A Graph-Spectral Approach to Shape-from-shading

Antonio Robles-Kelly, Member, IEEE, and Edwin Hancock

Abstract—In this paper we explore how graph-spectral methods can be used to develop a new shape-from-shading algorithm. We characterise the field of surface normals using a weight matrix whose elements are computed from the sectional curvature between different image locations and penalise large change in surface normal direction. Modeling the blocks of the weight matrix as distinct surface patches, we use a graph seriation method to find a surface integration path that maximises the sum of curvature dependent weights and that can be used for the purposes of height reconstruction. To smooth the reconstructed surface, we fit quadrics to the height data for each patch. The smoothed surface normal directions are updated ensuring compliance with Lambert's law. The processes of height recovery and surface normal adjustment are interleaved and iterated until a stable surface is obtained. We provide results on synthetic and real-world imagery.

Index Terms—Shape-from-shading, graph seriation, graphspectral methods.

I. INTRODUCTION

Shape-from-shading is concerned with recovering local estimates of the surface normal direction, and hence the surface height function, from single images of shaded surfaces. The problem is a classical one which has been studied in both the computer vision and photogrammetry literature for a number of decades. In photogrammetric remote sensing the problem is known as photoclinometry, and one of the earliest practical applications was to recover the shapes of lunar maria from optical images [1]. More recent work has extended the technique to the radar domain and has developed radar-clinometry methods to analyse the topography of the surface of Venus using images delivered by the Magellan probe [2]. Remote sensing has also lead to the development of the technique referred to as photometric stereo, which involves surface normal and height recovery from sets of images recovered under different light source directions [3]. In computer vision, some of the first work was reported in the PhD theses by Horn [4] and by Krakauer [5]. Early developments in the field are well documented by Horn in the chapter "Height and Gradient from Shading" which appears in the collection of papers entitled Shape-from-Shading [6]. A more up to date review and comparative study is presented by Zhang et al. [7].

Stated succinctly the aim is to recover surface normal direction from measured image brightness using the image radiance equation. According to physics, the measured brightness depends on a number of factors including the viewer and light source directions, and the physical characteristics of the surface. To simplify the reflectance model, a number of assumptions are made. Firstly, the surface is assumed to be illuminated by a single point light source placed at infinity.

The authors are with the University of York, Heslington, York YO105DD, UK. E-mail:{arobkell, erh}@cs.york.ac.uk

As a result, there is no dispersion in the light source direction, and effects due to variations in the inverse squared distance to the light source over the surface are ignored. Secondly, the surface is assumed to be perfectly matte, and the observed surface radiance to follow Lambert's law. As a result for a surface of constant or known albedo, the observed brightness is independent of viewer direction, and depends only on the angle between the surface normal and the light source direction. Even under these restrictions, the shape-from-shading process is under constrained since the surface normal has two degrees of freedom corresponding to the slant and tilt angles of the local tangent plane to the surface, and the measured brightness provides only a single piece of information, or constraint, at each image location. Hence, additional constraints must be supplied. The most important of these is that the surface normals vary smoothly in direction across the surface. In addition, the integrability constraint is also sometimes imposed, and this requires that the Hessian matrix is symmetric at every location on the surface. This constraint is important if the surface is to be reconstructed from the recovered field of surface normals, since the estimated height will depend on the integration path if the constraint is not satisfied. Finally, constraints are provided by locations on the surface, e.g. the occluding boundary, where the surface normal directions are known.

Of course, there has been a considerable effort aimed at overcoming the restrictions listed above. For instance, methods have been developed for estimating the light source direction [8] and the albedo [9]. There has also been work aimed at understanding the effect of inter-reflections [10] and multiple light source directions. Finally, alternatives to Lambert's law have been investigated to allow for specular reflection from shiny surfaces, or boundary brightening effects for surfaces that exhibit roughness.

A. Related Literature

There are a number of perspectives from which the shapefrom-shading literature can be viewed. For instance, Zhang et al [7] 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, and this 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 height constraints, or by minimisation of an energy function. Although more robust to noise, global methods can have a tendency to oversmooth the recovered surface. There have been attempts to overcome this problem. For instance Zheng and Chellappa [9] have imposed a gradient consistency constraint that penalises differences between the image intensity gradient and the brightness gradient for the recovered surface. Worthington and Hancock [11] 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 normalised image radiance.

Another important issue underpinning shape-from-shading process is that of how to reconstruct 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 wave front 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 the shape-from shading or shape-from-texture process. Horn and Brooks [12], [13] realise surface height recovery as a post-processing 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 [14] average the surface normal directions to obtain a height estimate. A more elegant solution is proposed by Frankot and Chellappa [2] who project the surface normals into the Fourier domain to impose integrability constraints and surface height is recovered using an inverse Fourier transform. Klette and his co-workers have enhanced this approach by showing how more complex regularisation constraints can be formulated in the Fourier domain [15]. Leclerc and Bobick [16] have developed a direct numerical method for height recovery from shading information which uses curvature consistency constraints. Dupuis and Oliensis [17] 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 [18] who compute the relative height of the surface with respect to the highest intensity point. However, despite this work relatively little effort has been devoted to the problem of selecting the surface integration path.

Existing shape-from-shading methods can also be viewed from the perspective of the method used to solve the image radiance equation. For instance, the recovery of surface normal and height information using the image radiance equation can be posed as the solution of a partial differential equation. Here a number of methods including the use of level-sets [19] and the Hamilton-Jacobi-Bellman [20] equations have proved to be effective. However, one of the most frequently used methods is to pose the problem of recovering the field of surface normals, or the surface, as that of energy minimisation. Here one of the best known methods is the regularisation approach of Horn and Brooks [12]. According to the regularisation method there are distinct terms in the cost function to encourage compliance with the image radiance equation, and smoothness of the recovered solution. Horn and Brooks [12] show how the solution may be recovered using an iterative process which propagates constraints on the surface normal direction from the boundary to the interior of the object. At each image location, the update process involves adjusting the surface normal direction by an amount determined by the brightness error in the light source direction, and an amount determined by the local average surface normal. In practice it is the surface normal averaging effect that dominates, and this has a tendency to oversmooth the recovered surface. To overcome this problem, Worthington and Hancock [11] have replaced the quadratic smoothness penalty function of the Horn and Brooks method with one dictated by robust statistics.

Of course, when posed as the problem of minimising an energy function, then the question of how to solve the underlying optimisation problem is also one of pivotal importance. The update process used by Horn and Brooks [12] is a local gradient descent method which relies on parallel local updates of the surface normal direction. It is hence reminiscent of relaxation labeling. There are alternative methods which can be used to perform the update process. For instance, Crouzil et al. [21] have cast the problem into a Markov field setting, and have developed a multiscale method which can overcome problems of convergence to a local minimum. Recent work has also explored the use of graph-cuts to find exact minima of an energy function [22]. However, one of the methods that has not received attention is that of posing the method in a graph spectral setting and using eigenvector methods to recover the solution.

The idea underpinning graph spectral methods is to abstract the problem in hand using a weighted graph. Here the nodes represent the basic image entities, and the weighted edges represent affinity relations between the entities. By computing the eigenvalues and eigenvectors of the weight matrix, it is possible to find groups or clusters of entities. The graph spectral method is in fact one of energy minimisation since the eigenvectors can be shown to be minimisers of a quadratic form. In fact, graph spectral methods have recently proved highly effective in image processing and computer vision. Perhaps the best known method is that of Shi and Malik [23] which has shown how to locate image regions by recursively bisecting a weighted graph that represents the affinity of pairs of pixels. The method is based on the normalised cut. This is a measure of the relative weight of the edges connecting the two parts of the partition (the cut) to the weight assigned to the edges within the two parts of the bisection (the association). A relaxed solution to the bisection problem is found by locating the eigenvector associated with the second smallest eigenvalue of the Laplacian matrix (the degree matrix minus the affinity weight matrix). The advantage of graph-spectral methods is that they can be used to find approximate or relaxed solutions without the need for parallel iterative updates at the pixel sites. The method also obviates the need for complex search algorithms. However, although they have been applied to region segmentation and grouping problems, graph spectral methods have been applied neither to shape-from-shading,

nor to curve detection problems of the sort that arise in the determination of the optimal integration path.

B. Contribution

The aim in this paper is to use graph spectral methods to locate integration paths through fields of surface normals, and hence perform surface reconstruction using shape-fromshading. Our shape-from-shading method builds on the work of Worthington and Hancock [11] by constraining the surface normals to fall on irradiance or Monge cones whose axes point in the light source direction and whose apex angles are determined by the normalised image brightness. Our aim is to use a graph-spectral method to find an integration path through the field of surface normals that can be used for surface height recovery. To do this we abstract the field of surface normals using a curvature dependant weight matrix computed from adjacent pairs of surface normals. The higher the curvature, i.e. the difference in surface normal direction, the smaller the weight. The aim is to recover the integration path that maximises the sum of weights across the field of surface normals. This can be viewed as the problem of ordering the set of nodes in a graph in a sequence such that strongly correlated elements are placed next to one another, and this is often referred to as graph-seriation [24]. The seriation problem can be approached in a number of ways. Clearly the problem of searching for a serial ordering of the nodes, which maximally preserves the edge ordering is one of exponential complexity. As a result approximate solution methods have been employed. These involve casting the problem in an optimisation setting. Hence techniques such as simulated annealing and mean field annealing have been applied to the problem. It may also be formulated using semidefinite programming, which is a technique closely akin to spectral graph theory since it relies on eigenvector methods. However, recently a graph-spectral solution has been found to the problem. Atkins, Boman and Hendrikson [24] have shown how to use the leading eigenvector of the Laplacian matrix to sequence relational data. The method has been successfully applied to the consecutive ones problem and a number of DNA sequencing tasks. There is an obvious parallel between this method and steady state random walks on graphs, which can be located using the leading eigenvector of the Markov chain transition probability matrix. However, in the case of a random walk the path is not guaranteed to encourage edge connectivity. The spectral seriation method on the other hand does impose edge connectivity constraints on the recovered path,

Unfortunately, the analysis of the seriation problem presented in the paper by Atkins, Bowman and Hendrikson [24] is not directly applicable to our surface reconstruction problem. We provide an analysis which establishes the relationship between the seriation path and the leading eigenvector of the weight matrix. We use the leading eigenvector of the weight matrix to find a surface integration path through the field of surface normals. This path provides a relatively stable surface reconstruction since it avoids high curvature surface structures such as ridges and ravines. Height recovery can be performed by traversing the integration path and incrementing the height function by applying the trapezoid rule to the local slant and tilt angles.

The eigenvector analysis also allows us to recursively segment the surface into patches in a manner that is reminiscent of the Shi and Malik algorithm [23]. The sites visited by the integration path can be viewed as constituting a single surface patch. By taking the outer product of the leading eigenvector, we can compute the curvature dependent weight matrix associated with the sites belonging to the patch. The elements of the weight matrix associated with the patch may be set to zero, and the leading eigenvector of the residual weight matrix computed. This process may be repeated until the remaining pixels form patches of negligible size.

We fit a bi-quadric surface to the height values in each patch. For each pixel location, we can estimate a smoothed surface normal by computing the derivatives of the fitted patch. The properties of the bi-quadric patch ensure that the smoothed surface normals satisfy the integrability constraint. We follow Worthington and Hancock [11] by projecting the smoothed surface normals onto the nearest location on the local radiance cone.

Hence, our shape-from-shading algorithm differs from those reported in the literature in a number of important ways. First, we impose the Lambertian radiance equation as a hard constraint by demanding that the surface normals fall on radiance cones aligned in the light source direction. The optimisation problem underpinning our method is the location of a surface integration path through the field of surface normals that maximises the sum of curvature dependent weights. We use this path to both adjust the estimated height and the surface normal directions. In this respect the method is close in spirit to the method of Dupuis and Oliensis [17] since it is concerned with the recovery of a path on the surface. Since the recovery of surface normal and height information is implicit, it differs from the regularisation method of Horn and Brooks [12]. The method differs from the level sets methods since, height recovery is based on an open-curve and not a closed contour. The main advantage of using the spectral method is that is does not involve search over a potentially complex surface geometry, and that the eigenvector analysis simplifies both the recovery of the integration path and the location of surface patches.

The outline of this paper is as follows. In Section 2 we briefly review Worthington and Hancock's geometric framework for Lambertian shape-from-shading. Section 3 describes our curvature-based characterisation of the weight matrix for sites on the integration path. In Section 4 we describe how ideas from spectral graph theory can be used to recover an integration path from the weight matrix. Section 5 describes the relationship between the leading eigenvectors and surface patches. Sections 6 addresses the problem of height recovery, while Section 7 describes the surface fitting and surface normal adjustment method. In Section 8 we provide an overview of the resulting iterative algorithm for surface recovery. Experiments on real-world and synthetic data are described in Section 9. Finally, Section 10 provides some conclusions and suggested directions for further investigation.



Fig. 1. Curvature and path illustration (see text).

II. LAMBERTIAN REFLECTANCE

The starting point for the development of our graph-spectral approach to shape-from-shading is the geometric framework recently reported by Worthington and Hancock [11]. Underpinning this work is the observation 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 irradiance cones. Initially, the surface normals are positioned so that their projections onto the image plane point in the direction of the image brightness 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.

To be more formal let \vec{L} be a unit vector in the light source direction. The light source direction can be arbitrary, i.e. it is not restricted to the viewer direction $(0, 0, 1)^T$. However, an estimate of the light source direction is required. In most of our experiments the direction is recorded at the time the images are collected. An alternative is to use one of the methods reported in the literature for estimating the direction [9]. Suppose that E_i be the normalised brightness at the image location indexed *i*. Further, suppose that $\vec{N}_{i}^{(k)}$ is the corresponding estimate of the surface normal at iteration k of our algorithm. According to Lambert's law under conditions of constant albedo the *normalised* image irradiance is $E_i = \vec{N}_i^{(k)} \cdot \vec{L}$.

The consequence of Lambert's law is that the surface normal is constrained for fall on an irradiance or Monge cone, whose axis is the direction on the light source \vec{L} and whose opening angle is $\cos^{-1} E_i$.

Our aim is to adjust the surface normals by first smoothing them and then projecting them back onto the local irradiance cone. This smoothing is achieved by fitting a quadric patch to the reconstructed height data, and estimating the local normal to the fitted surface patch. Details of this process are described later in Section 7 of this paper. However, suppose that after local smoothing, the off-cone surface normal is $\vec{N}_i^{(k)}$. The updated on-cone surface normal which satisfies the image irradiance equation as a hard constraint is obtained via the rotation $\vec{N}_{i}^{(k+1)} = \Phi^{(k)} \vec{N}_{i}^{(k)}$. The matrix $\Phi^{(k)}$ 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 offcone normal and the light source direction. This angle is equal to 300

$$\theta_i^{(k)} = \cos^{-1} E_i - \cos^{-1} \frac{\vec{N}_i^{(k)} \cdot \vec{L}}{||\vec{N}_j^{(k)}|| \cdot ||\vec{L}||}$$

This rotation takes place about the axis whose direction is given by the vector $(u, v, w)^T = \vec{N}_i^{(k)} \times \vec{L}$. This rotation axis is perpendicular to both the light source direction and the offcone normal. Hence, the rotation matrix is

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

where $c = \cos \theta_i^{(k)}$, c' = 1 - c, and $s = \sin \theta_i^{(k)}$.

This simple approach has a number of shortcomings. For instance, it needs known, though arbitrary, light source direction and constant albedo. Additionally, the method is not applicable when there are significant departures from Lambert's law. This may be the case for shiny or glossy surfaces when specularities are present, or for rough surfaces where limb brightening is observed. However, there are reflectance models that account for these effects. For instance the Torrance-Sparrow [25] model may be used for specularties, the Wolff model [26] for shiny surface reflectance and the Oren-Nayar model [27] for rough surface reflectance. In principle these models may be used to construct reflectance maps and may be used to recover surface normals and height using Horn's method [12]. However, when using the method of irradiance cones, the models may be used to recover corrected Lambertian reflectance as a preprocessing step. Alternatively, the method reported in our recent paper may be used to perform Lambertian correction using empirically estimated reflectance distributions [28]. Finally, we note that using the image gradient direction to initialise the positions of the surface normals on the irradiance cone has the effect of biasing the recovered needle-map in the favour of convex surfaces.

III. AFFINITY MATRIX

The method outlined in the previous section provides us with estimates of the field of surface normals at locations



Fig. 2. Effect of varying the light source direction on the shape-from-shading algorithm: Input images (top row); Views of the recovered surfaces (second and third rows).



Fig. 3. Surface reconstruction for noise free data: Initial images (top row); ground-truth surfaces (second row); needle-maps (third row); reconstructed surfaces (fourth row); absolute difference between ground truth and reconstructed surfaces.

on the image plane Π . Our aim is to reconstruct the height function for surface *S* from this planar field of surface normals, under the assumption that the image of the surface is formed under orthographic projection. To realise this goal, we require an integration path. This path must traverse or connect the sites

of the pixel lattice, By traversing the path the relative surface height function can be reconstructed. This is done using the trapezium rule to increment the height using the distance traveled on the path and the known slant and tilt angles of the surface normals at different locations on the image plane. In the work reported here the path is one that optimises a graph-spectral criterion that penalises high curvature. To this end, we require a means of gauging the affinity of pixels based on an image plane approximation to the surface curvature.

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 plane-curve Γ can be viewed as the projection of a curve Γ_S from the surface S onto the plane Π . Whereas the curve Γ_S has both curvature and torsion, the planar curve Γ has only curvature. Further, suppose that $\kappa(l)$ is the sectional curvature of the curve Γ_S at the point Q with parametric co-ordinate l. Dropping the iteration dependant superscript used in the previous section, for the surface S the field of unit surface normals is sampled on the image plane Π , and the surface normal at the location $Q \in \Pi$ is denoted by \vec{N}_Q . Accordingly, and following do Carmo [29], 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 $\kappa(l)$ we require the differential of the surface or Hessian matrix $d\vec{N}_q: T_Q(S) \to T_Q(S)$. The maximum and minimum eigenvalues κ_1 and κ_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_S is \vec{n} and the unit tangent vector is $t_Q \in T_Q(S)$. The sectional curvature of Γ_S at Q is given by

$$z(l) = \frac{(\vec{t}_Q \cdot \vec{e}_1)^2 (\kappa_1 - \kappa_2) + \kappa_2}{\vec{n} \cdot \vec{N}_Q}$$
(1)

where $(\vec{t}_Q \cdot \vec{e}_1)^2 (\kappa_1 - \kappa_2) + \kappa_2$ is the normal curvature and $\phi = \arccos \vec{n} \cdot \vec{N}_Q$ is the angle between the curve normal and the surface normal.

Our aim is to approximate the curvature $\kappa(l)$ using the field of surface normals on the plane Π . In practice, we deal 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. To approximate the sectional curvature we make use of the surface normal directions. To commence, we note that if the path between the locations i and j can be approximated by a circle of radius R on the surface, then the approximate sectional curvature is $|\hat{\kappa}_{i,j}| = \frac{1}{R_{i,j}}$. Suppose that the surface normal directions at the pixel locations i and j are respectively \vec{N}_i and \vec{N}_j . Further suppose that the line connecting the pixel sites on the image plane is of length $l_{i,j}$. The approximating circle connects the points i and j, and has $l_{i,j}$ as the connecting chord. Hence, the change in direction of the radius vector of the circle is $\theta_{i,j} = \arccos \vec{N_i} \cdot \vec{N_j}$, and as a result $\cos \theta_{i,j} = \vec{N_i} \cdot \vec{N_j}$. If the angle $\theta_{i,j}$ is small, then we can make the Maclaurin approximation $\cos \theta_{i,j} \simeq 1 - \frac{\theta_{i,j}^2}{2} = \vec{N_i} \cdot \vec{N_j}$. Moreover, the small angle approximation to the radius of curvature of the circle is $R_{i,j} = \frac{l_{i,j}}{\theta_{i,j}}$ and hence

$$\hat{\kappa}_{i,j}^2 = \frac{2(1 - \vec{N}_i \cdot \vec{N}_j)}{l_{i,j}^2}$$
(2)

The geometry outlined above is illustrated in Figure 1a.

In the next section of this paper, we will present a graph spectral technique for locating the integration path Γ . The starting point for this analysis is a weight matrix that can be used to represent the affinity of pixels based on the curvature of the connecting path, or the difference in surface normal directions. To compute the elements of the affinity matrix, we associate with the pair of pixels a cost or energy that is equal to the product of the distance between the sites and the squared sectional curvature of the connecting path, i.e. $\mathcal{E}_{i,j} = \hat{\kappa}_{i,j}^2 |_{i,j}$.

Using the approximation to the sectional curvature, we find that the cost associated with the step from the pixel i to the pixel j is

$$\mathcal{E}_{i,j} = \frac{2}{l_{i,j}} (1 - \vec{N}_i . \vec{N}_j)$$
(3)

To pursue the graph-spectral analysis of the field of surface normals, we require a weight matrix which reflects both the sectional curvature and the connectivity of the pixel lattice. For the pixels indexed i and j we define the weight matrix to have elements

$$W(i,j) = \begin{cases} \exp[-\beta \mathcal{E}_{i,j}] & \text{if } j \in \mathcal{N}_i \\ 0 & \text{otherwise} \end{cases}$$
(4)

where \mathcal{N}_i is the set of pixels-neighbours of the pixel *i* and β is a constant. Hence, the curvature dependent weight is only non-zero if pixels abut one-another. Moreover, the weights are unity if the sectional curvature is zero and tend to zero as the curvature increases.

IV. GRAPH SERIATION

In the previous section, we showed how the change in surface normal directions could be used to compute the elements of a curvature dependent weight matrix. In this section, we describe how the leading eigenvector of this matrix can be used to locate an integration path that maximises the total weight. We pose this as a process of graph-spectral seriation.

To commence, we pose the problem in a graph-based setting. The set of pixel sites can be viewed as a weighted graph G = (V, E, W) with index-set V, edge-set $E = \{(i, j) | (i, j) \in V \times V, i \neq j\}$ and weight function $W : E \rightarrow [0, 1]$. Let the integration path commence at the node j_1 and proceed via the sequence of nodes $\Gamma = \{j_1, j_2, j_3, ...\}$.

The seriation problem as stated by Atkins, Boman and Hendrickson [24] is as follows. The aim is to find a path sequence for the nodes in the graph using a permutation π . The permutation gives the order of the nodes in the sequence. The sequence is such that the edge weight decreases as the path is traversed. Hence, if $\pi(i) < \pi(j) < \pi(k)$, then W(i, j) > W(i, k) and W(j, k) > W(i, k). This behaviour can be captured using the penalty function

$$g(\pi) = \sum_{i=1}^{|V|} \sum_{j=1}^{|V|} W(i,j)(\pi(i) - \pi(j))^2$$

By minimising $g(\pi)$ it is possible to find the permutation that minimises the difference in edge weight between adjacent nodes in the path, and this in turn sorts the edge weights into magnitude order. Unfortunately, minimising $g(\pi)$ is potentially

Fig. 4. Results on noise-free image data for the algorithm of Bichsel and Pentland: reconstructed surfaces (top row); absolute difference between ground truth and reconstructed surfaces (bottom row)

NP complete due to the combinatorial nature of the discrete permutation π . To overcome this problem, a relaxed solution is sought that approximates the structure of $g(\pi)$ using a vector $\vec{x} = (x_1, x_2, ...)$ of continuous variables x_i . Hence, the penalty function considered is

$$\hat{g}(\vec{x}) = \sum_{i=1}^{|V|} \sum_{j=1}^{|V|} W(i,j)(x_i - x_j)^2$$

The value of $\hat{g}(\vec{x})$ does not change a constant amount is added to each of the components x_i . Hence, the minimisation problem must be subject to constraints on the components of the vector \vec{x} . The constraints are that

$$\sum_{i=1}^{|V|} x_i^2 = 1 \quad \text{and} \quad \sum_{i=1}^{|V|} x_i = 0 \tag{5}$$

Atkins, Bowman and Atkinson show that the solution to this relaxed problem may be obtained from the Laplacian matrix. Let D be the diagonal matrix with elements $D(i, i) = \sum_{j=1}^{|V|} W(i, j)$ equal to the total weight of the edges connected to the node i. The Laplacian matrix is L = D - A. If $\vec{e} = (1, 1, 1..., 1)^T$ is the all-ones vector, then the solution to the minimisation problem is the vector

$$\vec{x} = \arg \min_{\vec{x}_*^T.\vec{e}=0, \vec{x}_*^T\vec{x}_*=1} \vec{x}_*^T L \vec{x}_*$$

$$= \arg \min_{\vec{x}_*^T.\vec{e}=0, \vec{x}_*^T\vec{x}_*=1} \sum_{i>j} W(i,j) (x_{*i} - x_{*j})^2 \quad (6)$$

When W is positive definite, then the solution is the Fiedler vector, i.e. the vector associated with the smallest non-zero eigenvalue of L. In fact, the associated eigenvalue minimises the Rayleigh quotient

$$\lambda = \arg\min_{x_*} \frac{\vec{x}_*^T L \vec{x}_*}{\vec{x}_*^T \vec{x}_*}$$

This eigenvalue satisfies the condition

$$\sum_{j=1}^{|V|} \sum_{i>j} W(i,j) (x_i - x_j)^2 = \lambda \sum_{i=1}^{|V|} x_i^2$$

Unfortunately, the procedure described above does not meet our requirements for three reasons. First, the penalty function $\hat{g}(\vec{x})$, does not impose edge connectivity constraints on the ordering computed during the minimisation process. Second, it implies no directionality in the transition between the nodes. Third, with our choice of weight function, minimisation of the penalty will not result in a path that minimises the curvature. To overcome these shortcomings, we turn our attention instead to maximising the cost function

$$\hat{g}_E(\vec{x}) = \sum_{i=1}^{|V|-1} \sum_{k=1}^{|V|} (W(i,k) + W(i+1,k)) x_k^2$$
(7)

When we combine the modified cost function with the constraints, we have that

$$\sum_{i=1}^{|V|-1} \sum_{k=1}^{|V|} (W(i,k) + W(i+1,k)) x_k^2 = \lambda \sum_{i=1}^{|V|-1} (x_i^2 + x_{i+1}^2)$$
(8)

By introducing the matrix

$$\Omega = \begin{bmatrix}
1 & 0 & 0 & 0 & \dots & 0 \\
0 & 2 & 0 & 0 & \dots & 0 \\
0 & 0 & 2 & 0 & \dots & 0 \\
\vdots & \vdots & \ddots & \ddots & \ddots & \vdots \\
0 & 0 & \dots & 0 & 2 & 0 \\
0 & 0 & \dots & 0 & 0 & 1
\end{bmatrix}$$
(9)

we can make the path connectivity requirement more explicit, and the maximiser of $g_E(\vec{x})$ satisfies the condition

$$\lambda = \arg\max_{x_*} \frac{\vec{x}_*^T \Omega W \vec{x}_*}{\vec{x}_*^T \Omega \vec{x}_*}$$
(10)



Fig. 5. Surface reconstruction for noisy data: Initial images (top row); needle-maps (second row); reconstructed surfaces (third row); absolute difference between ground truth and reconstructed surfaces (bottom row).

As a result, the leading eigenvalue λ_* of the weight matrix W is the maximiser of the cost function $\hat{g}_E(\vec{x})$. From the Perron-Frobenius theorem [30], it is known that, maximiser of this utility function is the leading (left) eigenvector of the matrix W. Moreover, since W is a real positive definite symmetric matrix, the associated eigenvector ϕ_* is unique. The Perron-Frobenius theorem ensures that the maximum eigenvalue $\lambda_* > 0$ of W has multiplicity one and, moreover, the coefficients of the corresponding eigenvector ϕ_* are all positive. As a result the remaining eigenvectors of W have at least one negative co-efficient and one positive co-efficient. If W is substochastic, ϕ_* is also known to be linearly independent of the all-ones vector e.

The elements of the leading eigenvector $\phi_{\,\ast}$ of W can be

used to construct an integration path. With our choice of the weight matrix, the components of ϕ_* decrease with increasing curvature of the seriation path. We commence from the node associated with the largest component of ϕ_* . We then sort the elements of the leading eigenvector such that they are both in decreasing magnitude order, and satisfy neighbourhood connectivity constraints on the pixel lattice. The procedure is a recursive one that proceeds as follows. At each iteration, we maintain a list of sites visited. At iteration k let the list of sites be denoted by \mathcal{L}_k . Initially, $\mathcal{L}_0 = j_o$ where $j_0 = \arg \max_j \phi_*(j)$, i.e. j_0 is the component of ϕ_* with the largest magnitude. Next, we search through the set of 8-neighbours of j_o to find the pixel associated with the largest remaining component of ϕ_* . If \mathcal{N}_{j_0} is the set of 8-neighbours of j_0 ,



Fig. 6. Results on noisy image data for the algorithm of Bichsel and Pentland: reconstructed surfaces (top row); absolute difference between ground truth and reconstructed surfaces (bottom row).

the second element in the list is $j_1 = \arg \max_{l \in \mathcal{N}_{j_0}} \phi_*(l)$. The pixel index j_1 is appended to the list of sites visited and the result is \mathcal{L}_1 . In the *k*th (general) step of the algorithm we are at the pixel site indexed j_k and the list of sites visited by the path so far is \mathcal{L}_k . We search through those 8-neighbours of j_k that have not already been traversed by the path. The set of pixel sites is $C_k = \{l | l \in \mathcal{N}_{j_k} \land l \notin \mathcal{L}_k\}$. The next site to be appended to the path list is therefore $j_{k+1} = \arg \max_{l \in C_k} \phi_*(l)$. This process is repeated until no further moves can be made. This occurs when $C_k = \emptyset$ and we denote the index of the termination of the path by T. The integration path Γ is given by the list of pixel sites \mathcal{L}_T .

To conclude this section, we return to the cost function underpinning the recovery of the integration path, to make the role of the optimisation process more explicit. When positioned at the site k, the next pixel visited satisfies the condition $j_{k+1} = \arg \max_{l \in C_k} \phi_*(l)$. Since ϕ_* is the leading eigenvector of W it satisfies the equation $W\phi_* = \lambda_*\phi_*$, As a result

$$j_{k+1} = \arg\max_{l \in C_k} \sum_{m \in \mathcal{N}_l} W(l,m)\phi_*(m)$$

Substituting for the definition of the weight in terms of the surface normals from using Equations (4) and (5),

$$j_{k+1} = \arg\max_{l \in C_k} \sum_{m \in \mathcal{N}_l} \left\{ \exp\left[-2\beta (1 - \vec{N}_m \cdot \vec{N}_l)\right] \right\} \phi_*(m)$$

When the angle between the surface normals is small, then we can write

$$j_{k+1} = \arg\max_{l \in C_k} \sum_{m \in \mathcal{N}_l} \left\{ 1 - 2\beta (1 - \vec{N}_m \cdot \vec{N}_l) \right\} \phi_*(m)$$

Hence, the path minimises the change in surface normal direction, and hence curvature, on the image plane Π .

V. EXTRACTING PATCHES

In practice the surface under study may have a patch structure. The patches may be identified by finding the blocks of the weight matrix induced under a permutation of the nodes. We find the blocks by computing the leading eigenvector of the weight matrix. The algorithm proceeds in an iterative fashion. The leading eigenvector of the current weight matrix represents a patch. The nodes with non-zero components in the leading eigenvector belong to the patch. The nodes are identified, and are then removed from further consideration by nulling their associated elements in the weight matrix. This process is repeated until all the principal patches are identified. This is the case when only an insignificant number of unassigned and unconnected nodes remain.

We commence by constructing the thresholded weight matrix W whose elements are defined as follows

$$\hat{W}(i,j) = \begin{cases} 0 & \text{if } W(i,j) << 1\\ W(i,j) & \text{otherwise} \end{cases}$$
(11)

The matrix \hat{W} is simply a thresholded version of the weight matrix W in which the vanishingly small elements are set to zero.

Our aim is identify groups of surface normals from a potentially noisy or ambiguous adjacency matrix \hat{W} which correspond to surface patches. Stated formally, suppose that in an image with an adjacency matrix \hat{W} there are *m* disjoint patches. Each such group should appear as a sub-block of the matrix \hat{W} . However, as a consequence of noise or errors in the shape-from-shading method which delivers the field of surface normals, these distinct groups or patches may be merged together. In other words, their corresponding sub-blocks are no longer disjoint.

Suppose that there are m distinct surface patches, each associated with an adjacency matrix $B^{(i)}$ where i is the patch index. If C represents a noise matrix, then the relationship between the observed weight matrix \hat{W} and the underlying



Fig. 7. Plot of the error percentage as a function of the variance for the four synthetic basic shapes.

block-structured weight matrix is $\hat{W} = B + C$ where $B = \mathcal{P}B_D \mathcal{P}^T$, \mathcal{P} is a permutation matrix and

$$B_D = \begin{pmatrix} B^{(1)} & 0 & \dots & 0\\ 0 & B^{(2)} & \ddots & \vdots\\ \vdots & \ddots & \ddots & 0\\ 0 & \dots & 0 & B^{(m)} \end{pmatrix}$$
(12)

is a block diagonal matrix in which $B^{(i)}$ is the sub-block corresponding to the patch indexed *i*.

To recover the matrix B, we turn to the eigenvector expansion of the matrix \hat{W} and write

$$\hat{W} = \lambda_* \phi_* \phi_*^T + \sum_{i=2}^{|V|} \lambda_i \vec{\phi}_i \phi_i^T$$
(13)

To identify the patches, we use the following iterative procedure. We initialise the algorithm by letting $\hat{W}^{(1)} = \hat{W}$. Further suppose that $\lambda_*^{(1)}$ is the leading eigenvalue and $\phi_*^{(1)}$ is the corresponding leading eigenvector of $\hat{W}^{(1)}$. The matrix $B^{(1)} = \lambda_*^{(1)} \phi_*^{(1)} \phi_*^{(1)T}$ represents the first block of \hat{W} , i.e. the most significant surface patch. 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 weight matrix $\hat{W}^{(2)} = \hat{W}^{(1)} - B^{(1)}$ in which the elements of the first patch are nulled. The leading eigenvector $\phi_*^{(2)}$ of the residual weight matrix $\hat{W}^{(2)}$ is used to compute the second block $B^{(2)} = \lambda_*^{(2)} \phi_*^{(2)} \phi_*^{(2)T}$. The process is repeated iteratively to identify all of the principal blocks of \hat{W} . At iteration p, $\phi_*^{(p)}$ is the leading eigenvector of the residual weight matrix $\hat{W}^{(p)}$, and the p^{th} block is $B^{(p)} = \lambda_*^{(p)} \phi_*^{(p)} \phi_*^{(p)T}$. The patch indexed n is the set of nodes for which the components of the leading eigenvector $\phi_*^{(p)}$ are non-zero. Hence, the indexset for the p^{th} patch is $S_p = \{i | \phi_*^{(p)}(i) \neq 0\}$. 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 $\phi_{*}^{(p)} \cdot \phi_{*}^{(l)} = 0$, where $p \neq l$.

The process of patch extraction is iterated until the associated regions are of insignificant size. In practice, this is the case when the number of pixels associated with a path is less than 5.

VI. HEIGHT RECOVERY

Our surface height recovery algorithm proceeds along the sequence of pixel sites defined by the order of the coefficients of the scaled leading eigenvector associated with the separate patches. For the p^{th} patch, the path is $\Gamma_p = (j_1^p, j_2^p, j_3^p, ...)$ where the order is established using the method outlined in Section 4. As we move from pixel-site to pixel-site defined by this path we increment the surface height-function. In this section, we describe the trigonometry of the height incrementation process.

At step n of the algorithm, we make a transition from the pixel with path-index j_{n-1}^p to the pixel with path-index j_n^p . The distance between the pixel-centres associated with this transition is

$$d_n^p = \sqrt{(x_{j_n^p}^2 - x_{j_{n-1}^p})^2 + (y_{j_n^p} - y_{j_{n-1}^p})^2}$$
(14)

This distance, together with the surface normals $\vec{N}_{j_n^p} = (N_{j_n^p}(x), N_{j_n^p}(y), N_{j_n^p}(z))^T$ and $\vec{N}_{j_{n-1}^p} = (N_{j_{n-1}^p}(x), N_{j_{n-1}^p}(y), N_{j_{n-1}^p}(z))^T$ at the two pixel-sites may be used to compute the change in surface height associated with the transition. Assuming that the two pixel sites are connected by a plane, whose normal is $\vec{N}_{j_n^p}$, we can write $N_{j_n^p}(x)(x-x_{j_n^p}) + N_{j_n^p}(y)(y-y_{j_n^p}) + N_{j_n^p}(z)(z-z_{j_n^p}) = 0$. As a result, the height increment is given by

$$h_{n}^{p} = \frac{r_{n}^{p}}{2} \left((x_{j_{n-1}^{p}} - x_{j_{n}^{p}}) \left(\frac{N_{j_{n-1}^{p}}(x)}{N_{j_{n-1}^{p}}(z)} + \frac{N_{j_{n}^{p}}(x)}{N_{j_{n}^{p}}(z)} \right) + (y_{j_{n-1}^{p}} - y_{j_{n}^{p}}) \left(\frac{N_{j_{n-1}^{p}}(y)}{N_{j_{n-1}^{p}}(z)} + \frac{N_{j_{n}^{p}}(y)}{N_{j_{n}^{p}}(z)} \right) \right)$$
(15)

If the height-function is initialised by setting $z_{j_0^p} = 0$, then the centre-height for the pixel with path-index j_{n+1}^p is

$$z_{j_{n+1}^p} = z_{j_n^p} + h_n^p \tag{16}$$

Once the surface normals that belong to the individual patches have been integrated together, then we merge them together to form a global surface. Suppose that S_p is the integrated surface for the p^{th} patch. We compute the mean height for the pixels belonging to the boundary for this patch. We merge the patches together by ensuring that abutting patches have the same mean boundary height. The geometry of this procedure is illustrated in Figure 1.

VII. REGION QUADRIC PATCH FITTING AND SURFACE NORMAL UPDATING

Once the height values are available for each pixel site in a patch, then we perform smoothing. We do this by fitting a local bi-quadric function to the height data for the patch sites. To do this we employ a simple least-squares fitting method.

Suppose that the pixels belonging to the patch indexed p are denoted by the set S_p . We aim to fit the bi-quadric

$$Q_p(x_i, y_i) = a_{1,p} + a_{2,p}x + a_{3,p}y_i + a_{4,p}x_i^2 + a_{5,p}x_iy_i + a_{6,p}y_i^2$$
(17)

to the height data for the sites in the patch. Let $M_i = (1, x_i, y_i, x_i^2, x_i y_i, y_i^2)^T$ be a vector of co-ordinate moments and let $V_p = (a_{1,p}, \dots, a_{6,p})^T$ be a vector of parameters. The least-squares parameter vector for the quadric patch satisfies the condition

$$V_p^* = \arg\min_{V_p} \sum_{i \in S_p} \left[z_i - V_p^T M_i \right]^2$$
(18)

In matrix form the solution vector is given by $V_p^* = (F_p^T F_p)^{-1} F_p^T$ where $F_p = \sum_{i \in S_p} z_i M_i M_i^T$. The surface fitting process can be viewed as one of smooth-

The surface fitting process can be viewed as one of smoothing over the extent of the patches defined by the leading eigenvectors of the weight matrix. As well as smoothing the local height estimates, the process also implicitly smooths the surface normal directions, since these can be computed from the gradient of the fitted local quadric patch. The estimate of the unit surface normal at the point (x_i, y_i) on the patch p is

$$\vec{N}_{i} = \frac{1}{R_{i}} \left[\frac{\partial Q_{p}(x_{i}, y_{i})}{\partial x_{i}}, \frac{\partial Q_{p}(x_{i}, y_{i})}{\partial y_{i}}, 1 \right]^{T} = \frac{1}{R_{i}} \left[a_{2,p} + 2a_{4,p}x_{i} + a_{5,p}y_{i}, a_{3,p} + a_{5,p}x_{i} + 2a_{6,p}y_{i}, 1 \right]^{T}$$
(19)

where
$$R_i = \sqrt{\left(1 + \left(\frac{\partial Q_p^{(k)}(x_i, y_i)}{\partial x_i}\right)^2 + \left(\frac{\partial Q_p^{(k)}(x_i, y_i)}{\partial y_i}\right)^2}$$
.

It is interesting to note that since

$$\frac{\partial^2 Q_p(x_i, y_i)}{\partial x_i \partial y_i} = \frac{\partial^2 Q_p(x_i, y_i)}{\partial y_i \partial x_i} = a_{5,p} \tag{20}$$

the field of smoothed surface normals satisfy the integrability constraint.

To update the surface normal direction so that they satisfy the image radiance equation, we rotate them using the matrix Φ as described in Section 2. The updated surface normals are no longer guaranteed to satisfy the integrability constraint, but fall on the closest location on the appropriate radiance cone.



Fig. 8. Results of the algorithm on the image of Michelangelo's Moses.

VIII. ALGORITHM DESCRIPTION

Having described the different steps of the surface reconstruction process, in this section we summarise how the steps are combined together in an algorithm. The sequence of processing steps is as follows:

- Step 0: An estimate of the light source direction is made either from the known geometry of the imaging setup or from the statistics of the image gradient [9]. The surface normals are placed in their initial positions on the irradiance cones. To do this we align them in the directions of the image gradient. The gradient is computed by first smoothing the grey-scale image by fitting local quadric patches to the raw image intensity. The smoothed image gradient is found from the derivatives of the fitted patch.
- Step 1: From the initial field of surface normals, we compute the sectional curvatures and hence the weight matrix. The blocks of the matrix are surface patches. The leading eigenvector of the weight matrix for each block is the patch integration path. Using the patch integration paths, we recover estimates of the surface height.
- Step 2: For the sites in each patch, we fit a quadric patch to the available height estimates. The fitted surface patches are used to compute an estimate of the surface gradient. At each location, the gradient estimate is used to adjust the position of the surface normals on their respective irradiance cones.
- Iteration: Steps 1 and 2 are iterated until a stable set of surface height estimates are located.

IX. EXPERIMENTS

Our experiments with the new graph-spectral method for shape-from-shading are divided into two parts. We commence with a study based on synthetic data which is aimed at establishing the noise sensitivity and failure modes of the method. In the second part of the study we focus on the behaviour of the method when confronted with real world data. Here we study images of classical statues.

A. Synthetic Data

In this section we provide some experiments on synthetic data. We commence by turning our attention to the stability of the recovered surface under varying light source direction. To this end, we have generated a set of 6 synthetic images of a Lambertian teapot illuminated with a single light source positioned in the direction $\vec{s}_D = [\sin(\theta_s), 0, \cos(\theta_s)]^T$. For our experiments, we have varied the angular variable θ_s in increments of 6° between -30° and 0° .

In the top row of Figure 2, we show the images in our dataset. Here, we have ordered the images, from left-to-right, in decreasing θ_s . In the remaining rows, we show two views of the recovered surfaces for each of the input images. The views corresponding to different illuminant directions are in good agreement with one another. This is an important observation since it indicates that the surface is stable to variations in the light source direction. Furthermore, despite some errors across the surface boundaries, the shape of the teapot is well recovered.

Next, we determine the accuracy of the surface reconstruction method. To this end we have generated synthetic surface height data. From the surfaces, we have computed the field of surface normal directions. A light source direction is then selected and the surfaces have been rendered using the Lambertian reflectance process outlined in Section 2. We have then applied the graph-spectral shape-from-shading method to the resulting synthetic images. We compare the resulting height estimates with the height data for the original surfaces. We also provide a comparative study using the shape-fromshading height recovery method of Bichsel and Pentland [18] as an alternative to our graph-spectral shape-from-shading algorithm. We have chosen the algorithm of Bichsel and Pentland since, from Zhang et al [7], appears to deliver reasonable results on a wide variety of images. Moreover, this shape-from-shading algorithm is local a technique that implicitly recovers the surface.

In Figure 3 we show the results obtained for a series of different surfaces. In the first, second and third rows we show the Lambertian shading, height data and surface normals for the synthetic surfaces. In the fourth row of the figure we show the surface reconstructed by applying our shape-from-shading method to the images in the top row. The bottom row of the figure 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

Next, we compare our results with those for the shape-fromshading height recovery method of Bichsel and Pentland [18]. In the top row of Figure 4, we show the recovered surface height while the bottom row shows and the absolute error between the ground-truth and the reconstructed surface height. We do not show the needlemaps due to the fact that the Bichsel and Pentland's algorithm does not deliver a field of surface normals at output. From the plots, it is clear that our algorithm out-performs the one of Bichsel and Pentland. Furthermore, in contrast with our results, the mean-squared error is 6.1% for the dome, 11.9% for the ridge, 10.2% for the torus and 26.8% for the volcano.

We have repeated these experiments under conditions of controlled noise. To do this we have added random measurement errors to the raw image brightness. The measurement errors have been sampled from a Gaussian distribution with zero mean and known variance. In Figure 5 we show the result of reconstructing the surfaces shown in Figure 2, when brightness errors have been added. The order of the rows is the same as Figure 3. The results for the shape-from-shading height recovery method of Bichsel and Pentland are shown in Figure 6. The row-order is the same as in Figure 4.

From Figure 5, it is clear that, for all four surfaces, our algorithm preserves the gross structure of the surfaces under study. However, the recovered height is clearly noisy. The height difference plots are relatively unstructured. These are important observations. They mean that our graph-spectral method simply transfers errors in surface normal direction into errors in height, without producing structural noise artifacts. In contrast, the surfaces recovered by the Bichsel and Pentland's algorithm are poorer for all the shapes except the ridge. Furthermore, the algorithm of Bichsel and Pentland recovers broken, uneven surfaces due to convex-concavity surface inversion errors introduced by the noise added to the image brightness.

To investigate the effect of noise further, in Figure 7 we plot the mean-squared error for our graph-spectral shape-fromshading algorithm and the height recovery method of Bichsel and Pentland. Here, we show the mean-squared error for the reconstructed surface height as a function of the standard deviation of the added Gaussian noise. The different curves are for the different surfaces shown in Figures 3 and 4. From the plots for the different surfaces shown in Figure 7, it is clear that, for our algorithm, the mean-squared error grows slowly with increasing noise standard deviation. The torus and the volcano give the poorest errors, while the ridge and the dome give the smallest errors. This is a reflection of the fact that the torus and the volcano are the more structured surfaces. There are two features to note from this plot. First, with the exception of the dome, for each of the shapes the error is larger than in the case of our graph-spectral method. Second,



Fig. 9. Results of the algorithm on the image of the Three Graces Relief Fig fragment.

Fig. 10. Results of the algorithm on the image of the Beethoven bust.

the rate of increase in error with noise standard deviation is greater compared to the graph-spectral case.

B. Real World Data

We have experimented with a variety of real world images, but in this paper we concentrate mainly on images of classical statues. These objects are predominantly Lambertian, although there are local specular highlights and local albedo variations. In principle we can overcome both of these problems. In a recent paper, we have described a probabilistic method for specularity removal which uses the Torrance and Sparrow model to perform Lambertian reflectance correction for shiny objects [31]. Local albedo changes can be accommodated using brightness normalisation or histogram equalisation. The objects studied are a detail of Michaelangelo's Moses and a section of the relief "Three Graces". We have also used an image of a bust of Beethoven from the University of Central Florida shape-from-shading data-base [7].

In Figure 8 we show our first sequence of results. The panels of the figure are organised as follows. In Figure 8a we show the original image used as input to the shape-from-shading process. This is a side view of the head of the statue "Moses". Figures 8b and c show the arrangements of quadric patches after one and four iterations of the algorithm. In these images, the different quadric patches are coded in different colours. Figures 8 d and e show the initial and final needle maps used for the purposes of surface integration. Figures 8f and g show the initial and final reconstructed surfaces. Finally, Figures 8h

and 8i show views of the reconstructed surface from different viewpoints. Initially, the set of surface patches is fragmented, lack coherence and do not reproduce the surface detail well. However, after four iterations the reconstructed surface is more continuous and the fine detail of the object is well reproduced. There are a number of features that are worth noting from the panels of the figure. First, the organisation of the surface normals and the arrangement of patches both improve as the algorithm iterates. Second, from the different surface views it is clear that the surface structure is well reconstructed. For instance the shape of the nose, particularly in the proximity of the nostrils, is well reproduced. Moreover, the fine structure of the beard, the detail in the eye-sockets and the shape of the cheek bones are all well reconstructed.

In Figure 9, we repeat the sequence of panels for a section of the relief "The Three Graces". However, only two of the subjects of the original sculpture are visible in the section of image used. The same iterative improvement in the quality of the surface normals and the patch arrangement is clear. In this case the algorithm converged after 3 iterations. Moreover, the detail of the relief is well reproduced. The legs, buttocks and indentation in the back of the left-hand figure are all well reconstructed.

Figures 10 shows analogous results for an image of a bust of Beethoven. Here the path structure is particularly clear and corresponds well to the topographic structure of the surface. For instance, the eye sockets correspond to distinct patches. Both the patch structure and the structure of the needle maps are improved as the algorithm iterates. Initially, little of the surface structure is evident. However, after the algorithm has converged the structure of the hair and the boundary of the cheeks have become well defined.

Finally, we have compared our algorithm with two alternatives. The first of these is the shape-from-shading height recovery method of Bichsel and Pentland [18]. The second method used in our comparison is a purely geometric surface integration method, which like the technique reported in this paper takes the needle-maps delivered by the Worthington and Hancock algorithm as input. This method uses the trapezoid rule to increment the height from equal height reference contours. The method was developed in conjunction with a study of terrain reconstruction using radar shape-from-shading and a full description can be found in [32]. The results for these two alternatives are shown in Figures IX-B and IX-B.

For the Bichsel and Pentland method, the overall quality of surface detail is poorer, and the method is unstable in shadows. For instance, in the case of the Moses only the chin and mouth appear to be well reconstructed. The height of the cheek is overestimated, and convex parts of the surface have become concave. There are similar problems with convexity-concavity inversion for the Three Graces. The results obtained by the geometric surface integration method are better than those obtained by Bichsel and Pentland and do not suffer so badly from the problems of surface inversion. Although the detail is better than that obtained using the Bichsel and Pentland method, the results delivered by the geometric integration method do not contain the fine detail delivered by our graph spectral method. For instance, the hair and fabric details of Beethoven are not well reproduced and the structures of the beard of Moses is not well recovered.

X. CONCLUSIONS

In this paper, we have described a graph-spectral algorithm for shape-from-shading. We constrain the surface normals at each image location to fall on an irradiance cone whose



Fig. 11. Results for the Bichsel and Pentland's algorithm.



Fig. 12. Results for the geometric surface integration algorithm.

axis is the light source direction and whose apex angle is determined by the measured image brightness. The method uses the eigenvectors of a weight matrix to both locate surface patches and identify a curvature minimising path for surface integration, and hence height recovery. By fitting quadric surfaces to the height data for pixel sites contained within patches, we perform surface smoothing. We update the surface normal directions by rotating them so that they point in the direction of the fitted surface gradient. The surface integration and surface normal adjustment steps are iterated until stable height estimates are recovered. The method proves effective for reconstructing surfaces from single views of 3D objects, and gives subjectively better results than a number of alternative shape-from-shading methods.

There are a number of ways in which the reported method can be further extended and improved. First, we aim to incorporate constraints provided by the permitted differential structure of the field of principal curvature directions. Second, we are extending the Lambertian reflectance model to accommodate effects such as local specularity, coherent limb brightening and rough surface scattering. Third, we are investigating whether the iterative process developed in the paper can be posed as statistical estimation using the apparatus of the EM algorithm. This would provide a more principled basis for the interleaved iterative steps for height recovery and surface normal adjustment. Work aimed at addressing these points is underway and will be reported in due course.

REFERENCES

- T. Rindfleisch, "Photometric method for lunar topography," *Photogrammetric Engineering*, vol. 32, no. 2, pp. 262–277, 1966.
- [2] R. T. Frankot and R. Chellappa, "A method of enforcing integrability in shape from shading algorithms," *IEEE Transactions on Pattern Analysis* and Machine Intelligence, vol. 4, no. 10, pp. 439–451, 1988.
- [3] R. J. Woodham, "Analysing images of curved surfaces," Artificial Intelligence, no. 17, pp. 117–140, 1981.
- [4] B. K. P. Horn, "Shape from shading: A method for obtaining the shape of a smooth opaque object from one view," in *MIT AI-TR*, 1970.

- [5] J. L. Krakauer, "Computer analysis of visual properties of curved objects," in *Massachusetts Institute of Technology*, AI-TR, 1971.
- [6] B. K. P. Horn and M. J. Brooks, Shape from Shading. MIT Press, 1989.
- [7] R. Zhang, P. S. Tsai, J. E. Cryer, and M. Shah, "Shape from shading: A survery," *IEEE Trans. on Pattern Analysis and Machine Intelligence*, vol. 21, no. 8, pp. 690–706, 1999.
- [8] A. P. Pentland, "Linear shape from shading," Int. Journal of Comp. Vision, vol. 1, no. 4, pp. 153–162, 1990.
- [9] Q. Zheng and R. Chellappa, "Estimation of illuminant direction, albedo, and shape from shading," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 13, no. 7, pp. 680–702, 1991.
- [10] D. Forsyth and A. Zisserman, "Mutual illumination," in Proc. of the IEEE Conf. on Comp. Vision and Pat. Recognition, 1989, pp. 466–473.
- [11] P. L. Worthington and E. R. Hancock, "New constraints on datacloseness and needle map consistency for shape-from-shading," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 21, no. 12, pp. 1250–1267, 1999.
- [12] B. K. P. Horn and M. J. Brooks, "The variational approach to shape from shading," *CVGIP*, vol. 33, no. 2, pp. 174–208, 1986.
 [13] B. K. P. Horn, "Height and gradient from shading," *International Journal*
- [13] B. K. P. Horn, "Height and gradient from shading," *International Journal of Computer Vision*, vol. 5, no. 1, pp. 37–75, 1990.
- [14] Z. Wu and L. Li, "A line-integration based method for depth recovery from surface normals." *CVGIP*, vol. 43, no. 1, pp. 53–66, July 1988.
- [15] T. Wei and R. Klette, "Theoretical analysis of finite difference algorithms for linear shape from shading," in CAIP, 2001, pp. 638–645.
- [16] Y. G. Leclerc and A. F. Bobick, "The direct computation of height from shading," in *Proceedings of Computer Vision and Pattern Recognition*, 1991, pp. 552–558.
- [17] P. Dupuis and J. Oliensis, "Direct method for reconstructing shape from shading," in *Proc. of the IEEE Conf. on Comp. Vision and Pattern Recognition*, 1992, pp. 453–458.
- [18] M. Bichsel and A. P. Pentland, "A simple algorithm for shape from shading," in *Proc. of the IEEE Conf. on Comp. Vision and Pattern Recognition*, 1992, pp. 459–465.
- [19] R. Kimmel, K. Siddiqqi, B. B. Kimia, and A. M. Bruckstein, "Shape from shading: Level set propagation and viscosity solutions," *International Journal of Computer Vision*, no. 16, pp. 107–133, 1995.
- [20] E. Rouy and A. Tourin, "A viscosity solution approach to shape-fromshading," *SIAM Journal on Numerical Analysis*, vol. 29, no. 3, pp. 867– 884, 1992.
- [21] A. Crouzil, X. Descombes, and J. D. Durou, "A multiresolution approach for shape from shading coupling deterministic and stochastic optimization," 2003, to appear in *IEEE Trans. on Pattern Analysis and Machine Intelligence.*
- [22] J. Y. Chang, K. M. Lee, and S. U. Lee, "Shape from shading using graph cuts," in *Proc. of the Int. Conf. on Image Processing*, 2003.
- [23] J. Shi and J. Malik, "Normalized cuts and image segmentation," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 22, no. 8, pp. 888–905, 2000.
- [24] J. E. Atkins, E. G. Bowman, and B. Hendrickson, "A spectral algorithm for seriation and the consecutive ones problem," *SIAM Journal on Computing*, vol. 28, no. 1, pp. 297–310, 1998.
- [25] K. E. Torrance and E. M. Sparrow, "Theory for off-specular reflection from roughened surfaces," *Journal of Optical Society of America*, vol. 57, no. 9, pp. 1105–1114, 1967.
- [26] L. B. Wolff, "Polarization-based material classification from specular reflection," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 12, no. 11, pp. 1059–1071, 1990.
- [27] M. Oren and S. K. Nayar, "Generalization of the lambertian model and implications for machine vision," *Int. Journal of Comp. Vision*, vol. 14, no. 3, pp. 227–251, 1995.
- [28] A. Robles-Kelly and E. R. Hancock, "Lambertian correction for rough and specular surfaces," in *Video, Visualisation and Graphics*, 2003, pp. 213–220.
- [29] M. P. Do Carmo, Differential Geometry of Curves and Surfaces. Prentice Hall, 1976.
- [30] R. S. Varga, Matrix Iterative Analysis, 2nd ed. Springer, 2000.
- [31] H. Ragheb and E. R. Hancock, "A probabilistic framework for specular shape-from-shading," *Pattern Recognition*, vol. 36, no. 2, pp. 407–427, 2003.
- [32] A. G. Bors, E. R. Hancock, and R. C. Wilson, "Terrain analysis using radar shape-from-shading," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 25, no. 8, pp. 974–992, 2003.