BMVC 2002

A Mumford-Shah Diffusion Process for Shape-from-Shading

Antonio Robles-Kelly* Edwin R. Hancock
Department of Computer Science,
University of York,York YO1 5DD, UK.
{arobkell,erh}@minster.cs.york.ac.uk

Abstract

In this paper we show how the Mumford-Shah functional can be used to derive diffusion kernels that can be employed in the recovery of surface height in shape-from-shading. We commence from an initial field of surface normals which are constrained to fall on the Lambertian reflectance cone and to point in the direction of the local Canny edge gradient. We aim to find a path through this field of surface normals which can be used for surface height reconstruction. We demonstrate that the Mumford-Shah functional leads to a diffusion process which is a Markov chain on the field of surface normals. Moreover, the diffusion kernel for the Markov process depends on the rate of change of surface normal direction. As a result, we can find the steady state of the Markov chain using the leading eigenvector for the transition probability matrix computed from the diffusion kernels. We show how the steady state path can be used for height recovery and also for smoothing the initial field of surface normals.

1 Introduction

Shape-from-shading is a process which allows surface shape to be recovered from shading variations using the physics of light. Unfortunately the problem is an ill-posed one [1] since at each image location it requires the two degrees of freedom of the surface normal, i.e. the local slant and tilt angles of a surface, to be recovered from a single measured brightness value. To overcome this problem the process is frequently posed in a variational setting, in which the goal is to find the field of surface normals so as to minimise a regularised energy functional [2].

The main problem with this approach is that the smoothness constraints tend to dominate the data-closeness constraints, resulting in a recovered field of surface normals that is devoid of fine surface detail. To overcome this problem, Worthington and Hancock [3] have recently developed a new geometric framework for shape-from-shading. Data-closeness is guaranteed by constraining the surface normals to fall on the local irradiance cone defined by Lamberts law. That is to say they must reside on a cone whose axis is the light source direction and whose opening angle is the inverse cosine of the measured image intensity.

Although effective, this new framework leaves considerable scope for further development. In particular, it is concerned with the recovery of fields of surface normals and does not address the problem of recovering surface height. In this paper our aim is to develop a height recovery method in which the surface normals are constrained to fall on

*Supported by CONACYT, under grant No. 146475/151752.

BMVC 2002 doi:10.5244/C.16.69

the Lambertian reflectance cone. We pose the height recovery process as that of finding an integration path through a field of surface normals which satisfy Lambert's equation. Our starting point is the Mumford-Shah functional [4]. This introduces penalties associated with curve length and total squared curvature of the path. The required curvature may be computed from the change in surface normal direction between locations on the image plane. We use the Mumford-Shah energy to define a probability distribution for the integration path. We demonstrate that this is a Markov chain on the sites of the field of surface normals. The diffusion kernels for the Markov chain hence depend on the sectional curvature associated with the trace of the integration path on the surface. We make use of the Markov property and develop a graph-spectral method for recovering an integration path for surface height recovery. To do this we borrow ideas from spectral graph theory concerning random walks on graphs [5, 6, 7]. The steady state random walk on the graph can be located by finding the leading eigenvector of the associated transition probability matrix. Moreover, the eigenstructure of the transition matrix also allows us to locate surface patches.

We use this property to find a curvature minimizing path across the field of available surface normals associated with each surface patch. We adjust the initial field of surface normals by traversing these paths and adjusting the surface normal direction. Each surface normal is rotated on the cone so that it points in the direction of minimum curvature along the path. This approach offers a number of advantages. First, we avoid smoothing across patch boundaries. Second, smoothing only proceeds in the direction of minimum curvature, and is hence regulated. Third, the propagation process is less computationally demanding since it is uni-dimensional.

2 Lambertian Reflectance

In the case of Lambertian reflectance from a matte surface of constant albedo illuminated with a single point light-source, the observed intensity is independent of the viewing direction. The observed intensity depends only on the quantity of absorbed light, and this in turn is proportional to the cosine of the incidence angle. Suppose that $\vec{L}$ is the unit-vector in the direction of the light source and that $\vec{N}_i$ is the unit-vector in the surface normal direction at the pixel i. According to Lambert's law, the observed image intensity at the pixel indexed i is $E_i = \vec{N}_i \cdot \vec{L}$.

Lambert's equation provides insufficient information to uniquely determine the surface normal direction. However, as recently observed by Worthington and Hancock [3], the equation does have a simple geometric interpretation which can be used to constrain the direction of the surface normal. The equation specifies that the surface normal must fall on the surface of a right-cone whose axis is aligned in the light-source direction $\vec{L}$ and whose apex angle is $\arccos(E)$.

Worthington and Hancock [3] exploit this property to develop a two-step iterative process for SFS. The process commences from a configuration in which the surface normals are placed on the position on the irradiance cone where their projections onto the image plane are aligned in the direction of the local (Canny) image gradient.

3 Diffusion Kernels

Stated formally, our goal is the recovery of height information from the field of surface normals. ¿From a computational standpoint the aim is to find a path along which simple trigonometry may be applied to increment the estimated height function. To be more

formal suppose that the surface under study is S and that the field of surface normals is sampled on the plane Π. Our aim here is to find a curve Γ across the plane Π that can be used as an integration path to reconstruct the height-function of the surface S. The projection of the curve Γ onto the surface S is denoted by Γ_S. Further, suppose that $\kappa(s)$ is the sectional curvature of the curve Γ_S at the point Q with parametric co-ordinate s. We seek the path Γ_S that minimises the Mumford-Shah [4] functional

$$\mathcal{E}(\Gamma_S) = \int_{\Gamma_S} \Big\{ \alpha + \beta\kappa(s)^2 \Big\} ds \tag{1}$$

where α and β are constants. The probability of the path can be written as $P_{\Gamma_S} = \exp[-\mathcal{E}(\Gamma_S)]$.

The field of unit surface normals for the surface S on the plane Π is denoted by $\vec{N}$. Accordingly, and following do Carmo [8], we let $T_Q(S)$ represent the tangent plane to the surface S at the point Q which belongs to the curve Γ_S. To compute the sectional curvature $\kappa(s)$ we require the differential of the surface or Hessian matrix $d\vec{N}_Q : T_Q(S) \rightarrow T_Q(S)$. The maximum and minimum eigenvalues λ_1 and λ_2 of $d\vec{N}_Q$ are the principal curvatures at the point Q. The corresponding eigenvectors $\vec{e}_1 \in T_Q(S)$ and $\vec{e}_2 \in T_Q(S)$ form an orthogonal basis on the tangent plane $T_Q(S)$. At the point Q the unit normal vector to the curve Γ is $\vec{n}_\Gamma$ and the unit tangent vector is $t_Q \in T_Q(S)$. The sectional curvature of Γ at Q is given by

$$\kappa(s) = \frac{(\vec{t}_Q.\vec{e}_1)^2(\lambda_1 - \lambda_2) + \lambda_2}{\vec{n}_\Gamma.\vec{N}_Q} \tag{2}$$

where $(\vec{t}_Q.\vec{e}_1)^2(\lambda_1 - \lambda_2) + \lambda_2$ is the normal curvature and $\psi = \arccos \vec{n}_\Gamma.\vec{N}_Q$ is the angle between the curve normal and the surface normal.

In practice, we will be dealing with points which are positioned at discrete positions on the pixel lattice. Suppose that i and j are the pixel indices of neighbouring points sampled on the pixel lattice along the path Γ_S. With this discrete notation, the path energy and its probability are given by

$$\mathcal{E}(\Gamma_S) = \sum_{(i,j)\in\Gamma_S} \Big\{ \alpha + \beta\kappa^2_{i,j} \Big\} s_{i,j} \quad \text{and} \quad P_{\Gamma_S} = \prod_{(i,j)\in\Gamma_S} \exp\Big[-\Big\{ \alpha + \beta\kappa^2_{i,j} \Big\} s_{i,j} \Big] \tag{3}$$

where $\kappa_{i,j}$ is an estimate of the curvature based on the surface normal directions at the pixel locations i and j, and $s_{i,j}$ is the path distance between these points. Hence, we can view the integration path Γ as Markov chain across the field of surface normals. The diffusion kernel for the process, i.e. probability of migrating between sites i and j is

$$K = \exp[-\mathcal{E}_{i,j}] = \exp\Big[-\Big\{ \alpha + \beta\kappa^2_{i,j} \Big\} s_{i,j} \Big] \tag{4}$$

In order to compute the path curvature appearing in the diffusion kernel, we make use of the surface normal directions. To commence, we note that $|\kappa_{i,j}| = \frac{1}{R_{i,j}}$ where $R_{i,j}$ is the radius of the local circular approximation to the integration curve on the surface. Suppose that the surface normal directions at the pixel locations i and j are respectively $\vec{N}_i$ and $\vec{N}_j$. The approximating circle connects the points i and j, and has the path segment

$s_{i,j}$ as the connecting chord. The change in direction of the radius vector of the circle is $\theta_{i,j} = \arccos \vec{N}_i.\vec{N}_j$ and hence $\cos\theta_{i,j} = \vec{N}_i.\vec{N}_j$. If the angle $\theta_{i,j}$ is small, then we can make the Maclaurin approximation $\cos\theta_{i,j} \simeq 1 - \frac{\theta_{i,j}^2}{2} = \vec{N}_i.\vec{N}_j$. Moreover, the small angle approximation to the radius of curvature of the circle is $R_{i,j} = \frac{s_{i,j}}{\theta_{i,j}}$. Hence, $\kappa_{i,j}^2 = \frac{2}{s_{i,j}^2(1-\vec{N}_i.\vec{N}_j)}$. As a result, we find that the cost associated with the step from the pixel i to the pixel j is $\mathcal{E}_{i,j} = \alpha s_{i,j} + \frac{2\beta}{s_{i,j}}(1 - \vec{N}_i.\vec{N}_j)$. The total cost associated with the integration path Γ_S is hence

$$\mathcal{E}_{\Gamma_S} = \alpha L_{\Gamma_S} + \sum_{(i,j)\in\Gamma_S} \frac{2\beta}{s_{i,j}}(1 - \vec{N}_i.\vec{N}_j) \tag{5}$$

where L_{Γ_S} is the length of the path. Hence, the integration path is a form of elastica which attempts to find an energy minimising path through the surface normals. The energy function is a variant of the Mumford-Shah functional. It has two terms. The first encourages the integration path to be one of minimum length. The second term encourages a path which minimises the total change in surface normal direction. There are clearly a number of ways in which the energy can be minimised. However, here we choose to make use of the fact that from Equation 6 it is clear that the energy function specifies a Markov chain on the sites of the Gauss map. In other words, the path can be viewed as a diffusion on the Gauss map. The steady state random walk for this diffusion can be found by locating the leading eigenvector of the transition probability matrix for the Markov chain.

To persue the graph-spectral analysis of the Gauss map, we require a transition probability matrix. For the pixels indexed i and j we commence by computing the weight

$$w_{i,j} = \begin{cases} \exp[-\mathcal{E}_{i,j}] & \text{if } j \in \mathcal{N}_i \\ 0 & \text{otherwise} \end{cases} \tag{6}$$

where $\mathcal{N}_i$ is the set of pixels-neighbours of the pixel i. Hence, the curvature weight is only non-zero if pixels abut one-another.

From the curvature dependant weights, we construct a transition probability matrix in which the upper diagonal elements sum to unity. The element in the row indexed i and column indexed j is

$$P(j,i) = P(i,j) = \frac{w_{i,j}}{\sum_{l=1;l\in V}^{|V\times V|}\sum_{k=l;k\in V}^{|V\times V|} w_{k,l}} \tag{7}$$

In next section, we describe how the leading eigenvector of this matrix can be used to determine steady state random walk for the Markov chain and how this path may be used for the purposes of surface integration.

4 Random Walks and Markov Random Chains

The set of pixel sites can be viewed as a weighted graph $G = (V, E, P)$ with index-set V and edge-set $E = \{(i,j)|(i,j) \in V \times V, i \neq j\}$. The off-diagonal elements of the transition probability matrix P are the weights associated with the edges. In this paper, we exploit a graph-spectral property of the transition matrix P to develop a surface height recovery method. This requires that we have the eigenvalues and eigenvectors of the matrix P to hand. To find the eigenvectors of the transition probability matrix, P, we first

solve polynomial equation $|P - \lambda I| = 0$. The unit eigenvector $\vec{\phi}_i$ associated with the eigenvalue λ_i is found by solving the system of linear equations $P\vec{\phi}_i = \lambda_i\vec{\phi}_i$ and satisfies the condition $\vec{\phi}_i^T\vec{\phi} = 1$.

Consider a random walk on the graph G. The walk commences at the node j_1 and proceeds via the sequence of edge-connected nodes $\Gamma = \{j_1, j_2, j_3, ...\}$ where $(j_i, j_{i-1}) \in E$. Suppose that the transition probability associated with the move between the nodes j_l and j_m is $P(l, m)$. If the random walk can be represented by a Markov chain, then the probability of visiting the nodes in the sequence above is $P_{\Gamma_S} = P(j_1)\prod_{l=1}^{|V|} P(j_{l+1}, j_l)$. This Markov chain can be represented using the transition probability matrix P whose element with row l and column m is $P(l, m)$. Further, let $Q_t(i)$ be the probability of visiting the node indexed i after t-steps of the random walk and let $Q_t = (Q_t(1), Q_t(2), ...)^T$ be the vector of probabilities. After t time steps $Q_t = (P^T)^t Q_0$. If λ_i are the eigenvalues of P and $\vec{\phi}_i$ are the corresponding eigenvectors of unit length, then $P = \sum_{i=1}^{|V|} \lambda_i\vec{\phi}_i\vec{\phi}_i^T$. As a result, after t applications of the Markov transition probability matrix $P^t = \sum_{i=1}^{|V|} \lambda_i^t\vec{\phi}_i\vec{\phi}_i^T$. If the row and columns of the matrix P sum to unity, then $\lambda_1 = 1$. Furthermore, from spectral graph theory [6] provided that the graph G is not a bipartite graph, then the smallest eigenvalue $\lambda_{|V|} > -1$. As a result, when the Markov chain approaches its steady state, i.e. $t \to \infty$, then all but the first term in the above series become negligible. Hence, $\lim_{t\to\infty} P^t = \vec{\phi}_1\vec{\phi}_1^T$. This establishes that the leading eigenvector of the transition probability matrix is the steady state of the Markov chain. For a more complete proof of this result see the book by Varga [9] or the review of Lovasz [5]. As a result, if we visit the nodes of the graph in the order defined by the magnitudes of the co-efficients of the leading eigenvector of the transition probability matrix, then the path is the steady state Markov chain. In this paper we aim to perform surface height recovery by co-joining pixel sites pixel sites along the path specified by the magnitude order of the components of the leading eigenvector.

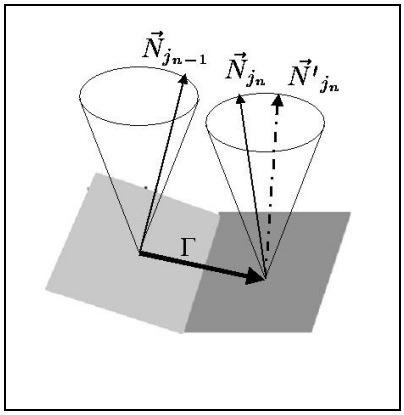

Figure 1: Geometrical meaning of the normal correcting step.

Suppose that the leading eigenvector for the transition probability matrix is denoted by $\vec{\phi}^* = (\phi^*(1), ..., \phi^*(|V|))^T$. Our aim is to use the sequence of nodes defined by the rank order of the magnitudes of the components of the leading eigenvector to define an integration path through the set of pixel sites. The pixel path defined by the rank order of the co-efficients of the leading eigenvector is given by the list of sorted node-indices $\Gamma = (j_1, j_2, j_3, ..., j_{|V|})$ where $\phi^*(j_1) > \phi^*(j_2) > \phi^*(j_3) > ... > \phi^*(j_{|V|})$. The subscript n of the pixel with node-index $j_n \in V$ is hence the rank-order of the eigenvector component $\phi^*(j_n)$.

5 Extracting Patches

In practice the surface under study may have a patch structure. The patches may be identified by finding the blocks of the transition probability matrix induced under a permutation

of the nodes. We commence by constructing the thresholded transition probability matrix A whose elements are defined as follows

$$A(i,j) = \begin{cases} 0 & \text{if } P(i,j) << 1 \\ P(i,j) & \text{otherwise} \end{cases} \tag{8}$$

Suppose that there are m distinct surface patches, each associated with an adjacency matrix $B^{(i)}$. If C represents a noise matrix, then the relationship between the observed transition matrix A and the underlying block-structured transition matrix is $A = B + C$ where $B = \mathcal{P}B_D\mathcal{P}^T$, $\mathcal{P}$ is a permutation matrix and $B_D = diag(B^{(1)}, B^{(2)}, .., B^{(i)}...)$ is a block diagonal matrix in which $B^{(i)}$ is the subblock corresponding to the patch indexed i. To recover the matrix B_D, we turn to the eigenvector expansion of the matrix A and write

$$A = \vec{b}_*\vec{b}_*^T + \sum_{i=2}^{|V|} \lambda_i \vec{b}_i\vec{b}_i^T \tag{9}$$

where the leading eigenvalue is unity i.e. $\lambda_1 = 1$, $\vec{b}_*$ is the leading eigenvector and the eigenvectors are normalised to be of unit length, i.e. $|\vec{b}_i| = 1$. To identify the patches, we use the following iterative procedure. We initialise the algorithm by letting $A^{(1)} = A$. Further suppose that $\vec{b}_*^{(1)}$ is the leading eigenvector of $A^{(1)}$. The matrix $B^{(1)} = \vec{b}_*^{(1)}\vec{b}_*^{(1)T}$ represents the first block of A. The nodes with non-zero entries belong to the patch. These nodes may be identified and removed from further consideration. To do this we compute the residual transition matrix $A^{(2)} = A^{(1)} - B^{(1)}$ in which the elements of the first patch are nulled. The leading eigenvector $\vec{b}_*^{(2)}$ of the residual transition matrix $A^{(2)}$ is used to compute the second block $B^{(2)} = \vec{b}_*^{(2)}\vec{b}_*^{(2)T}$. The process is repeated iteratively to identify all of the principal blocks of A. At iteration n, $\vec{b}_*^{(n)}$ is the leading eigenvector of the residual transition probability matrix $A^{(n)}$, and the n^{th} block is $B^{(n)} = \vec{b}_*^{(n)}\vec{b}_*^{(n)T}$. The index set of the patch indexed n is the set of nodes for which the components of the leading eigenvector $\vec{b}_*^{(n)}$ are non-zero. Hence, the index-set for the i^{th} patch is $S_i = \{k|\vec{b}_*^{(i)}(k) \neq 0\}$. It is important to stress that the patches are non-overlapping, i.e. the inner product of the block eigenvectors for different patches is zero $\vec{b}_*^{(k)}.\vec{b}_*^{(l)} = 0$, where $k \neq l$.

6 Normal Correction and Height Recovery

Our surface height recovery algorithm proceeds along the sequence of pixel sites defined by the order of the co-efficients of the leading eigenvector associated with the separate patches. For the k^{th} patch, the path is $\Gamma_k = (j_k^1, j_k^2, j_k^3, ...,)$ where the order of the co-efficients of the leading eigenvector for this patch is such that $\phi_k^*(j_k^1) > \phi_k^*(j_k^2) > \phi_k^*(j_k^3) > ....$ As we move from pixel-site to pixel-site defined by this path we perform two computations. First, we adjust the position of the surface normals on the irradiance cone so that they are smoothed in the local direction of minimum curvature. Second, we recover the surface from the Gauss map by incrementing the height function using the local slope parameters. At step n of the algorithm, we make a transition from the pixel with path-index j_{n-1} to the pixel with path-index j_n. To perform smoothing and height recovery, we make use of the surface normals $\vec{N}_{j_n}$ and $\vec{N}_{j_{n-1}}$ at the two pixel-sites.

Turning our attention first to the smoothing process, our aim is to adjust the surface normal directions so that they are consistent with the direction of the curvature minimis-

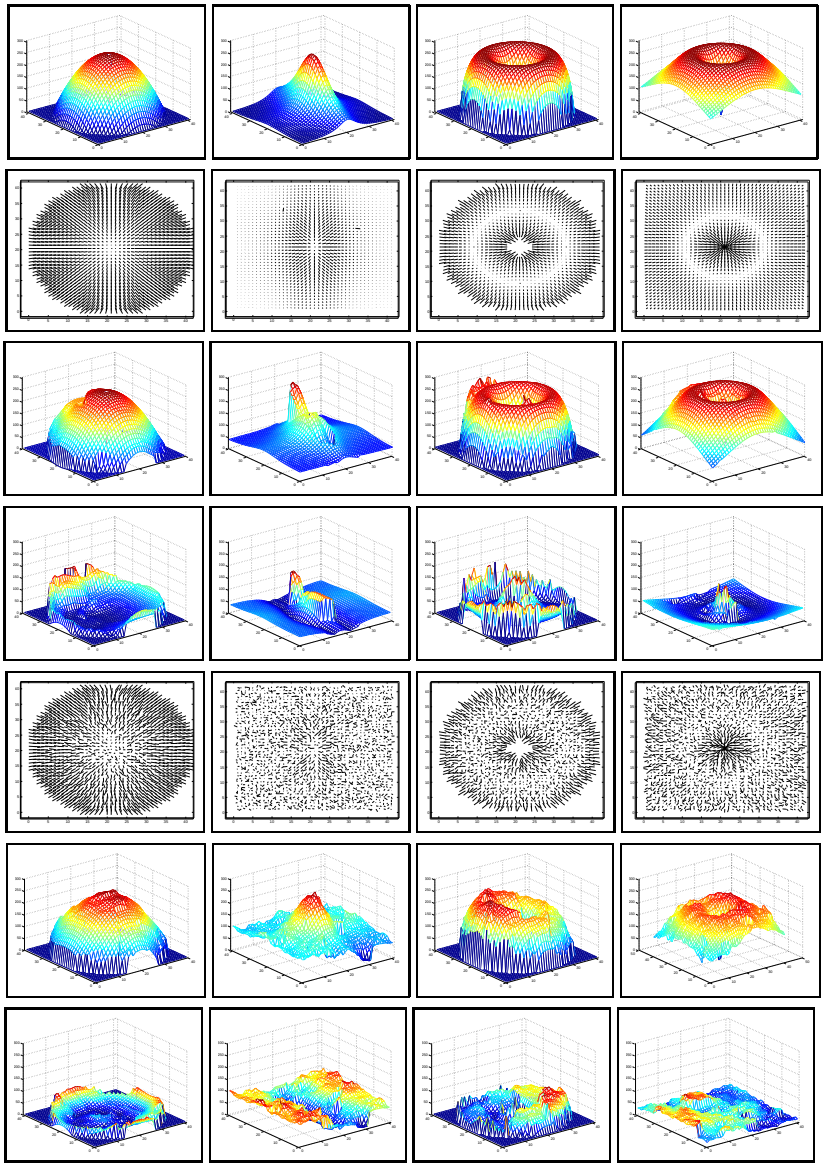

Figure 2: Ground-truth, needlemaps, error plots and surface-height recovery results for four basic shapes under noise-controlled conditions.

ing integration path and also remain on their respective irradiance cone. We therefore parameterise the surface normal directions using two angles. The first of these is the polar angle θ_{j_n} between the surface normal and the light source direction. This is simply the opening angle of the irradiance cone, and this angle must be kept fixed in order to ensure that the recovered surface normal satisfies Lambert's law. The second is the azimuthal angle φ_{j_n}, which measures the position of the surface normal on the irradiance cone. The angles φ_{j_n} and θ_{j_n} can be defined in terms of the surface normal components. Using

some simple trigonometry we can write

$$\varphi_{j_n} = \begin{cases} 2\pi - \arccos\frac{\vec{N}_{j_n}(x)}{R_{j_n}} & \text{if } \vec{N}_{j_n}(y) \\ \arccos\frac{\vec{N}_{j_n}(x)}{R_{j_n}} & \text{otherwise} \end{cases}; \qquad \theta_{j_n} = \arccos(\vec{N}_{j_n}(z))$$

where R_{j_n} is the projection of the surface normal $\vec{N}_{j_n}$ onto the image plane given by $R_{j_n} = \sin(\arccos \vec{N}_{j_n}(z))$.

These angles are adjusted in the smoothing step. This smoothing step uses the following equations to update the tilt and azimuthal angles of the surface normals

$$\varphi'_{j_n} = \varphi_{j_n}(1 - \eta_{j_n}) + \varphi_{j_{n-1}}\eta_{j_n}; \qquad \theta'_{j_n} = \theta_{j_n}$$

where η_{j_n} is a weight which measures the consistency of the surface normal directions with the direction of the curvature minimising path and $\varphi'_{j_n}, \theta'_{j_n}$ are the updated angles. For the smoothing step, we make use of the thresholded change in surface normal direction

$$\Delta N_{j_n} = \begin{cases} \mid \vec{N}_{j_n} - \vec{N}_{j_{n-1}} \mid & \text{if } \mid \vec{N}_{j_n} - \vec{N}_{j_{n-1}} \mid \geq \epsilon \\ 0 & \text{otherwise} \end{cases} \tag{10}$$

where the threshold ϵ is a constant that has been found empirically to have 0.5 as its optimum value. We compute the smoothing weight η_{j_n} using the hyperbolic tangent function as suggested by robust statistics and as used by Worthington and Hancock in their work on needle-map smoothing. The weight is given by

$$\eta_{j_n} = (\log\cosh(\mid \varphi_{j_n} - \varphi_{j_{n-1}} \mid \Delta N_{j_n}))^2 \tag{11}$$

and takes the product of the change in azimuthal angle and surface normal direction as its arguement. Once the updated angles have been computed, then the surface normals may be updated using the equation $\vec{N}'_{j_n} = (R_{j_n}\cos\varphi'_{j_n}, R_{j_n}\sin\varphi'_{j_n}, \cos\theta'_{j_n})^T$. With the smoothed surface normals to hand, we can compute the height increment. Using the trapezium rule the increment may be computed using the components of the surface normals and is given by

$$h_n = \frac{d_n}{2}\left\{\frac{\vec{N}'_{j_n}(x)}{\vec{N}'_{j_n}(y)} + \frac{\vec{N}_{j_{n-1}}(x)}{\vec{N}_{j_{n-1}}(y)}\right\} \tag{12}$$

where d_n is the Euclidean distance on the x-y plane between the two pixel sites associated with the nth transition. If the height-function is initialised by setting $z_{j_0} = 0$, then the centre-height for the pixel with path-index j_n is $z_{j_{n+1}} = z_{j_n} + h_n$. We merge abutting patches by assuring they have the same mean boundary height.

7 Experiments

The experimental evaluation of the new surface reconstruction method is divided into two parts. We commence with a sensitivity study aimed at evaluating the method on synthetic data. In the second part of the study, we focus on real-world data.

7.1 Synthetic Data

In this section we provide some experiments on synthetic data. The aim here is to determine the accuracy of the surface reconstruction method. To this end we have generated synthetic surfaces. From the surfaces, we have computed the field of surface normal directions. We have then applied the graph-spectral method to the 2D field of surface normals to recover an estimate of the surface height.

In Figure 2 we show the results obtained for a series of different surfaces under conditions of controlled noise. To do this we have added random measurement errors to the surface height. The measurement errors have been sampled from a Gaussian distribution with zero mean and known variance. In the top row we show the original noise-free synthetic surface, i.e. ground-truth. In the second row of the figure we show the field of surface normals for the noise-free surface. The third row shows the surface reconstructed from the field of surface normals when no Gaussian noise has been added. The fourth row shows the absolute error between the ground-truth and reconstructed surface height. From left-to-right the surfaces studied are a dome, a sharp ridge, a torus and a volcano. In all four cases the surface reconstructions are qualitatively good. For the dome the height errors are greater at the edges of the surface where the slope is largest. In the case of the ridge, there are errors at the crest. For the volcano, there are some problems with the recovery of the correct depth of the "caldera", i.e. the depression in the centre. For the reconstructed surfaces, the relative mean-squared errors are 5.6% for the dome, 10.8% for the ridge, 7.8% for the torus and 4.7% for the volcano. Hence, the method seems to have greater difficulty for surfaces containing sharp creases.

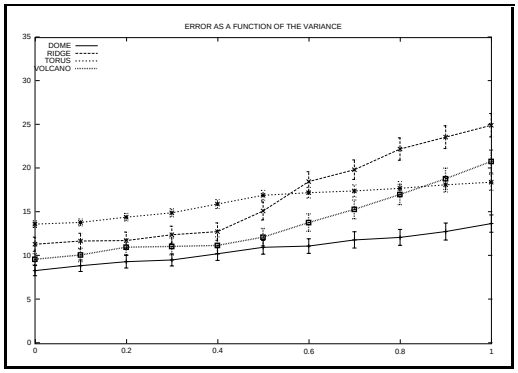

Figure 3: Error versus variance

In the fifth row of Figure 2 we show the field of surface normals for the noise-corrupted surface. In the sixth row, we show the result of reconstructing the surfaces shown in the top row, when random height errors have been added. The seventh , i.e. bottom, row shows the difference between the height of the surface reconstructed from the noisy surface normals and the ground-truth height function. In the case of all four surfaces, the gross structure is maintained. However, the recovered height is clearly noisy. The height difference plots are relatively unstructured. These are important

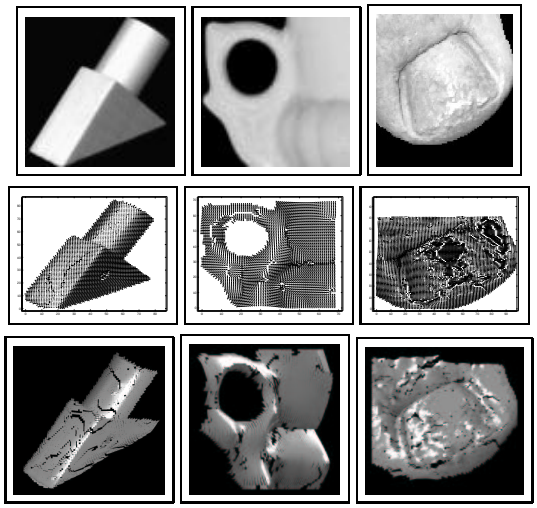
Figure 4: Results of the surface integration process.

observations. They mean that our graph-spectral method simply transfers errors in surface normal direction into errors in height, without producing structural noise artifacts.

To investigate the effect of noise further, Figure 4 we plot the mean-squared error for the reconstructed surface height as a function of the standard deviation of the added Gaussian noise. From the plots for the different surfaces shown in Figure 3, it is clear that the mean-squared error grows slowly with increasing noise standard deviation.

7.2 Real World Experiments

We have performed experiments using real world imagery. The images used in our study have been selected because they have proved problematic when conventional shape-from-shading algorithms are used. In the top panel of Figure 4 we show the three gray-scale images studied. These are a wooden block, an urn handle and a marble toe. The wooden block image illustrates the usefulness of the new method when dealing with ambiguous initial normal estimates (i.e. the planar faces of the block) while the other two images show the ability of the new method for smoothing on the minimum curvature direction keeping patch boundaries fixed. In the second row, we show the final needle-maps. The smoothing process improves the quality of the field of surface normals, but does not over-smooth across the patch boundaries. In the third row, we show the surfaces reconstructed from the smoothed fields of surface normals. It is evident that the smoothing process leads to surfaces which are in good subjective agreement with the image contents.

8 Conclusions

In this paper, we have described a graph-spectral algorithm shape-from-shading. The method commences from an initial surface normal estimate and performs smoothing along a curvature minimising path defined by the leading eigenvector of a curvature-based transition probability matrix. By traversing this path and applying some simple trigonometry we reconstruct the surface height function. Results on real world images reveal that the method is able to deliver accurate height reconstructions for complex surfaces.

References

[1] B. K. P. Horn and M. J. Brooks. Height and gradient from shading. *International Journal of Computer Vision*, 5(1):37–75, 1986.

[2] B. K. P. Horn and M. J. Brooks. The variational approach to shape from shading. *CVGIP*, 33(2):174–208, 1986.

[3] P. L. Worthington and E. R. Hancock. Needle map recovery using robust regularizers. *Image and Vision Computing*, 17:545–557, 1999.

[4] D. Mumford and J. Shah. Optimal approximations by piecewise smooth functions and associated variational problems. *Comm. in Pure and Appl. Math.*, 42(5):577–685, 1989.

[5] L. Lovász. Random walks on graphs: a survey. *Bolyai Society Mathematical Studies*, 2(2):1–46, 1993.

[6] Fan R. K. Chung. *Spectral Graph Theory*. American Mathematical Society, 1997.

[7] M. Doob D. Cvetković and H. Sachs. *Spectra of Graphs:Theory and Application*. Academic Press, 1980.

[8] M. P. Do Carmo. *Differential Geometry of Curves and Surfaces*. Prentice Hall, 1976.

[9] R. S. Varga. *Matrix Iterative Analysis*. Springer, second edition, 2000.