of varying the light source direction is to magnify the errors for both methods, our algorithm is more robust and delivers better results. Fig. 17 provides a quantitative investigation of this effect. Here we plot the mean-squared height error for the three methods as a function of the angle between viewer and light source directions. In the left-hand panel, we show the mean-squared error when our heat flux method is used to recover the surface normal directions. The right-hand panel shows the error plots when the Worthington and Hancock method is used.

Again, the error for the alternatives is consistently larger than the one given by our method.

Finally, we present timing statistics for our method and the other two alternatives. All our experiments were performed on a Pentium 4, 3-GHz PC. It is worth noting that the implementations of our method and those of Horn and Brooks and Frankot and Chellappa are not fully optimized, and, hence, the timings shown below are provided as an illustration of the algorithms performance. The method of Horn and Brooks [10] hinges around solving a Poisson equation subject to the natural boundary condition. In this case, a solver based on the finite element method was used. In the method developed by Frankot and Chellappa [21], integrability is enforced making use of an optimisation process in the Fourier domain. Here, we have used a fast Fourier transform (FFT).

In Table I, we show the processing times for the results presented in Fig. 15. In the table, we show the mean and variance for the height recovery timings when processing the imagery with different levels of added noise. From the table, it is clear that the algorithm of Horn and Brooks is the one whose computational complexity is largest. This is due to the differential equation solver invoked by the method at runtime. Our algorithm is slightly more computationally intensive than the algorithm of Frankot and Chellappa, but much more efficient than the one of Horn and Brooks.

VI. CONCLUSION

In this paper, we have presented a novel approach for recovering the 3-D representation of the object from a single image making use of shading information. The contributions are twofold. First, we have shown how the heat equation can be used to recover the field of surface normals. To do this, we have introduced a scalar field that evolves with time under the heat equation. The z component of the surface normal is the normalized image brightness, and the azimuthal component of the surface normal is computed using the gradient of the scalar potential. We demonstrate how the scalar potential can be recovered by solving a simple difference equation using a computationally efficient finite difference method. This approach leads to a parameter-free model which constrains the surface normal to satisfy Lambert’s law. Second, we have

posed the problem of recovering the surface height as a that of embedding the field of surface normals into a manifold that resides in a Euclidean space. We have also shown how the complexity of the embedding process may be greatly reduced by a subsampling approach. We have performed experiments on synthetic and real-world imagery.

There are a number of ways in which the work reported in this paper may be further extended. First, we aim to explore how constraints from differential geometry can be incorporated into the heat-flow smoothing process. This is a problem that has been extensively studied in the context of Beltrami flows [25] and may allow us to incorporate information concerning surface topography. Second, we aim to explore more deeply the relationship between the scalar potential and the structure of the underlying surface. For instance, it would be interesting to explore the relationship between the divergence of the scalar potential and the surface curvature.

REFERENCES

- [1] E. Mingolla and J. Todd, "Perception of solid shape from shading," *Biol. Cybern.*, vol. 53, pp. 137–151, 1986.
- [2] P. N. Belhumeur and D. J. Kriegman, "What is the set of images of an object under all possible lighting conditions?," *Comput. Vis. Pattern Recognit.*, pp. 270–277, 1996.
- [3] M. Bichsel and A. P. Pentland, "A simple algorithm for shape from shading," in *Proc. IEEE Conf. Computer Vision and Pattern Recognition*, 1992, pp. 459–465.
- [4] R. Kimmel, K. Siddiqi, B. B. Kimia, and A. M. Bruckstein, "Shape from shading: Level set propagation and viscosity solutions," *Int. J. Comput. Vis.*, vol. 16, pp. 107–133, 1995.
- [5] A. M. Bruckstein, "On shape from shading," *Comput. Vis., Graph., Image Process.*, vol. 44, no. 2, pp. 139–154, 1988.
- [6] P. Dupuis and J. Oliensis, "Direct method for reconstructing shape from shading," in *Proc. IEEE Conf. Computer Vision and Pattern Recognition*, 1992, pp. 453–458.
- [7] A. Robles-Kelly and E. R. Hancock, "A graph-spectral approach to shape-from-shading," *IEEE Trans. Image Process.*, vol. 13, no. 7, pp. 912–926, Jul. 2004.
- [8] R. Zhang, P. S. Tsai, J. E. Cryer, and M. Shah, "Shape from shading: A survey," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 21, no. 8, pp. 690–706, Aug. 1999.
- [9] K. Ikeuchi and B. K. P. Horn, "Numerical shape from shading and occluding boundaries," *Artif. Intell.*, vol. 17, no. 1–3, pp. 141–184, 1981.
- [10] B. K. P. Horn and M. J. Brooks, "The variational approach to shape from shading," *CVGIP*, vol. 33, no. 2, pp. 174–208, 1986.
- [11] Q. Zheng and R. Chellappa, "Estimation of illuminant direction, albedo, and shape from shading," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 13, no. 7, pp. 680–702, Jul. 1991.
- [12] P. L. Worthington and E. R. Hancock, "New constraints on data-closeness and needle map consistency for shape-from-shading," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 21, no. 12, pp. 1250–1267, Dec. 1999.
- [13] S. Osher and J. Sethian, "Fronts propagating with curvature-dependent speed: Algorithms based on hamilton," *J. Comput. Phys.*, vol. 79, pp. 12–49, 1988.
- [14] J. D. Durou and H. Maitre, "On convergence in the methods of strat and of smith for shape from shading," *Int. J. Comput. Vis.*, vol. 17, no. 3, pp. 273–289, 1996.
- [15] E. Rouy and A. Tourin, "A viscosity solution approach to shape-from-shading," *SIAM J. Numer. Anal.*, vol. 29, no. 3, pp. 867–884, 1992.
- [16] E. Prados and O. Faugeras, "Perspective shape from shading and viscosity solutions," in *Proc. IEEE Int. Conf. Computer Vision*, 2003, pp. II:826–831.
- [17] R. Kimmel and A. M. Bruckstein, "Tracking level sets by level sets: A method for solving the shape from shading problem," *Comput. Vis. Image Understand.*, vol. 62, no. 2, pp. 47–48, Jul. 1995.
- [18] R. Kimmel, R. Malladi, and N. Sochen, "Images as embedded maps and minimal surfaces: Movies, color, texture, and volumetric medical images," *Int. J. Comput. Vis.*, vol. 39, no. 2, pp. 111–129, 2000.
- [19] B. K. P. Horn, "Height and gradient from shading," *Int. J. Comput. Vis.*, vol. 5, no. 1, pp. 37–75, 1990.
- [20] Z. Wu and L. Li, "A line-integration based method for depth recovery from surface normals," *CVGIP*, vol. 43, no. 1, pp. 53–66, Jul. 1988.
- [21] R. T. Frankot and R. Chellappa, "A method of enforcing integrability in shape from shading algorithms," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 4, no. 10, pp. 439–451, Oct. 1988.
- [22] T. Wei and R. Klette, "Theoretical analysis of finite difference algorithms for linear shape from shading," *CAIP*, pp. 638–645, 2001.
- [23] A. Papoulis, *Probability, Random Variables, and Stochastic Processes*. New York: Mc Graw-Hill, 1984.
- [24] A. Robles-Kelly and E. R. Hancock, "Estimating the surface radiance function from single images," *Graph. Models*, to be published.
- [25] N. Sochen, "Affine invariant flows via the beltrami framework," *J. Math. Imag. Vis.*, vol. 20, pp. 133–145, 2004.



Antonio Robles-Kelly (M'96) received the B.Eng. degree in electronics and telecommunications from the Instituto Tecnológico y de Estudios Superiores de Monterrey, Monterrey, Mexico (with honors) in 1998, and the Ph.D. degree in computer science from the University of York, York, U.K., in 2003.

In 2001, as a graduate student at the University of York, he visited the University of South Florida as part of the William Gibbs/Plessey Award for the best research proposal to visit an overseas research laboratory. After receiving his doctorate, he was a Research Associate under the MathFit-EPSCRC framework at the University of York until December 2004. Currently, he is a Research Scientist with National ICT Australia, Australian National University, Canberra Laboratory, Canberra, where he leads the Spectral Imaging and Source Mapping (SISM) project of the Vision Systems, Technologies, and Applications (ViSTA) programme. His research interests are in the areas of computer vision, pattern recognition, and computer graphics. Along these lines, he has done work on segmentation and grouping, graph matching, shape-from-X, hyperspectral image understanding, and reflectance models. He is also interested in the differential structure of surfaces.



Edwin Hancock received the degree in physics and the Ph.D. degree in high-energy physics from the University of Durham, Durham, U.K., in 1977 and 1981, respectively.

For ten years, he was a Researcher in the fields of high-energy nuclear physics and pattern recognition at the Rutherford-Appleton Laboratory (now the Central Research Laboratory of the Research Councils), Cheshire, U.K. During this period, he also held adjunct teaching posts at the University of Surrey, Surrey, U.K., and the Open University, Milton Keynes, U.K. In 1991, he joined the University of York, York, U.K., as a Lecturer in the Department of Computer Science. He was promoted to Senior Lecturer in 1997 and to Reader in 1998. In 1998, he was appointed to a Chair in Computer Vision. He now leads a group of some 15 faculty, research staff, and Ph.D. students working in the areas of computer vision and pattern recognition. He has been a Guest Editor for special issues of the journals *Image and Vision Computing* and *Pattern Recognition*. He has been a member of the Editorial Board for *Pattern Recognition*. He has published over 80 journal papers and 300 refereed conference publications. He has been on the program committees for numerous national and international meetings. In 1997, with Marcello Pelillo, he established a new series of international meetings on energy minimization methods in computer vision and pattern recognition. His main research interests are in the use of optimization and probabilistic methods for high- and intermediate-level vision. He is also interested in the methodology of structural and statistical pattern recognition. He is currently working on graph matching, shape-from-X, image databases, and statistical learning theory. His work has found applications in areas such as radar terrain analysis, seismic section analysis, remote sensing, and medical imaging.

Prof. Hancock became a Fellow of the International Association for Pattern Recognition in 1998. He has been a member of the Editorial Board of the IEEE TRANSACTIONS ON PATTERN ANALYSIS AND MACHINE INTELLIGENCE. He was awarded the Pattern Recognition Society medal in 1991 and an Outstanding Paper Award in 1997 by the journal *Pattern Recognition*.