



ELSEVIER

Available at  
www.ComputerScienceWeb.com  
POWERED BY SCIENCE @ DIRECT®

Pattern Recognition  
Letters

Pattern Recognition Letters 24 (2003) 579–595

www.elsevier.com/locate/patrec

# Darboux smoothing for shape-from-shading

Hossein Ragheb <sup>\*,1</sup>, Edwin R. Hancock

*Department of Computer Science, University of York, York YO1 5DD, UK*

Received 11 October 2001; received in revised form 6 June 2002

## Abstract

This paper describes a new surface normal smoothing process which can be used in conjunction with shape-from-shading. Rather than directly smoothing the surface normal vectors, we exert control over their directions by smoothing the field of principal curvature vectors. To do this we develop a topography sensitive smoothing process which overcomes the problems of singularities in the field of principal curvature directions at the locations of umbilics and saddles. The method is evaluated on both synthetic and real-world images.

© 2002 Elsevier Science B.V. All rights reserved.

*Keywords:* Shape-from-shading; Darboux smoothing; Surface normal; Principal curvature; Lambertian reflectance

## 1. Introduction

Shape-from-shading is concerned with recovering surface orientation from local variations in measured brightness. There is strong psychophysical evidence for its role in surface perception and recognition (Belhumeur and Kriegman, 1996; Koenderink et al., 1992; Koenderink and van Doorn, 1979). However, despite considerable effort over the past two decades, reliable shape-from-shading recovery has proved to be an elusive goal (Horn and Brooks, 1986; Brooks and Horn, 1985). The reasons for this are two-fold. Firstly, the recovery of surface orientation from the image

irradiance equation is an under-constrained process which requires the provision of boundary conditions and constraints on surface smoothness to be rendered tractable. Secondly, real-world imagery rarely satisfies these constraints. Several authors have attempted to develop shape-from-shading methods which overcome these shortcomings. For instance, Oliensis and Dupuis (1993) and Bichsel and Pentland (1992) have developed solutions for which shape-from-shading is not under-constrained, but which require prior knowledge of the heights of singular points of the surface. Meanwhile, Kimmel and Brookstein have shown how the apparatus of level-set theory can be used to solve the image irradiance equation as a boundary value problem (Bruckstein, 1988; Kimmel and Bruckstein, 1995). The paper by Zhang et al. (1999) provides a comprehensive review of the literature on shape-from-shading and also performs a detailed comparative study of the available methods.

\* Corresponding author. Tel.: +44-1904-432795; fax: +44-1904-432767.

E-mail addresses: [hossein@cs.york.ac.uk](mailto:hossein@cs.york.ac.uk) (H. Ragheb), [erh@cs.york.ac.uk](mailto:erh@cs.york.ac.uk) (E.R. Hancock).

<sup>1</sup> Sponsored by the university of Bu-Ali Sina, Hamedan, Iran.

However, in general shape-from-shading has suffered from the dual problems of model dominance (which results in oversmoothing of the recovered surface) and poor data-closeness. Recently, Worthington and Hancock have developed a new shape-from-shading scheme which has gone some way to overcoming some of the shortcomings of existing, iterative shape-from-shading algorithms (Worthington and Hancock, 1999). Specifically, they have shown that the image irradiance equation can be treated as a hard constraint (Worthington and Hancock, 1999). The idea is a simple one. For a Lambertian surface, the image irradiance equation constrains the local surface normal to fall on a cone. The axis of this cone points in the light-source direction while the cosine of the apex angle of the cone is proportional to the local image brightness. Worthington and Hancock exploit this property to develop a geometric method for iteratively updating the set of local surface normal directions. At each image location the surface normal is initially positioned on the irradiance cone so that its projection onto the image plane points in the direction of the local image gradient. The field of surface normal directions is then subjected to a smoothing process which is sensitive to the local topographic surface structure. Smoothing is effected using robust error kernels whose width is governed by the variance of the local shape-index (an angular measure of local surface topography first suggested by Koenderink and van Doorn (1992)). This smoothing process results in surface normals which violate the image irradiance equation, i.e. do not fall on the irradiance cone. To restore data-closeness, the smoothed surface normals are projected onto the closest position on the cone. This process is iterated to convergence.

The observation underpinning this paper is that although the use of topographic information improves the pattern of surface normals, it overlooks a considerable body of information concerning the differential structure of the underlying surface residing in the local Darboux frame. At each point on the surface the  $z$ -axis of the Darboux frame points in the direction of the local surface normal, while the  $x$  and  $y$  axes are aligned in the directions of the local maximum and minimum surface curvatures. There have been several attempts to

improve the consistency of recovered surface information by smoothing Darboux frames. For instance Sander and Zucker (1990) have incorporated a least-squares process for smoothing the principal curvature directions in the inference of surfaces from 3D position data. This smoothing process has been applied as a post-processing step in shape-from-shading by Ferrie and Lagarde (1990). However, there are limitations to the use of least-squares smoothing of the Darboux frames. These relate to fact that the principal curvature direction is undefined at umbilic surface points and that the direction is discontinuous at saddles.

Our aim in this paper is to refine the information returned by shape-from-shading, not by smoothing the field of surface normal directions, but by smoothing the field of principal curvature directions. There are two novel contributions: First, to overcome the problems encountered at umbilics (the centres of cups and domes) and saddles, we develop a topography sensitive weighting process. Second, we incorporate the smoothing process into the Worthington and Hancock (1999) framework which ensures that the recovered surface normals satisfy the data-closeness constraints imposed by the image irradiance equation. By smoothing in this way we avoid using surface normal directions in parabolic surface regions, i.e. at the locations of ridges and ravines. These topographic features are associated with rapid rates of change in surface normal direction and, hence, rapid variations in the local image intensity (i.e. they are locations where the image gradient is large). Since the change in principal curvature direction is smaller than the change in surface normal direction at such features, Darboux smoothing may moderate the over-smoothing of ridge and ravine topography. Experiments on synthetic and real-world data reveal that the method delivers surface orientation information that is both accurate and stable.

## 2. Shape-from-shading

The shape-from-shading algorithm of Worthington and Hancock has been demonstrated to

deliver needle-maps which preserve fine surface detail (Worthington and Hancock, 1999). The observation underpinning the method is that for Lambertian reflectance from a matte surface, the image irradiance equation defines a cone of possible surface normal directions. The axis of this cone points in the light-source direction and the opening angle is determined by the measured brightness. If the recovered needle-map is to satisfy the image irradiance equation as a hard constraint, then the surface normals must each fall on their respective reflectance cones. Initially, the surface normals are positioned so that their projections onto the image plane point in the direction of the image gradient. Subsequently there is iterative adjustment of the surface normal directions so as to improve the consistency of the needle-map. In other words, each surface normal is free to rotate about its reflectance cone in such a way as to improve its consistency with its neighbours. This rotation is a two-step process. First, we apply a smoothing process to the current surface normal estimates. This may be done in a number of ways. The simplest is local averaging. More sophisticated alternatives include robust smoothing with outlier reject and, smoothing with curvature or image gradient consistency constraints. This results in an off-cone direction for the surface normal. The hard data-closeness constraint of the image irradiance equation is restored by projecting the smoothed off-cone surface normal back onto the nearest position on the reflectance cone.

To be more formal let  $\vec{L}$  be a unit vector in the light-source direction and let  $E(i, j)$  be the brightness at the image location  $(i, j)$ . Further, suppose that  $\vec{N}^k(i, j)$  is the corresponding estimate of the surface normal at iteration  $k$  of the algorithm. The image irradiance equation is

$$E(i, j) = \vec{N}^k(i, j) \cdot \vec{L} \quad (1)$$

As a result, the reflectance cone has opening angle  $\cos^{-1}(E(i, j))$ . After local smoothing, the off-cone surface normal is  $\vec{N}_S^k(i, j)$ . The updated on-cone surface normal which satisfies the image irradiance equation as a hard constraint is obtained via the rotation

$$\vec{N}^{k+1}(i, j) = \Phi \vec{N}_S^k(i, j) \quad (2)$$




Fig. 1. Computing the updated surface normal direction by rotating the smoothed surface normal direction to the closest position on the irradiance cone.

The matrix  $\Phi$  rotates the smoothed off-cone surface normal estimate by the angle difference between the apex angle of the cone, and the angle subtended between the off-cone normal and the light-source direction. This angle is

$$\theta = \cos^{-1}(E(i, j)) - \cos^{-1}(\vec{N}_S^k(i, j) \cdot \vec{L}) \quad (3)$$

This rotation (Fig. 1) takes place about the axis whose direction is given by the vector

$$(u, v, w)^T = \vec{N}_S^k(i, j) \times \vec{L} \quad (4)$$

where  $\times$  denotes the vector cross product operator. The axis of rotation is perpendicular to both the light-source direction and the off-cone normal. Hence, the rotation matrix is

$$\Phi = \begin{pmatrix} c + u^2c' & -ws + uvc' & vs + uwc' \\ ws + uvc' & c + v^2c' & -us + vvc' \\ -vs + uvc' & us + vvc' & c + w^2c' \end{pmatrix}$$

where  $c = \cos \theta$ ,  $c' = 1 - c$ , and  $s = \sin \theta$ .

### 3. Local Darboux frames

Once estimates of the surface normal directions are to hand, we can represent the local surface structure using a Darboux frame. To compute the Darboux frame, we commence from the Hessian matrix

$$H = \begin{bmatrix} \frac{\partial}{\partial x}(\vec{N}^k)_x & \frac{\partial}{\partial x}(\vec{N}^k)_y \\ \frac{\partial}{\partial y}(\vec{N}^k)_x & \frac{\partial}{\partial y}(\vec{N}^k)_y \end{bmatrix} = \begin{bmatrix} h_{11} & h_{12} \\ h_{21} & h_{22} \end{bmatrix} \quad (5)$$

where for simplicity we have suppressed the coordinate suffixes on the surface normals. In practice, we approximate the Hessian matrix using directional first derivatives computed from the current estimates of the surface normals. At the pixel with row and column indices  $(i, j)$ , the estimate of the Hessian is

$$H = \begin{bmatrix} \frac{(\vec{N}_{i+1,j}^k)_x - (\vec{N}_{i,j}^k)_x}{(x_{i+1,j}) - (x_{i,j})} & \frac{(\vec{N}_{i+1,j}^k)_y - (\vec{N}_{i,j}^k)_y}{(y_{i+1,j}) - (y_{i,j})} \\ \frac{(\vec{N}_{i,j+1}^k)_x - (\vec{N}_{i,j}^k)_x}{(x_{i,j+1}) - (x_{i,j})} & \frac{(\vec{N}_{i,j+1}^k)_y - (\vec{N}_{i,j}^k)_y}{(y_{i,j+1}) - (y_{i,j})} \end{bmatrix} \quad (6)$$

where  $(x_{i,j}, y_{i,j})$  are the co-ordinates of the pixel  $(i, j)$ .

The two eigenvalues of the Hessian matrix are the maximum and minimum curvatures (Woodham, 1994):

$$\lambda_M^k = -\frac{1}{2}(h_{11} + h_{22} + S)$$

and

$$\lambda_m^k = -\frac{1}{2}(h_{11} + h_{22} - S) \quad (7)$$

where  $S^2 = (h_{11} - h_{22})^2 + 4(h_{12}h_{21})$ .

The integrability condition ( $h_{12} = h_{21}$ ) is frequently used as an additional smoothness constraint in shape-from-shading to render the recovery of surface height tractable (Frankot and Chellappa, 1988). The estimate of the Hessian given in Eq. (6) does not in general result in values of the quantities  $h_{12}$  and  $h_{21}$  which are identical. Hence, we do not enforce integrability at every point on the surface. However, if the quantity  $S^2$  is negative (i.e.  $S$  is undefined) then surface integrability is enforced using the equation  $h_{12} = h_{21} = (h_{12} + h_{21})/2$  which results in a positive value for  $h_{12}h_{21}$  and hence a real value for  $S$ . When this is the case, then the two curvatures  $\lambda_M^k$  and  $\lambda_m^k$  are also guaranteed to be real-valued.

On the tangent plane to the surface, the eigenvectors associated with the two curvatures are given by the two-component vectors

$$\vec{M}_2^k = \begin{cases} \left(\frac{1}{2}(h_{11} - h_{22} + S), h_{21}\right)^T & h_{11} \geq h_{22} \\ \left(-h_{12}, \frac{1}{2}(h_{11} - h_{22} - S)\right)^T & h_{11} < h_{22} \end{cases} \quad (8)$$

$$\vec{m}_2^k = \begin{cases} \left(h_{12}, -\frac{1}{2}(h_{11} - h_{22} + S)\right)^T & h_{11} \geq h_{22} \\ \left(\frac{1}{2}(h_{11} - h_{22} - S), h_{21}\right)^T & h_{11} < h_{22} \end{cases} \quad (9)$$

We compute the eigenvector associated with the maximum curvature  $\lambda_M^k$ . This is referred to as the principal curvature direction. From the known direction of the surface normal and the direction of the principal curvature direction on the tangent plane, we compute the three-component unit-vector  $\vec{M}^k$  in the image co-ordinate system. The direction of the surface normal together with the principal curvature direction define a local frame of reference referred to as a Darboux frame.

From the two eigenvalues of the Hessian, we can also compute the shape-index of Koenderink and van Doorn (1992). The quantity is angular in nature and is defined to be

$$\phi^k = \frac{2}{\pi} \arctan \left( \frac{\lambda_m^k + \lambda_M^k}{\lambda_m^k - \lambda_M^k} \right) \quad \lambda_M^k \geq \lambda_m^k \quad (10)$$

The shape-index varies smoothly in the range between  $-1$  and  $+1$  as the topography of the surface changes. The shape-index does not depend on the magnitude of the surface curvature. To capture the degree of bending or magnitude of the curvature, the curvedness  $K = \sqrt{(\lambda_m^k)^2 + (\lambda_M^k)^2}$  is used. In our experiments, we use this measure to visualise the whereabouts of curved surface features.

The relationship between the shape-index and the topographic class of the underlying surface is summarised in Table 1. The table lists the topographic classes (i.e. dome, ridge, saddle ridge etc.) and the corresponding shape-index interval. We visualise surface topography by mapping the shape-index to a grey-scale value. As we move from top to bottom in the table, then the associated grey-scale value decreases.

Table 1  
Topographic classes

Class	Region-type	Shape-index
Dome	Elliptic	[5/8, 1)
Ridge	Parabolic	[3/8, 5/8)
Saddle ridge	Hyperbolic	[1/8, 3/8)
Plane	Hyperbolic	Undefined
Saddle-point	Hyperbolic	[-1/8, 1/8)
Saddle-rut	Hyperbolic	[-3/8, -1/8)
Rut	Parabolic	[-5/8, -3/8)
Cup	Elliptic	[-5/8, -1)

#### 4. Smoothing principal curvature directions

We would like to smooth the needle-map using the field of principal curvature directions. However, this is not a straightforward task. The reason for this is that certain classes of local surface topography are characterised by rapid changes in principal curvature direction. There are two important cases. Firstly, for elliptic structures, i.e. domes and cups, there is a singularity since the principal curvature is not defined at the associated umbilic or spherical point, and the surrounding vector field is radial in structure. Secondly, for hyperbolic structures, i.e. saddles, there is a discontinuity in the principal curvature direction. This is due to the fact that the directions of maximum and minimum curvature flip as the saddle is traversed.

To overcome these problems we weight against domes and saddles in the smoothing process. To do this we note that the shape-index associated with parabolic structures is  $\pm 1/2$ . We therefore use the topographic weight

$$w_{i,j}^{(k)} = \exp(-\mu(|\phi_{i,j}^k| - \frac{1}{2})^2) \tag{11}$$

where  $\mu$  is a constant. The weight is used to compute the smoothed principal curvature direction

$$\vec{M}_S^k(i,j) = \sum_{(l,m) \in R_{ij}} w_{l,m}^{(k)} \vec{M}^k(l,m) \tag{12}$$

over the neighbourhood  $R_{i,j}$  of the pixel indexed  $(i,j)$ .

Our aim is to use the smoothed principal curvature direction to compute an updated Darboux frame.

#### 5. Surface normal adjustment

Once the smoothing of the field of principal curvature directions is complete, then we can use their direction vectors ( $\vec{M}^k$ ) to compute revised estimates of the surface normal directions ( $\vec{N}_D^k$ ). This is a two-step process. First, rotate the unsmoothed surface normals around the irradiance cone in a manner which is consistent with the rotation of the principal curvature directions which results from the Darboux smoothing process. To do this, we project the smoothed and unsmoothed principal curvature directions onto the plane which is perpendicular to the axis of the irradiance cone (i.e. the light-source direction). Once these projections are to hand, then we compute their angle-difference on the plane. We rotate the on-cone surface normal around the cone by an amount equal to this angle. The second step is to restore the orthogonality of the Darboux frame by rotating the surface normal away from the irradiance cone so that it is perpendicular to the smoothed principal curvature direction. The procedure adopted is illustrated in Figs. 2 and 3.




Fig. 2. Rotating the Darboux frame around the light-source direction: (a) rotating the current surface normal around the cone; (b) finding the rotation angle using the principal curvature directions.




Fig. 3. Computing the updated surface normal direction: (a) finding the direction of the local surface normal to the rotated Darboux frame; (b) rotating the closest surface normal onto the irradiance cone.

### 5.1. Rotation about the irradiance cone

To be more formal, the projection of the unsmoothed principal curvature direction onto the plane perpendicular to the axis of the irradiance cone is

$$\vec{M}_P^k = \vec{L} \times \vec{M}^k \times \vec{L} \quad (13)$$

while the projection of the smoothed principal curvature direction is

$$\vec{M}_{SP}^k = \vec{L} \times \vec{M}_S^k \times \vec{L} \quad (14)$$

On the plane perpendicular to the irradiance cone, the unsigned angle between the projected principal curvature directions is

$$|\alpha^k| = \cos^{-1}(\vec{M}_P^k \cdot \vec{M}_{SP}^k) \quad (15)$$

To perform the rotation of the on-cone surface normal in a consistent manner, we require the sign or sense of this angle. Because the vectors  $\vec{M}_P^k$  and  $\vec{M}_{SP}^k$  are both perpendicular to the light-source direction  $\vec{L}$ , the cross product  $\vec{M}_P^k \times \vec{M}_{SP}^k$  is either in the direction of the light-source direction  $\vec{L}$  or in the opposite direction. So, we can compute the signed rotation angle  $\alpha^k$  using the following equation:

$$\alpha^k = \overbrace{(\vec{L} \cdot (\vec{M}_P^k \times \vec{M}_{SP}^k))}^{\text{sign}} |\alpha^k| \quad (16)$$

We rotate the current surface normal around the irradiance cone defined by the light-source di-

rection  $\vec{L}$  and the apex angle  $\cos^{-1}(E)$  by the angle  $\alpha^k$ . The updated surface normal direction is given by the formula

$$\begin{aligned} \vec{N}_R^k = \vec{N}^k \cos(\alpha^k) + (\vec{N}^k \cdot \vec{L})\vec{L}(1 - \cos(\alpha^k)) \\ + (\vec{L} \times \vec{N}^k) \sin(\alpha^k) \end{aligned} \quad (17)$$

The rotation is illustrated in Fig. 2.

The surface normals  $\vec{N}_R^k$  computed in this way are not necessarily perpendicular to the corresponding smoothed principal curvature directions.

### 5.2. Computing the Darboux surface normal direction

The surface normals computed using the method outlined in the previous subsection are no longer perpendicular to the two principal curvature directions. We must therefore provide a means of restoring the geometry of the Darboux frame.

To do this, we use the field of rotated surface normals  $\vec{N}_R^k$  to compute an estimate of the principal curvature direction, which we denote by  $\vec{M}_R^k$  (Fig. 3a). The vector that is perpendicular to  $\vec{M}_R^k$  and the smoothed principal curvature direction  $\vec{M}_S^k$  provides an average surface normal. However, we must test the sign of the z-component of the resulting vector to ensure that the handedness of the co-ordinate system does not flip. The surface normal used is

$$\vec{N}_D^k = \begin{cases} \vec{M}_S^K \times \vec{M}_R^K & \text{if } (\vec{M}_S^K \times \vec{M}_R^K)_z \geq 0 \\ \vec{M}_R^K \times \vec{M}_S^K & \text{otherwise} \end{cases} \quad (18)$$

5.3. Imposing the irradiance constraint

To restore the irradiance constraint we can clearly apply the procedure of Worthington and Hancock and project the normals onto the nearest position on the irradiance cone. However, there are situations where the principal curvature directions are unavailable. For instance, the principal curvature direction is not defined at the location of umbilics. Also, it can not be computed when the Hessian matrix is singular.

There is a further reason to use smoothed surface normals  $\vec{N}_S^k$  in conjunction with Darboux surface normals  $\vec{N}_D^k$ . The reason for this is that there are certain surface regions for which the Darboux surface normals  $\vec{N}_D^k$  may be more accurate and may be closer to the irradiance cone than the smoothed surface normals  $\vec{N}_S^k$ . The cases in question are (a) locations where image intensity gradients are large; (b) at the locations of ridges and ruts. On the other hand, the smoothed surface normals  $\vec{N}_S^k$  may be more accurate when (a) where the intensity gradients are small and (b) at the locations of domes.

To overcome these problems, we have used a switching process. For each location in the image we have computed two surface normal directions. The first of these is the surface normal  $\vec{N}_D^k$  computed from the principal curvature direction. The second is the surface normal  $\vec{N}_S^k$  found by applying the robust smoothing method of Worthington and Hancock to the field of current surface normal estimates ( $\vec{N}^k$ ). From these two alternatives, we select the surface normal which is closest to the irradiance cone. In other words, the surface normal used for projection onto the irradiance cone is

$$\vec{N}_C^k = \begin{cases} \vec{N}_D^k & \text{if } |\cos^{-1}(\vec{N}_D^k \cdot \vec{L}) - \cos^{-1}(E)| \\ & \leq |\cos^{-1}(\vec{N}_S^k \cdot \vec{L}) - \cos^{-1}(E)| \\ \vec{N}_S^k & \text{otherwise} \end{cases} \quad (19)$$

With the closest surface normal to hand, we can effect projection onto the irradiance cone. The rotation angle required for the projection is

$$\beta^k = \cos^{-1}(\vec{N}_C^k \cdot \vec{L}) - \cos^{-1}(E) \quad (20)$$

After projection the updated surface normal is

$$\vec{N}^{k+1} = \vec{N}_C^k \cos(\beta^k) + \left( \left( \frac{\vec{N}_C^k \times \vec{L}}{\|\vec{N}_C^k \times \vec{L}\|} \right) \times \vec{N}_C^k \right) \sin(\beta^k) \quad (21)$$

The geometry of this procedure is described in Fig. 3b. Furthermore, we provide a flow chart for the

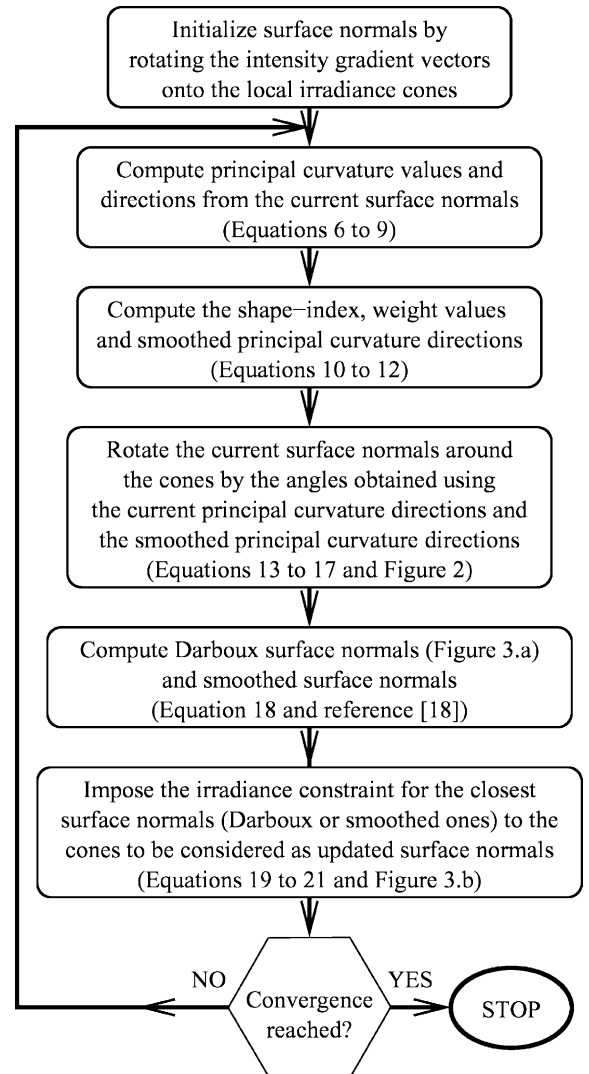


Fig. 4. Steps of the new shape-from-shading algorithm (see Sections 2–5 for details).

steps of the new Darboux smoothing shape-from-shading algorithm in Fig. 4.

## 6. Experiments

We have experimented with our new shape-from-shading method on synthetic and real-world imagery. The experimentation with synthetic data is designed to measure the effectiveness of the method on surfaces which contain a variety of topographic structure. Since our method relies on well defined principal curvature directions, we may anticipate that the main advantages of our method are most noticeable at ridge or ravine structures. We have therefore conducted some experiments on real and synthetic data aimed at evaluating our Darboux smoothing process at such structures and comparing its performance with the method described by Worthington and Hancock (1999).

To visualise the shape-index  $\phi$  (computed using Eq. (10)) for the field of surface normal directions obtained using shape-from-shading, the quantity needs to be rescaled from the interval  $[-1, +1]$  to the grey-scale interval  $[0, 255]$ . Hence, we display the quantity  $255(1 + \phi)/2$  at each image location. To aid visualisation, we also scale the curvedness values (which are always positive) to fall in the grey-scale interval  $[0, 255]$ .

We have compared our Darboux smoothing method with that of Worthington and Hancock (1999). This method does not make use of the principal curvature directions. However, the principal curvature directions can easily be computed from the field of surface normals using the method outlined in Section 3.

In our experiments, we provide comparison by computing the difference in surface normal directions. These are displayed as fields of differences in unit surface normal directions. The vectors in the resulting vector fields show the direction and magnitude of the vector difference in the surface normals. We also visualise the differences in surface normal directions alone by displaying the dot-product between the unit surface normals (this quantity takes on its maximum value of unity when the directions of the surface normal are identical).

The dot-products are mapped onto the grey-scale interval, and hence the brighter the pixel, the smaller the directional error.

### 6.1. Synthetic data

Figs. 5–10 show the results obtained with synthetic surfaces. In Figs. 5–7, in the top row, we show the elevation data for the surface (left panel) and a rendering of the local surface orientation information (right panel). The second, third and fourth rows respectively show ground-truth, the results obtained with the Worthington and Hancock method, and, those obtained with our new Darboux smoothing method. In these three rows, the left column shows the surface normal data, the second column shows the field of minimum curvature directions, the third column is the total surface shape-index.

The surface studied in Fig. 5 is obtained by placing four elliptical domes on a parabolic ridge. In Fig. 6 the surface consists of three narrow paraboloid ridges placed on top of perpendicular wide ridge. The surface in Fig. 7 is similar to that in Fig. 6, except that the four surfaces have their maxima at the same height.

For the surface shown in Fig. 5, the minimum curvature vectors radiate from the four domes, and is parallel to the crest line of the parabolic ridge. There should be a sharp discontinuity in direction where the domes intersect the ridge. Both smoothing methods blur this transition region. However, in the case of our new Darboux smoothing method the blurring is less pronounced. This feature emerges more clearly in the shape-index images, where the blurring associated with the Worthington and Hancock method is very strong. By contrast the Darboux smoothing method produces results which are in good agreement with ground-truth.

In Figs. 6 and 7, the pattern of minimum curvature directions is more straightforward. In each of the ridge surfaces the minimum curvature direction is parallel to the crest line. Hence, for the narrow ridges, the minimum curvature direction points in the  $y$ -direction, while for the wide ridge, it points in the  $x$ -direction. Again, the Darboux smoothing method gives better agreement with

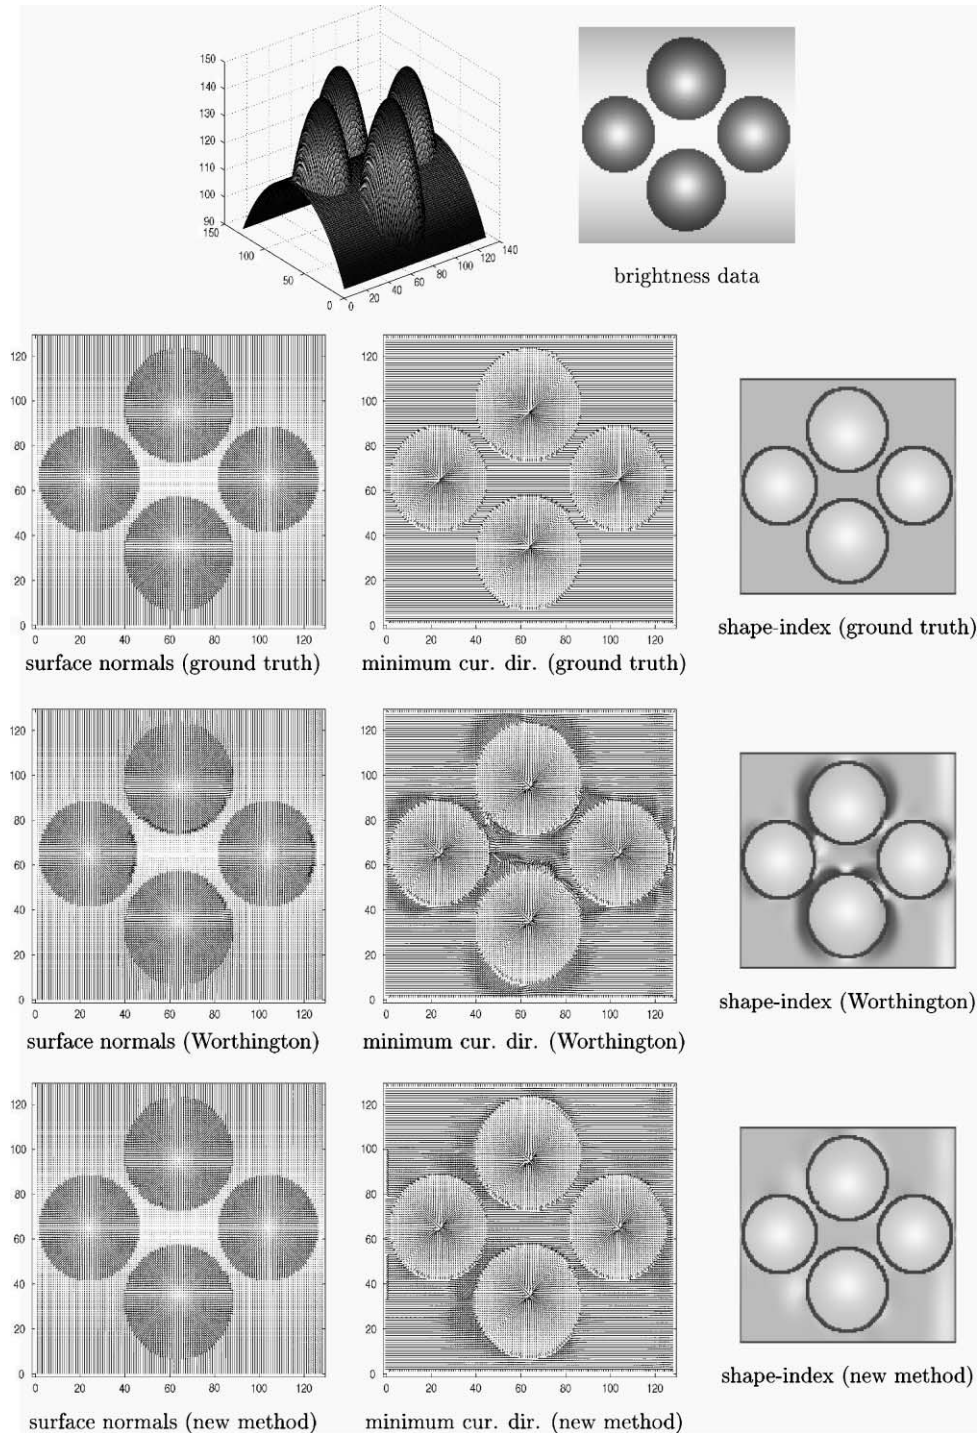


Fig. 5. Results for paraboloid-ellipsoids; ground-truth data (first and second rows); the Worthington and Hancock method (third row); the new Darboux smoothing method (bottom row).

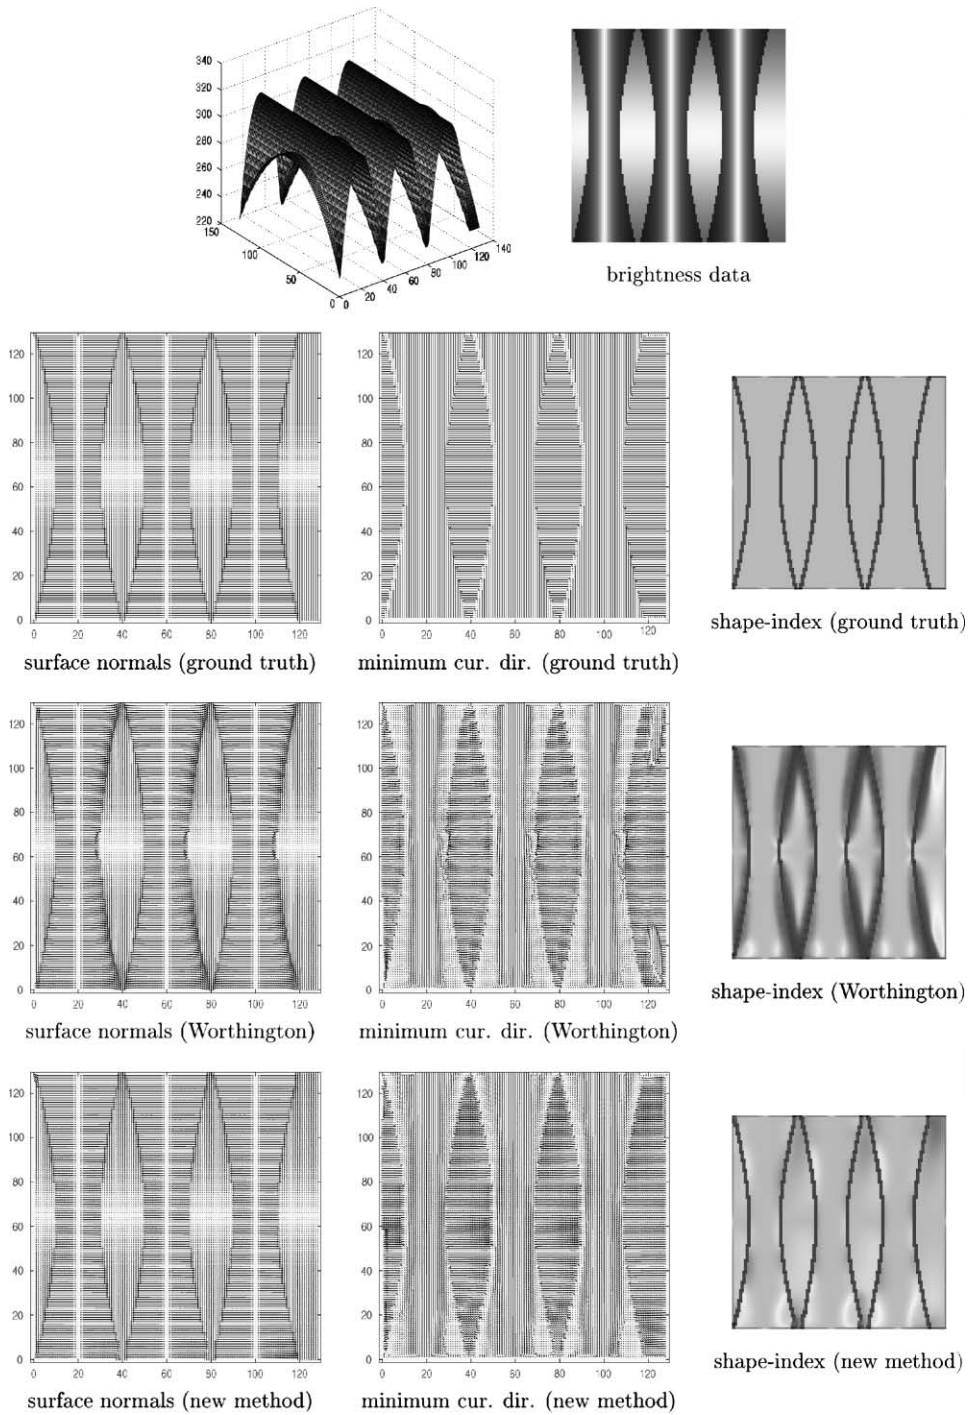


Fig. 6. Results for crossed paraboloids; ground-truth data (first and second rows); the Worthington and Hancock method (third row); the new Darboux smoothing method (bottom row).

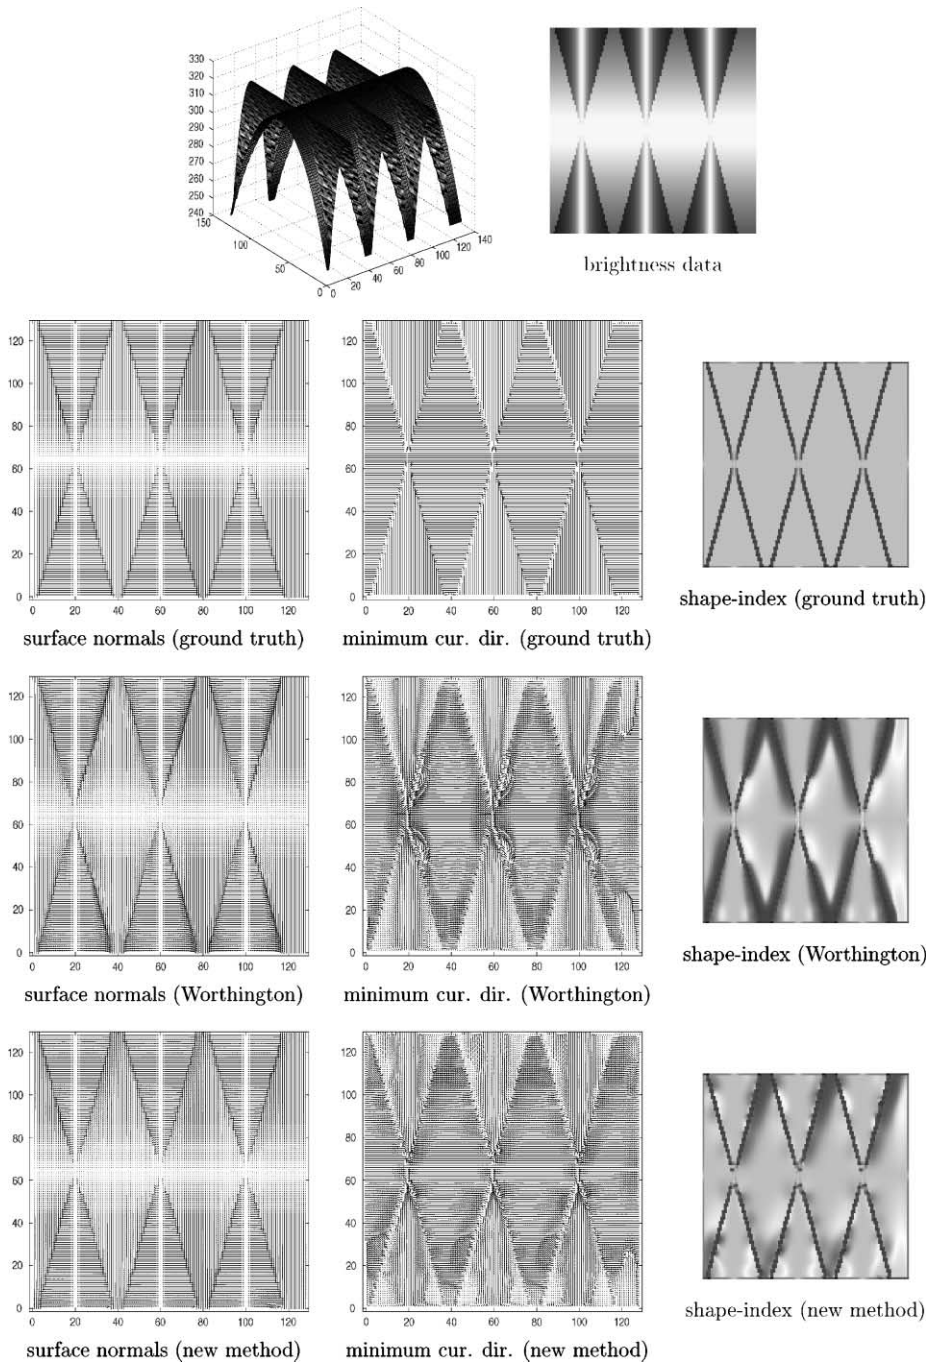


Fig. 7. Results for crossed paraboloids; ground-truth data (first and second rows); the Worthington and Hancock method (third row); the new Darboux smoothing method (bottom row).

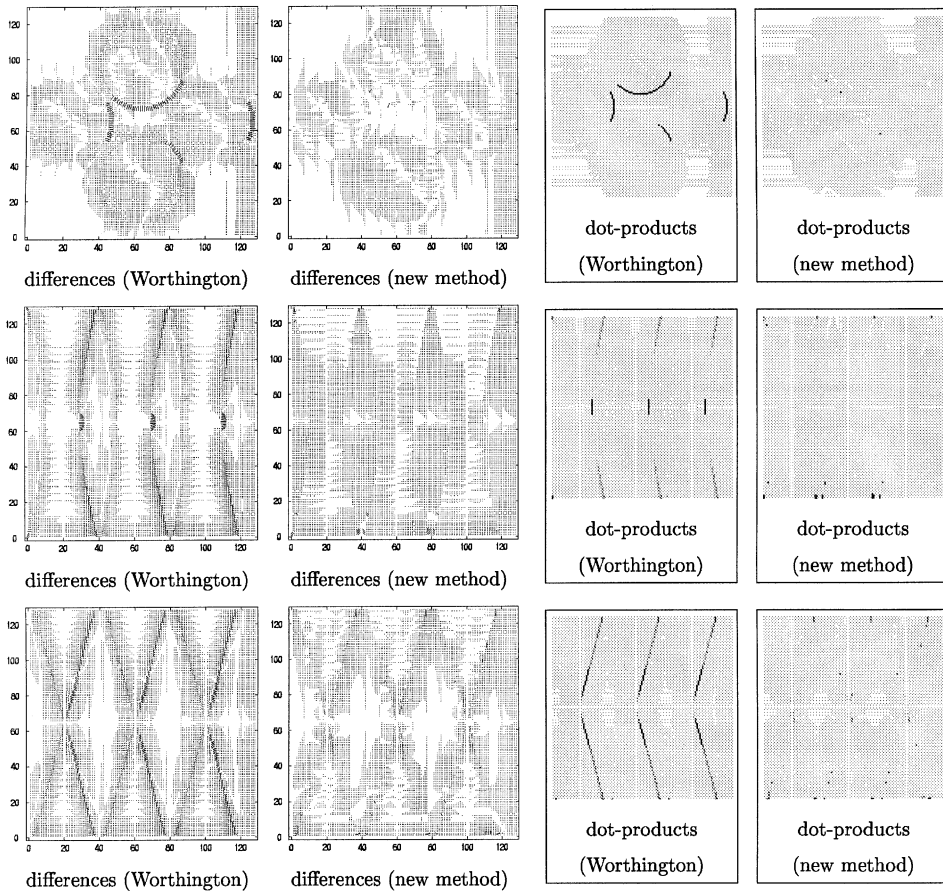


Fig. 8. Results for the synthetic images: rows from top to bottom are for the surfaces used in Figs. 5–7 respectively; field of difference for the results obtained using the Worthington and Hancock method (first column) and using the new Darboux smoothing method (second column); dot-product images for the former (third column) and the latter (fourth column) methods.




Fig. 9. Comparison between the total surface normal errors computed at each iteration for the results obtained using the Worthington and Hancock (dashed curves) and the new Darboux smoothing (solid curves) methods; (panels from left to right are for the Figs. 5–7 respectively).

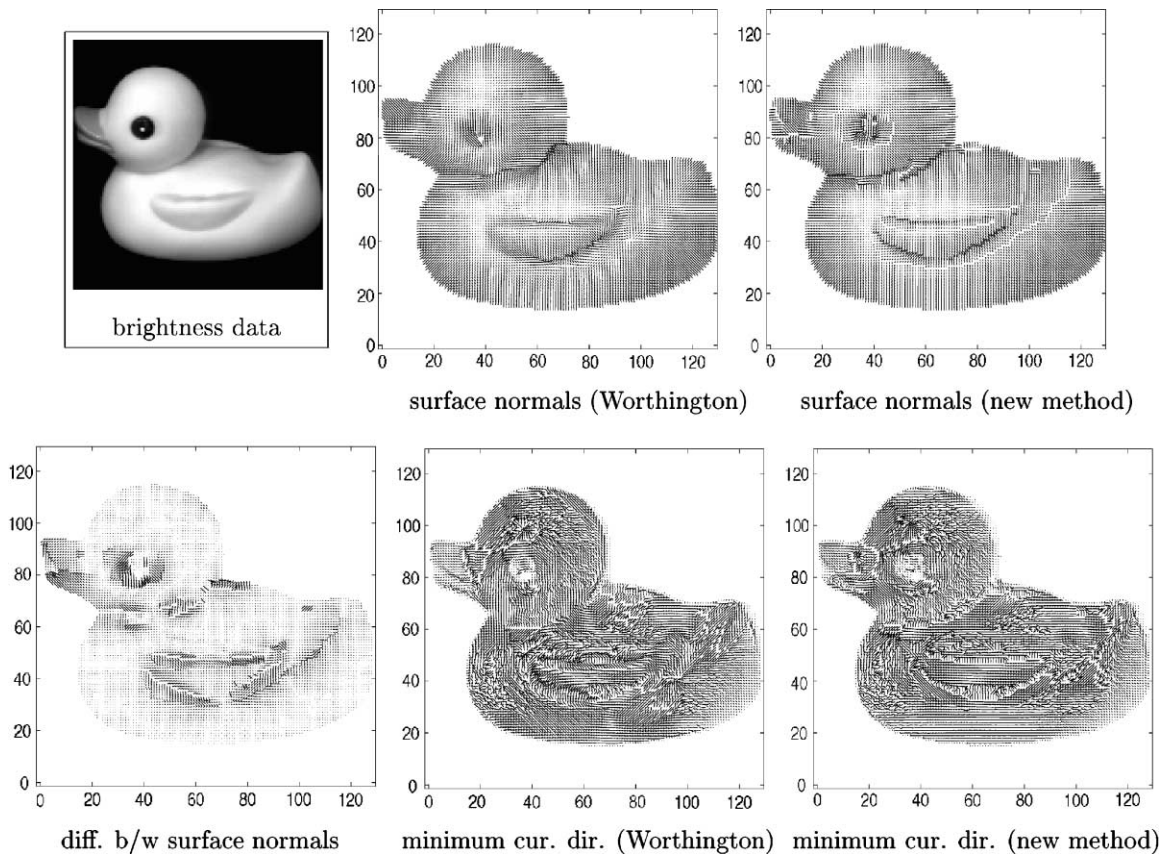


Fig. 10. Results for the toy duck (top-left): surface normals obtained using the two methods (first row); field of difference and minimum curvature directions (second row).

ground-truth and results in less blurring at the intersection of the different surfaces.

In Fig. 8, the rows from top to bottom are for the three synthetic surfaces used in Figs. 5–7 respectively. the first and the second columns show the field of difference vectors between the ground-truth and the surface normals obtained using the Worthington and Hancock method (first column) and our Darboux smoothing method (second column). Furthermore, the third and the fourth columns show the images of dot-products between the estimated and the ground-truth surface normal directions. For the first row, from the difference fields and dot-product panels it is clear that the main departures between the two smoothing methods occur in the dome regions of the surface.

The plots in Fig. 9 present some data to illustrate the iterative qualities of the two algorithms.

The curves show the total error between the ground-truth and the estimated surface normals as a function of iteration number. If  $\vec{N}_{GT}(i, j)$  is the field of unit ground-truth surface normals and  $\vec{N}(i, j)$  is the field of estimated unit surface normals, then the error is defined to be  $E = \sum_{i,j} [||\vec{N}_{GT}(i, j) - \vec{N}(i, j)||]$ . Here the solid curve shows the error obtained with Darboux smoothing, while the dotted curve is the result obtained with the Worthington and Hancock method. The Darboux smoothing method converges monotonically to a lower value of the total error in fewer iterations. In the case of the Worthington and Hancock method, the method initially converges, but then the error rapidly increases with iteration number. This convergence problem with the Worthington and Hancock method is overcome by terminating the method at the

optimal number of iterations. When our method is run with this optimal number of iterations, then the error is comparable, and in some cases better than, the Worthington and Hancock method.

Finally, we discuss the additional computational overheads introduced by our Darboux smoothing method when compared with that of Worthington and Hancock. At each iteration, our method involves the additional computation of the eigenvalues and eigenvectors of the Hessian matrix. This is straightforward since the Hessian matrix is  $2 \times 2$  and, as shown in Section 3, there are simple closed-form expressions for the quantities. In addition, there is the computation of the weighted mean principal curvature direction and the rotation of the Darboux frame. Since the amount of Darboux smoothing depends on the topographic complexity of the surfaces under

study, the additional computation required is image dependant. The main effect here is that the number of iterations required increases with the complexity of the topographic structure present. However, as an example for the image shown in Fig. 5 which is of size  $128 \times 128$  pixels, the Worthington and Hancock method took 7 s, while our new method took 11 s (both in 33 iterations) on a Pentium III (700 MHz).

## 6.2. Real-world data

We have studied a variety of real-world images. The first of these is the toy duck (Fig. 10) which comes from the Columbia University Object Library (COIL-20) database Nene et al. (1996). This database contains a number of images of relatively complex shapes viewed from different directions

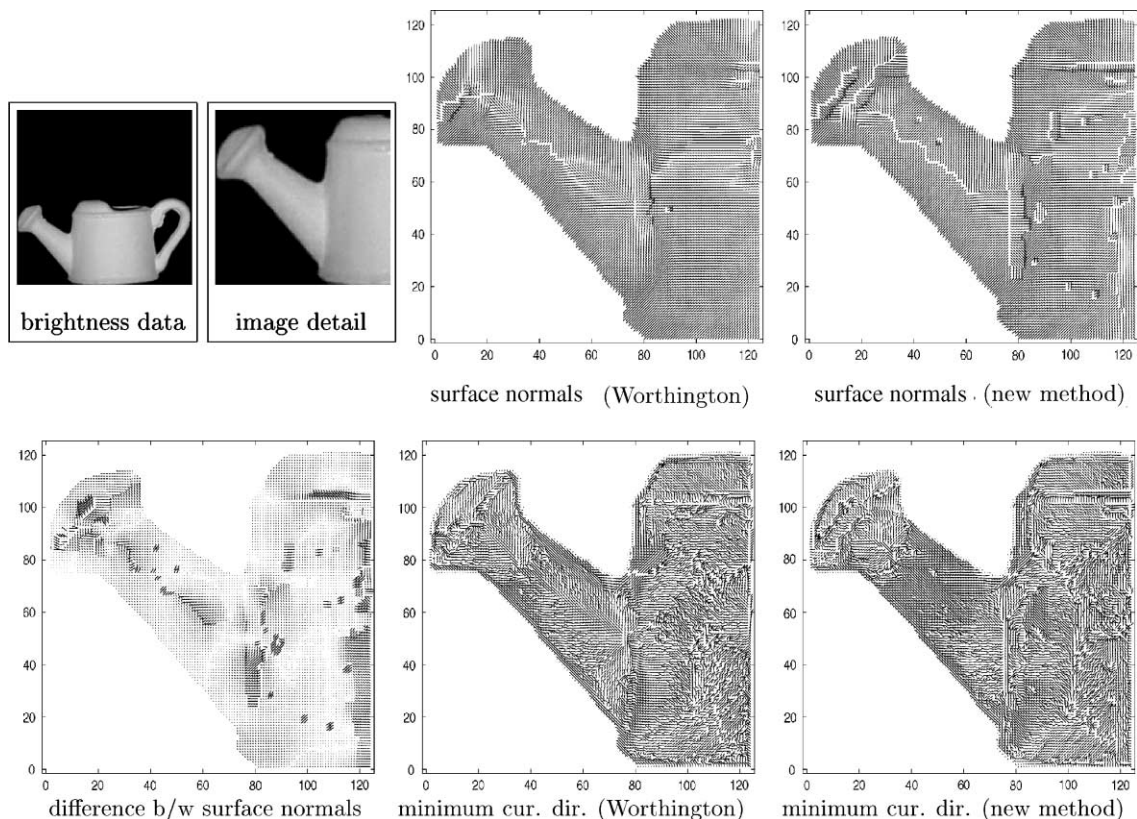


Fig. 11. Results for the terra-cotta watering can (top-left): surface normals obtained using the two methods (first row); field of difference and minimum curvature directions (second row).

with known light-source direction. We have also captured images of terra-cotta objects (Figs. 11, 12) using an Olympus 10E camera. These objects are made of clay and although their surfaces are slightly rough (Ragheb and Hancock, 2002) (non-Lambertian), no filtering has been applied to them. Each object has been imaged under controlled lighting conditions in a darkroom. The objects have been illuminated using a single collimated tungsten light-source. The light-source direction is recorded at the time the images are captured. However, if the light-source direction is not to hand, there are some well known methods in the literature for estimating the light-source direction from a single image (Pentland, 1982).

In Figs. 10–13 we study images of a toy duck, a terra-cotta watering can and a terra-cotta teapot.

In Figs. 10–12, the left panel in the top row is the original image. The results shown in the first row are the estimated field of surface normals obtained using the two methods. In the second row, from left to right the panels show the following results. In the first panel we show the field of difference vectors between the two sets of surface normals. The second and third panels show the fields of minimum curvature directions for the Darboux smoothing method and the Worthington and Hancock method. Finally, in Fig. 13, we use the surface normals obtained using the two methods to compute the shape-index images (the first and the second columns), the curvedness images (the third and the fourth columns) and the dot-product images (the fifth column) for the three objects shown in Figs. 10–12.

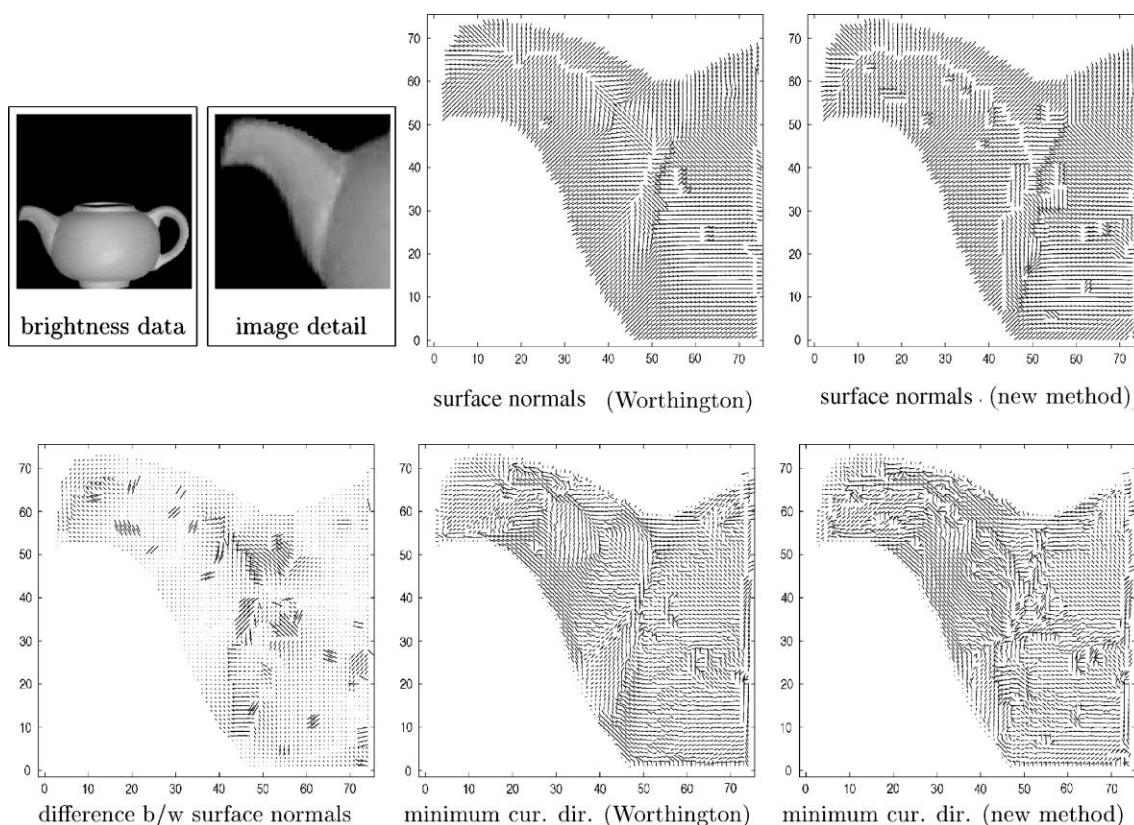


Fig. 12. Results for the terra-cotta teapot (top-left): surface normals obtained using the two methods (first row); field of difference and minimum curvature directions (second row).

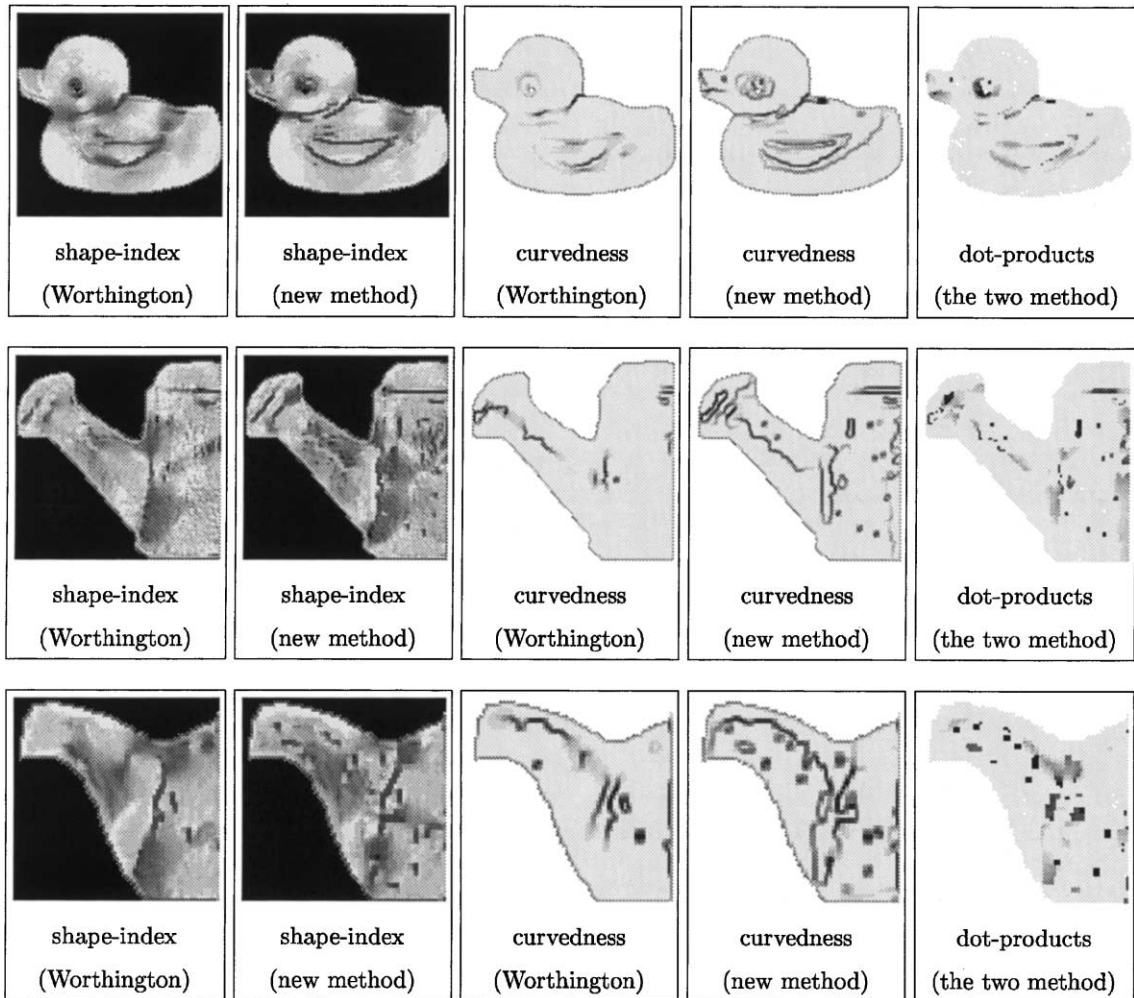


Fig. 13. Results for the toy duck, the watering can and the coffee-pot: shape-index, curvedness and dot-product images for the surface normal directions obtained using the Worthington and Hancock method and the new Darboux smoothing method (shown in Figs. 10–12 respectively).

In the case of the toy duck image shown in Fig. 10, there are a number of features worth mentioning. First, from the curvedness data, the Darboux smoothing method locates the ridge line which extends from the tail of the duck to its chest. This structure is not detected by the Worthington and Hancock method. Second, from the shape-index data, the Darboux smoothing method provides cleaner detail around, the neck, wing and beak. These features are also evident in the surface normal difference field.

In the case of the watering can image of Fig. 11, there are similar differences. These are mainly associated with the intersection of the spout with the body, and the detail of the spout. From the curvedness data, the Darboux smoothing method locates the main crest line of the spout, which is a weak feature in the Worthington and Hancock method. In the teapot image of Fig. 12, we have a similar surface structure around the spout. The Darboux smoothing method again gives better location of the surface intersection and the structure of the ribbed spout.

## 7. Conclusions

To conclude, we have presented a new smoothing method for use in conjunction with shape-from-shading. The method adjusts the surface normal direction in a manner which is consistent with the principal curvature directions. Since smoothing of the principal curvature field is not appropriate at saddles or umbilics, we control the process using a topography sensitive process. The net effect is a smoothing process which allows the field of principal curvature directions to take control in parabolic surface regions. The net result is a shape-from-shading method which has improved effectiveness in detecting crest lines.

There are clearly a number of ways in which the work described in this paper can be extended. In the first instance, the smoothing process can be extended to elliptic and hyperbolic surface regions using curl and divergence criteria to control the regularisation of the principal curvature directions. A second line of investigation will be to incorporate the smoothing process into situations where the reflectance is not completely Lambertian and different geometric constraints need to be applied to the surface normals. For instance for specular reflectance, the surface normal is the bisector of the light-source and viewer directions. Studies aimed at addressing these issues are in hand and will be reported in due course.

## References

- Belhumeur, P.N., 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, San Francisco, USA. pp. 270–277.
- Bichsel, M., Pentland, A.P., 1992. A simple algorithm for shape from shading. Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, Champaign, USA. pp. 459–465.
- Brooks, M.J., Horn, B.K.P., 1985. Shape and source from shading. Proceedings of the International Joint Conference on Artificial Intelligence. pp. 932–936.
- Bruckstein, A.M., 1988. On shape from shading. Computer Vision Graphics Image Process. 44, 139–154.
- Nene, S.A., Nayar, S.K., Murase, H., 1996. Columbia Object Image Library (COIL-20), Technical Report CUCS-005-96.
- Ferrie, F.P., Lagarde, J., 1990. Curvature consistency improves local shading analysis. International Conference on Pattern Recognition. Vol. 1, Atlantic City, USA. pp. 70–76.
- Frankot, R.T., Chellappa, R., 1988. A method for enforcing integrability in shape from shading algorithms. IEEE T. Pattern Anal. 10 (4), 439–451.
- Horn, B.K.P., Brooks, M.J., 1986. The variational approach to shape from shading. Computer Vision Graphics Image Process. 33 (2), 174–208.
- Kimmel, R., Bruckstein, A.M., 1995. Tracking level-sets by level-sets: a method for solving the shape from shading problem. Comput. Vis. Image Und. 62 (1), 47–58.
- Koenderink, J.J., van Doorn, A.J., 1992. Surface shape and curvature scales. Image Vis. Comput. 10, 557–565.
- Koenderink, J.J., van Doorn, A.J., Kappers, A.M.L., 1992. Surface perception in pictures. Percept. Psychophys. 52 (5), 487–496.
- Koenderink, J.J., van Doorn, A.J., 1979. The internal representation of solid shape with respect to vision. Biol. Cybern. 32, 211–216.
- Oliensis, J., Dupuis, P., 1993. A global algorithm for shape from shading. Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. New York, USA. pp. 692–701.
- Pentland, A.P., 1982. Finding the illuminant direction. J. Opt. Soc. Am. 72 (4), 448–455.
- Ragheb, H., Hancock, E.R., 2002. Highlight Removal using Shape-from-Shading. The Seventh European Conference on Computer Vision. Vol. 2. Copenhagen, Denmark. pp. 626–641.
- Sander, P.T., Zucker, S.W., 1990. Inferring surface structure and differential structure from 3D images. IEEE T. Pattern Anal. 12 (9), 833–854.
- Woodham, R.J., 1994. Gradient and curvature from the photometric-stereo method, including local confidence estimation. J. Opt. Soc. Am. 11 (11), 3050–3068.
- Worthington, P.L., Hancock, E.R., 1999. New constraints on data-closeness and needle-map consistency for SFS. IEEE T. Pattern Anal. 21 (11), 1250–1267.
- Zhang, R., Tsai, P., Cryer, J.E., Shah, M., 1999. Shape from shading: a survey. IEEE T. Pattern Anal. 21 (8), 690–706.