PERGAMON

Pattern Recognition 32 (1999) 1211–1223

PATTERN RECOGNITION

THE JOURNAL OF THE PATTERN RECOGNITION SOCIETY

Consistent topographic surface labelling

Richard C. Wilson*, Edwin R. Hancock

Department of Computer Science, University of York, York, Y01 5DD, UK

Received 26 August 1998; received in revised form 14 October 1998; accepted 14 October 1998

Abstract

This paper describes work aimed at consistently labelling surface facets using topographic classes derived from mean and Gaussian curvature measurements. There are two distinct contributions. Firstly, we develop a statistical model which allows label probabilities to be assigned to the different topographic classes. These probabilities capture uncertainties in the computation of surface curvature from raw surface normal information. The probabilities are computed using propagation of variance from the surface normal measurements. The second contribution is to demonstrate how topographic surface labelling can be realised using probabilistic relaxation. The key ingredient is to develop a constraint dictionary for the feasible configurations of the topographic labels that can occur on neighbouring faces of the surface mesh. These constraints relate to the legal adjacency of different topographic structures together with the smoothness and continuity of uniform regions.

Keywords: *H-K* curvature; Relaxation labelling; Curvature dictionary

1. Introduction

Topographic labels derived from the mean and Gaussian curvatures are widely exploited as a means of representing the differential structure of surfaces [1–5]. The surface labelling task commences by computing the Hessian matrix from the directional second derivatives of the surface. The mean curvature is related to the trace of the Hessian, while the Gaussian curvature is computed from its determinant. Once the curvatures are to hand, then topographic labels are assigned on the basis of whether the two curvatures are positive, negative or consistent with zero. Based on this coarse quantisation of the curvature information, the surface data may be segmented into meaningful topographic structures such as ridges or valleys, saddle points or lines, and, domes or cups. These structures can be further organised into simply connected elliptical or hyperbolic regions which are separated from one-another by parabolic lines. Unfortunately, because the Hessian matrix is based on second-derivatives the reliable estimation of surface curvature, and hence the extraction of topographic structure, has proved to be a task of notorious difficulty in the analysis and range or volumetric imagery [6–8]. Some of the limitations of the alternative strategies for curvature estimation were unearthed in the comparative study of Flynn and Jain [9].

It is for these reasons that strategies aimed at circumventing the direct estimation of the second-derivatives have been developed. One of the most popular approaches is to approximate the surface by a low-order piecewise continuous surface [10–12]. The parameters of the best-fit local patches are used to estimate curvature.

*Corresponding author. Tel.: +44 1904 432762; Fax: +44 1904 432767; E-mail: wilson@minster.york.ac.uk

0031-3203/99/$20.00
PII: S0031-3203(98)00146-0

Although, it was originally developed as a means of extracting a topographic description of gray-scale edge features, Haralick's facet model has found widespread application in the interpretation of range imagery [10]. Haralick's idea is to fit a set of orthogonal polynomials to image gray-scales. The fit parameters are used to identify salient topographic features such as edges (ravines) and lines (ridges). Moreover, there is a simple error analysis for the extracted model parameters. In a genuine surface interpretation problem, Bolle and Cooper [11] fit Haralick's orthogonal polynomials to 3D data to identify topographic primitives. The parameters of these polynomials are used to calculate the Hessian. A Bayesian classifier is used to segment the surface into scene primitives such as cylinders and spheres.

An alternative approach is to extract curvature information using a bi-quadric patch. Several authors have developed variants of this idea. For instance, Besl and Jain adopt a hierarchical fitting technique [1]. Firstly, a local tangent plane is extracted by identifying the principal component axes for the distribution of surface datapoints using a local support neighbourhood, Next, the plane-fit is refined using a cubic patch.

Several authors have directed their attention towards understanding the statistical uncertainties involved in the assignment of mean and Gaussian curvature labels. For instance, Abdelmalek [12] focusses on the uncertainties that exist when the Hessian is estimated using orthogonal polynomials. This analysis determines an upper error bound on the values of mean and Gaussian curvature. The analysis commences by imposing the condition that the depth error at each point must be smaller than twice the noise-variance. However the resulting error bounds are an order of magnitude greater that the actual errors. A more faithful description of realistic data-point error distributions is obtained by Hilton et al. [13] who address the issue of analysis of variance to improve the statistical fidelity of the fitting process.

Despite these efforts aimed at improving the reliability of curvature estimation, the problem of how to refine inconsistent topography has received less attention. In essence, local surface-fitting does not guarantee that the extracted curvature estimates are globally consistent when viewed from the constraints imposed by the topographic structure of smooth continuous surfaces. Interrogation of the literature reveals that it is only Sander and Zucker [4] who have made any serious attempt at exploiting the idea of curvature consistency to improve the recovery of a consistent differential surface structure. Their idea has been to iteratively update local Darboux frames by imposing the constraint that the principal curvature directions should vary smoothly across the surface. The initial estimates of the Hessian required in this analysis are derived from the least squares fitting of bi-quadric patches. Consistency is measured by the smoothness of the field of principle curvature directions. However, there is no attempt to reconcile the quality of the recovered surface description with the underlying statistical uncertainties in the raw surface data. Although their subsequent work has focussed on how to recover singularities in the principal curvature field [5], there is no real attempt to exploit the topographic structure of the surface to improve curvature consistency.

1.1. Paper outline

Based on this review of the literature our observations in this paper are two-fold. Firstly, the majority of topographic labelling techniques opt to make statistical estimates of the Hessian by local surface fitting. The main disadvantage of surface fitting is that the process may have an over-smoothing effect on fine surface detail. In particular, over-smoothing due to the use of an excessively large support neighbourhood can erode important features such as ridge-lines or surface cusps. Our second observation is that there has been no consolidated attempt to incorporate statistical information into the refinement of topographic surface labels. In other words, the process is again driven almost exclusively by the smoothness or continuity assumption. This is a disappointing omission. It means that important constraints provided by the adjacency structure of the topographic labels are invariably overlooked when attempts are made to impose curvature consistency.

These two observations provide the basic motivation for the novel contributions of the study reported in this paper. In the first instance, rather than using a conventional surface-fitting technique, we make a direct estimate of the Hessian matrix using statistics derived from surface normals. The idea is to use the method of least squares to estimate the directional second derivatives using differences in the normal components over a local support neighbourhood. This offers the advantage that it simplifies the problem of estimating the statistical uncertainties inherent in the elements of the Hessian matrix. Once the least-squares estimates are to hand, then their associated covariance parameters can be computed from the fit-residuals. We propagate the measured covariance structure for the normals through to the estimation of mean and Gaussian curvature.

Our motivation in embarking on this statistical analysis is to provide probability estimates for the various topographic-labels at individual points on the surface. These label-probabilities are an essential pre-requisite for the subsequent consistent labelling of the surface. It is here that we make our second novel contribution. We are interested in improving the consistency of the labelled surface using relaxation labelling. Although the literature contains a plethora of alternative schemes, these share the common feature of requiring a representation of the constraints which pertain to the labelling problem in-hand. When the labelling problem is of a highly

structured nature, as is the case in assigning mean and Gaussian curvature labels, then a particularly powerful way of representing the compatibility between adjacent labels is to compile a dictionary of legal neighbourhood label configurations. As recently demonstrated by Hancock and Kittler [14,15] this constraint representation can offer significant performance gains when used in conjunction with both discrete and probabilistic relaxation. Our second novel ingredient is therefore to compile a dictionary which represents the valid configurations of topographic labels that can be consistently assigned to neighbouring sites on the surface. In this way we tap the rich source of constraints provided by the highly structured nature of the topographic surface labels. Furthermore these configurations embody smoothness and continuity constraints on the surface label configuration. Since we adopt a triangulated mesh structure to represent the surface, these neighbourhoods consist of triangular configurations of four sub-triangles.

The framework adopted for combining these two sources of information is the dictionary-based relaxation scheme of Hancock and Kittler [14]. This relaxation process is Bayesian and requires the specification of two model components. The first of these are a set of initial label-probabilities that represent the uncertainties in the topographic label assignments. The second is a dictionary of legitimate label configurations. Although we choose to use the dictionary-based probabilistic relaxation process, the model ingredients reported in the paper can be utilised by a wide variety of consistent-labelling schemes.

The outline of this paper is as follows. In Section 2, we review the representation of differential surface structure using the Hessian matrix. Section 3 describes how we estimate the Hessian using surface normals. In Section 4, we detail our modelling of label uncertainty through the propagation of variance. Section 5 outlines the construction of a dictionary of label configurations. In Section 6 we provide some experimental evaluation of our method. Finally, Section 7 offers some conclusions and provides some directions for future research.

2. Representing differential surface structure

In this paper we are interested in estimating the local differential structure of surfaces using computed estimates of the surface normal directions. This is to be contrasted with the fitting of a local surface patch and estimating curvature from the computed parameters of the patch. We commence by providing some of the formal ingredients of our surface representation. Fig. 1 illustrates the basic geometry used in estimating curvature from surface normal information. In particular, the local surface orientation is determined by the direction of the surface normal $\mathbf{n} = (n_x, n_y, 1)^{\mathrm{T}}$. When the surface is represented by a twice differentiable function $z = f(x, y)$, then the components of the normal are related to the surface gradient in the following manner:

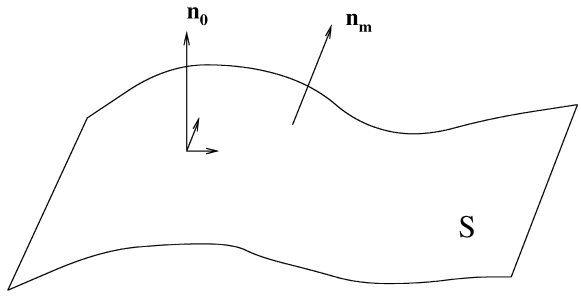

Fig. 1. A local co-ordinate system on the surface.

$$\mathbf{n} = \begin{pmatrix} \dfrac{\partial f}{\partial x} \\ \dfrac{\partial f}{\partial y} \\ 1 \end{pmatrix}. \tag{1}$$

In this continuous case, the differential structure of the surface is captured by the Hessian matrix

$$\mathcal{H} = \begin{pmatrix} \dfrac{\partial^2 f}{\partial x^2} & \dfrac{\partial^2 f}{\partial x \partial y} \\ \dfrac{\partial^2 f}{\partial x \partial y} & \dfrac{\partial^2 f}{\partial y^2} \end{pmatrix} \tag{2}$$

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

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

$$\mathcal{H} = \begin{pmatrix} \alpha & \beta \\ \beta & \gamma \end{pmatrix}. \tag{3}$$

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

$$\alpha = \frac{\partial n_x}{\partial x}, \tag{4}$$

$$\gamma = \frac{\partial n_y}{\partial y}. \tag{5}$$

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

$$\beta = \frac{\partial n_x}{\partial y} = \frac{\partial n_y}{\partial x}. \tag{6}$$

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

With estimates of the elements of the Hessian to-hand, we can compute the mean (H) and Gaussian (K) curvatures of the surface. According to the definitions given above

$$H = \frac{1}{2}(\alpha + \gamma), \tag{7}$$

$$K = \alpha\gamma - \beta^2. \tag{8}$$

The signs and zeros of these two quantities can be used to label the surface according to topographic class. The different classes are defined in Table 1. There are adjacency constraints applying to the curvature labels. In particular, the cup (C) and dome (D) surface types may not appear adjacent to each other on a surface. Moreover, elliptic regions on the surface (those for which K is positive) must be separated from hyperbolic regions (those for which K is negative) by a parabolic line (where $K = 0$). Parabolic lines are effectively zero crossings of the mean or Gaussian curvatures. In other words, domes and cups are enclosed by ridge or valley-lines. Moreover, domes or cups cannot be adjacent to saddle-structures. In Section 5 we will exploit these constraints to construct a dictionary to represent the spatial organisation of the topographic labels.

Table 1
Topographic classes

Class	Symbol	H	K	Region-type
Dome	D	−	+	Elliptic
Ridge	R	−	0	Parabolic
Saddle ridge	SR	−	−	Hyperbolic
Plane	P	0	0	Hyperbolic
Saddle-point	S	0	−	Hyperbolic
Cup	C	+	+	Elliptic
Valley	V	+	0	Parabolic
Saddle-valley	SV	+	−	Hyperbolic

3. Computing the Hessian using sampled normals

In this section we describe how to make a statistical estimate of the Hessian matrix from a sample of surface normals. Specifically, we use the method of least squares to estimate the elements of $\mathscr{H}$ and to compute the errors associated with these estimates.

We commence by assuming that we have a set of surface normal measurements associated with a tentative surface. Moreover, we assume that the variance in the surface normal components is known. For instance, in the case of volumetric intensity images with the surface normals estimated using directional edge-detection operators, then Sharp and Hancock [16,17] have shown how the variance-covariance matrix for the surface normals is determined by the intensity noise-variance together with the autocorrelations of the filter kernels. Let $\mathbf{n}_0$ represent the surface normal at the position (x_0, y_0, z_0) and let $\mathbf{n}_m$ be a neighbouring surface normal with position (x_m, y_m, z_m). If the normals are close to each other, then we can approximate the change in the components of the surface normal using a first-order Taylor expansion. Accordingly,

$$\Delta n_x^{(m)} = \frac{\partial n_x}{\partial x}\Delta x^{(m)} + \frac{\partial n_x}{\partial y}\Delta y^{(m)}, \tag{9}$$

$$\Delta n_y^{(m)} = \frac{\partial n_y}{\partial x}\Delta x^{(m)} + \frac{\partial n_y}{\partial y}\Delta y^{(m)}, \tag{10}$$

where the measured change in the components of the surface normal is given by $\mathbf{n}_m - \mathbf{n}_0 = (\Delta n_x^{(m)}, \Delta n_y^{(m)}, 0)^{\mathrm{T}}$. The displacements in point co-ordinates are $\Delta x^{(m)} = x_m - x_0$ and $\Delta y^{(m)} = y_m - y_0$. Substituting from

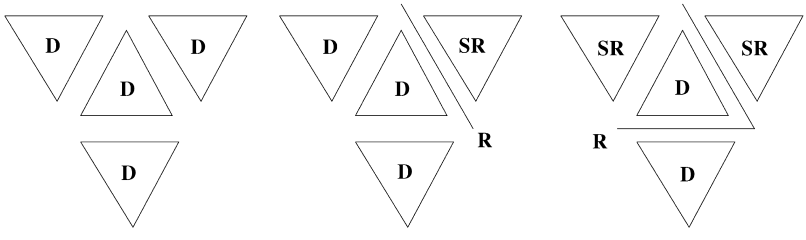

Fig. 2. Elements of the dictionary for the D-label.

Eqs. (4)–(6), we can rewrite the Taylor expansion in terms of elements of the Hessian matrix, i.e.

$$\Delta n_x^{(m)} = \alpha \Delta x^{(m)} + \beta \Delta y^{(m)}, \tag{11}$$

$$\Delta n_y^{(m)} = \beta \Delta x^{(m)} + \gamma \Delta y^{(m)}. \tag{12}$$

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

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

$$\Delta n_x^{(m)} = \Delta x^{(m)} \cdot \alpha + \Delta y^{(m)} \cdot \beta + 0 \cdot \gamma,$$

$$\Delta n_y^{(m)} = 0 \cdot \alpha + \Delta x^{(m)} \cdot \beta + \Delta y^{(m)} \cdot \gamma. \tag{13}$$

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

$$\mathbf{N} = \mathbf{XP} \tag{14}$$

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

$$\mathbf{N} = \begin{pmatrix} \Delta n_x^{(1)} \\ \Delta n_y^{(1)} \\ \Delta n_x^{(2)} \\ \vdots \end{pmatrix},$$

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

$$\mathbf{X} = \begin{pmatrix} \Delta x^{(1)} & \Delta y^{(1)} & 0 \\ 0 & \Delta x^{(1)} & \Delta y^{(1)} \\ \Delta x^{(2)} & \Delta y^{(2)} & 0 \\ & \vdots & \end{pmatrix}$$

and $\mathbf{P}$ is the parameter vector

$$\mathbf{P} = \begin{pmatrix} \alpha \\ \beta \\ \gamma \end{pmatrix}.$$

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

$$\hat{\mathbf{P}} = \arg\min_{\mathbf{P}} (\mathbf{N} - \mathbf{XP})^{\mathrm{T}} (\mathbf{N} - \mathbf{XP}). \tag{15}$$

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

$$\hat{\mathbf{P}} = (\mathbf{X}^{\mathrm{T}}\mathbf{X})^{-1}\mathbf{X}^{\mathrm{T}}\mathbf{N}. \tag{16}$$

The resulting least-squares estimates of the elements of the Hessian can be used to compute the mean and Gaussian curvatures using Eqs. (7) and (8).

4. Covariance propagation

Rather than assigning the topographic labels defined in Table 1 by thresholding the estimated values of the mean and Gaussian curvature, we adopt a relaxation method. Specifically, we aim to iteratively refine a set of initial label probabilities using dictionary-based probabilistic relaxation. In order to embark on this endeavor, we will require a probabilistic characterisation of the uncertainties in the mean and Gaussian curvatures. We use the variance-covariance matrix for the estimated values of H and K to provide the necessary measure of uncertainty. This matrix is computed by propagating the variance from the surface-normal components. In our experiments we will use surfaces segmented from volumetric intensity images to evaluate our labelling process. We commence from the assumption that the components of the surface normal are un-correlated and have the same variance v. The changes $\Delta n_x, \Delta n_y$ are calculated from the difference of a normal to some central reference normal thus

$$\Delta n_x^{(m)} = n_x^{(m)} - n_x^{(0)},$$

$$\Delta n_y^{(m)} = n_y^{(m)} - n_y^{(0)}.$$

In consequence, the covariance matrix for the normal changes is given by

$$\Sigma_N = v \begin{pmatrix} 1 & 0 & 2 & 0 & 2 & \cdots \\ 0 & 1 & 0 & 2 & 0 & \cdots \\ 2 & 0 & 1 & 0 & 2 & \cdots \\ & & \vdots & & & \end{pmatrix}. \tag{17}$$

Since the parameter-vector $\hat{\mathbf{P}}$ is related to the computed surface normals in a linear fashion according to Eq. (16), the parameter covariance matrix is given by

$$\mathbf{\Sigma}_p = (\mathbf{X}^{\mathrm{T}}\mathbf{X})^{-1}\mathbf{X}^{\mathrm{T}}\mathbf{\Sigma}_{\mathbf{N}}\mathbf{X}(\mathbf{X}^{\mathrm{T}}\mathbf{X})^{-1}. \tag{18}$$

For the sake of convenience, we denote the elements of $\mathbf{\Sigma}_p$ as follows:

$$\mathbf{\Sigma}_p = \begin{pmatrix} v_{\alpha\alpha} & v_{\alpha\beta} & v_{\alpha\gamma} \\ v_{\beta\alpha} & v_{\beta\beta} & v_{\beta\gamma} \\ v_{\gamma\alpha} & v_{\gamma\beta} & v_{\gamma\gamma} \end{pmatrix}. \tag{19}$$

Since the mean and Gaussian curvatures are computed from the trace and determinant of the Hessian matrix, we can proceed no further with the linear propagation of variance. However, if we perform a first-order error analysis, then the variances and covariance for the mean and Gaussian curvatures are approximately given by

$$v_{HH} = v_{\alpha\alpha} + v_{\gamma\gamma} + 2v_{\alpha\gamma}, \tag{20}$$

$$v_{KK} = \gamma^2 v_{\alpha\alpha} + \alpha^2 v_{\gamma\gamma} + 4\beta^2 v_{\beta\beta} + 2\alpha\gamma v_{\alpha\gamma} - 4\gamma\beta v_{\alpha\beta} - 4\alpha\beta v_{\beta\gamma}, \tag{21}$$

$$v_{KH} = \gamma v_{\alpha\alpha} + \alpha v_{\gamma\gamma} + (\alpha + \gamma) v_{\alpha\gamma} - 2\beta v_{\alpha\beta} - 2\beta v_{\beta\gamma}. \tag{22}$$

Specifically, the $H - K$ covariance matrix is

$$\mathbf{\Sigma}_l = \begin{pmatrix} v_{HH} & v_{KH} \\ v_{KH} & v_{KK} \end{pmatrix}. \tag{23}$$

4.1. Label probabilities

Our motivation for performing the analysis of errors on the values of H and K was to provide probability estimates for the different curvature label types at a point on the surface. This is achieved by adopting a Gaussian distribution for the mean and Gaussian curvatures. The basic model assumes that the observed values of the curvatures h_m and k_m deviate from their true values under Gaussian measurement errors. The distribution of errors is controlled by the $H - K$ covariance matrix Σ_l. If $\mathbf{x} = (h - h_m, k - k_m)^{\mathrm{T}}$ is the difference between the true measurement vector $(h, k)^{\mathrm{T}}$ and the observed measurement vector $(h_m, k_m)^{\mathrm{T}}$, then the assumed distribution is as follows

$$p(H = h_m, K = k_m) = \frac{1}{2\pi\sqrt{|\mathbf{\Sigma}_l|}}\, \mathrm{e}^{-\frac{1}{2}\mathbf{x}^{\mathrm{T}}\mathbf{\Sigma}_1^{-1}\mathbf{x}}. \tag{24}$$

We now provide an illustration of how this probability distribution can be used to assign curvature label probabilities. We take as our example the saddle class S which requires $H = 0$ and $K < 0$. The probability of the saddle label is taken to be equal to the cumulative probability distribution over the range of H and K for which the saddle condition holds. In other words,

$$P^{(0)}(S) = \int_{-\infty}^{0} p(H = h_m, K = k)\, \mathrm{d}k. \tag{25}$$

To proceed, suppose that the components of the inverse covariance matrix are A_{hh}, A_{hk} and A_{kk}, i.e.

$$\mathbf{\Sigma}_l^{-1} = \begin{pmatrix} A_{hh} & A_{hk} \\ A_{hk} & A_{kk} \end{pmatrix}. \tag{26}$$

On substituting for the inverse covariance matrix in the probability distribution of Eq. (24) and performing the integration in Eq. (25), the initial probability for the saddle-class is given by

$$P^{(0)}(S) = \frac{\sqrt{A_{hh}}}{|\Sigma_l|^{1/2}}\, \mathrm{e}^{-\frac{k^2}{2}(A_{hk}^2 - A_{hh}A_{kk})} \times \frac{1}{2}\left\{1 + \mathrm{erf}\left[\sqrt{\frac{2}{A_{hh}}}\left(h_m + \frac{A_{hk}}{A_{hh}}\, k_m\right)\right]\right\}. \tag{27}$$

The remaining initial label probabilities are obtained by computing the relevant cumulative distribution function over the range of H and K values for which the relevant topographic label applies (see Table 1).

5. Labelling the surface

Equipped with the values of K and H, the local surface structure can be characterised using the standard eight topographic label classes from Table 1. However, because of uncertainties in the surface normals, these labels are likely to be noisy and subject to classification error. In particular neighbouring surface labels may be inconsistent with one-another when viewed from the perspective of legal differential structure. In other words, they may violate the constraint that elliptical and hyperbolic regions must be simply connected, and be separated from one-another by thin parabolic lines. Furthermore, regions of the same label type may be fragmented by noise. It should be noted that these observations only apply to smoothly and continuously curving surfaces.

In order to quantify these observations, we construct an exhaustive dictionary of allowed label configurations which satisfy these constraints. It should be noted that if we wish to construct a dictionary which effectively smoothes noisy surface labels to create contiguous regions, we must guarantee that the scale over which the smoothing takes place is appropriate. In other words, if we construct a dictionary which allows only smooth label configurations, we must ensure that the surface varies smoothly over the scale of a dictionary item. It is for this reason that we adopt the surface mesh representation of Wilson and Hancock [18] in which the density of surface elements is proportional to the surface curvature and the surface is smooth over adjacent elements of the mesh.

5.1. Dictionary

As mentioned above, there are strong adjacency constraints on the valid label configurations appearing on the neighbourhoods of some classes of surface. One of the goals of the work reported in this paper is to describe a methodology for enumerating and encoding these constraints in a dictionary for the topographic labels. Of course, the construction of a dictionary depends critically on the choice of neighbourhood topology. In the experimental section later we will be demonstrating the utility of the method on triangular surface meshes. In this case the natural neighbourhood consists of a triangle and it's three directly adjoining elements. We will therefore confine our attention exclusively to this arrangement of objects in constructing the dictionary.

We commence by considering the dome-class (D) for which $K > 0$ and $H < 0$. This is an example of an elliptic region. It can therefore be connected to other dome labels. In fact if smoothness and continuity considerations are to be satisfied, the label must be connected to one or more other dome labels. In other words, smoothness considerations rule out the case in which a dome label is surrounded by non-dome labels. It cannot be adjacent to any class other than the saddle-ridge for which $K < 0$ and $H < 0$, and these two labels are separated by a ridge line. The ridge and valley labels fall into the category of parabolic lines. In other words, they must form the boundaries between hyperbolic regions and elliptical regions. Specifically, they are effectively zero crossings of Gaussian curvature. In consequence the ridge label intercedes between elliptical domes and hyperbolic saddle-ridges, and resides on the mesh edges separating these two labels. The configurations satisfying these two constraints are shown in Fig. 2. It should be noted that all sections of the dictionary are rotation invariant. The cup-class (C) is symmetric with the dome under reversal of the sign of the mean curvature (H). Under this transformation, the ridge-label is replaced with the valley label. The dictionary for the cup class is therefore constructed by performing the mappings $D \rightarrow C$, $R \rightarrow V$ and $SR \rightarrow SV$ in Fig. 2.

The two hyperbolic region labels have a more complicated neighbourhood structure. The saddle-valley and the saddle-ridge labels must again form contiguous patches. However, they can be bounded by both the converse saddle surface and the dome or cup of appropriate mean-curvature. For instance the saddle-ridge can be adjacent to the dome and the saddle-valley. Again smoothness considerations rule out an unconnected central label. The set of legal dictionary-items for the saddle ridge which satisfy these three constraints are shown in Fig. 3 (excluding the rotational symmetries). The saddle-valley is again symmetric under reversal of the sign of the mean-curvature.

5.2. Probabilistic relaxation

In order exploit the highly structured nature of the surface labelling constraints, we have chosen to employ the technique of probabilistic relaxation. However, it should be pointed out that the model ingredients developed in this paper could equally be applied to other consistent labelling schemes. Our reasons for choosing probabilistic relaxation are two-fold. In the first instance, the method draws on labelling constraints using an exhaustive list or compilation of valid neighbourhood label configurations. In our surface labelling application the dictionary is simply that described in the previous section. The second reason is that the framework is Bayesian and combines evidence for label assignments. Rather than commencing from hard and potentially erroneous

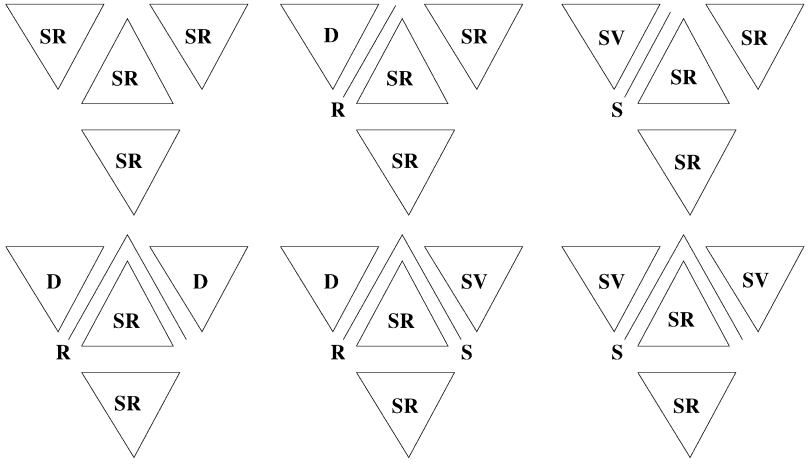

Fig. 3. Elements of the dictionary for the SR-label.

label assignments, the initial characterisation is in terms of *a posteriori* label probabilities. In other words, dictionary-based relaxation allows us to exploit and combine both consistent surface label structure and the covariance structure of the H-K curvatures.

According to the original formulation of Hancock and Kittler [14], the local labelling is described by a set of probabilities. Specifically, $P^{(n)}(S_i = \omega)$ is the weight of evidence assigned to the label assignment ω at site S_i for the iterative epoch n of the algorithm. Initially, these probabilities are calculated using the computed values of H and K, together with their known covariance structure. The label-probabilities are iteratively updated using the following non-linear relaxation rule:

$$P^{(n+1)}(S_i = \omega) = \frac{P^{(n)}(S_i = \omega)Q^{(n)}(S_i = \omega)}{\sum_{\omega' \in \Omega} P^{(n)}(S_i = \omega')Q^{(n)}(S_i = \omega')}. \quad (28)$$

The critical ingredient in the update formula is the support function $Q^{(n)}(S_i = \omega)$ which combines evidence from the context-conveying neighbourhood K_i of the surface-site S_i for the label assignment $\omega \in \Omega$. Here $\Omega = \{D, R, SR, S, SV, V, C\}$ denotes the complete set of curvature labels. According to Hancock and Kittler [14], the support function takes on the following product-form:

$$Q^{(n)}(S_i = \omega) = \sum_{\Lambda \in \Theta(\omega)} \left\{ \prod_{k \in K_i} \frac{P^{(n)}(S_k = \lambda_k)}{P(S_k = \lambda_k)} \right\} \times P(S_l = \lambda_l \ \forall l \in K_i). \quad (29)$$

In the above formula $\Theta(\omega)$ denotes the set of label configurations from the dictionary that assign the label ω to the centre-object, i.e. S_i. The complete dictionary is the set of configurations

$$\Theta = \bigcup_{\omega \in \Omega} \Theta(\omega). \quad (30)$$

We have used K_i to denote the index-set of the objects that form the local neighbourhood of the site S_i. With this notation, the dictionary item $\Lambda = \{\lambda_k;\ k \in K_i\}$ is a configuration of valid topographic labels on the neighbourhood K_i. In the work reported here, the dictionary consists of the set of curvature label configurations described in Section 5.1. The structure of the label model is represented by the prior probabilities, $P(S_l = \lambda_l \ \forall l \in K_i)$ for the legal label configurations belonging to the dictionary. Here we assume that the dictionary items occur with uniform a priori probabilities, i.e.

$$P(S_l = \lambda_l \ \forall l \in K_i) = \frac{1}{\sum_{\omega' \in \Omega} |\Theta(\omega)|}. \quad (31)$$

Finally, $P(S_i = \omega)$ is the single-object prior for assigning the label ω to the site S_i. Here we model the priors by counting the number of times the dictionary makes the assignment of the label ω to the centre-object. In other words,

$$P(S_i = \omega) = \frac{|\Theta(\omega)|}{\sum_{\omega' \in \Omega} |\Theta(\omega)|}. \quad (32)$$

With these ingredients in place the relaxation scheme is iterated to locate the final consistent labelling of the surface.

6. Experimental evaluation

In this section we offer some experimental validation of our surface-labelling algorithm. There are two aspects to this study. In the first instance, we provide some real-world evaluation of the method on surface-data extracted from MRI data. In order to evaluate the method under controlled conditions, the second set of experiments are conducted using synthetic surfaces. The aim here is to compare the labellings obtained with ground-truth data. In particular, we illustrate some of the iterative properties of curvature labelling process. We also measure the sensitivity of the surface labels to controlled levels of added noise.

6.1. Real world data

The real world application of our curvature labelling technique concerns surface-data extracted from a volumetric MRI image of a human head. The processing chain used to extract the surface data is as follows. We commence by extracting connected contours on the individual slices in the MRI volume by applying an edge following snake to the raw grey-scale values. The snake attaches itself to boundaries closest to the perimeter of the slices. As a result, the recovered 3D data-points lie on the boundary surface of the head. Once the surface points are to hand, we extract a triangulated surface mesh. The mesh used in our experiments is described in recent papers by Wilson and Hancock [18,19]. This mesh is fitted to the raw surface data-points using a series of split and merge operations aimed at optimising a local variance-bias criterion. The neighbourhoods used in our curvature labelling experiments consist of a centre-triangle together with the three triangles that share an edge.

The data used in our study is a frontal view of the head which contains fine surface detail around the nose and eye-sockets. Fig. 4 illustrates the surface-segmentation process. Initially, the surface-mesh is relatively noisy. The triangulation provides only a rough initial surface representation. The mesh adapts to this surface by decreasing the density on the relatively smooth forehead, and increasing the surface density around the highly curved nose and eye-sockets (see Fig. 4). Noisy surface structure

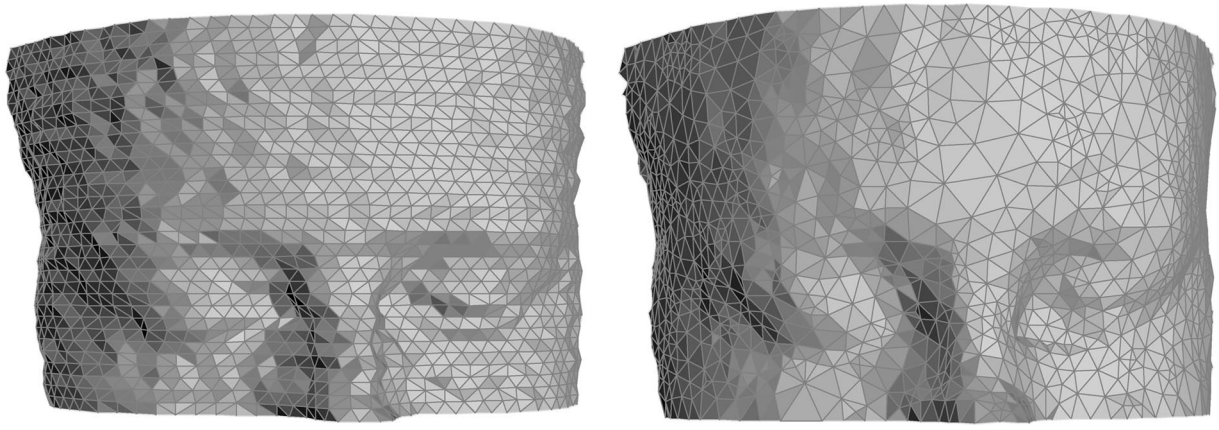

Fig. 4. Initial and final configuration of the mesh over a surface extracted from 3D slice data.

is smoothed away. There is no evidence for over-fitting with a spuriously high mesh-density in the proximity of surface-noise.

We illustrate the labelling of the segmented surface in Fig. 5. The figure shows an iterative sequence of winner-take-all labellings derived from the probabilities delivered by the relaxation process. Initially, the labelling contains many local inconsistencies. These are most marked in the highly curved regions around the eye sockets and on the ridge of the nose. After application of the relaxation process, most of the local inconsistencies are corrected. The hyperbolic and elliptical regions display impressive connectivity. The algorithm reaches stable convergence after approximately 20 iterations.

6.2. Simulated surface data

The real-world experiments described in the previous subsection reveal in a qualitative way that our labelling scheme has some practical utility. However, since ground-truth is not available it does not allow sensitivity limits to be established. In order to furnish operating limits for our curvature labelling process, we provide some experiments on simulation data with known ground-truth.

We commence by generating surfaces with controlled levels of additive Gaussian noise. The surfaces simulate range images. The additive Gaussian noise models sensing errors in the surface height distribution. The raw height data is triangulated using sample points from the surface. The mesh optimisation algorithm of Wilson and Hancock [18] is then applied to produce a surface triangulation. The points within these triangles then provide the surface normal information required for the initial surface-curvature estimation.

We begin by demonstrating the ability of the algorithm to recover from noise using a simple spherically curved surface of radius 200 units and with additive Gaussian noise of variance 1 unit in the z-direction. Fig. 6 shows the initial curvature labelling on the left and the final labelling after application of the algorithm on the right. The initial labelling is corrupted with large numbers of spurious labels and inconsistencies between adjacent labels. After relabelling, the labels on the cup surface are completely consistent; initial some 61% of labels are correct and after application of the algorithm all labels are correct.

Our next set of experiments use a more complex surface which contains all the different label classes. Fig. 7 shows the initial labelling of a cosine surface with additive Gaussian noise of variance 1 unit. The surface labels are colour coded according to the following convention. Domes appear as red, saddle ridges are yellow, saddle valleys are mauve and cups are blue. Notice that in the initial labelling the abutting boundaries of the different regions are intrinsically inconsistent when viewed from the perspective of the topographic label-set. Fig. 8 shows the labelling of the surface after 12 iterations of the dictionary-based relaxation scheme. The green edges appearing in the updated labelling are parabolic lines (i.e. either ridges, valleys or saddles), which are consistently labelled between the appropriate regions. The overall consistency of the labels from the standpoint of smoothness constraints is considerably improved. Initially some 61% of labels are correct; after application of the scheme, this improves to 81%. Study of the ground truth labelling (Fig. 9) reveals that most of these errors occur in the saddle label classes. Although these saddle regions appear in qualitatively correct positions relative to the elliptic classes, mis-location of the region boundaries results in incorrect labels relative to ground truth.

We have performed a more systematic study of the labelling scheme under controlled noise conditions. In

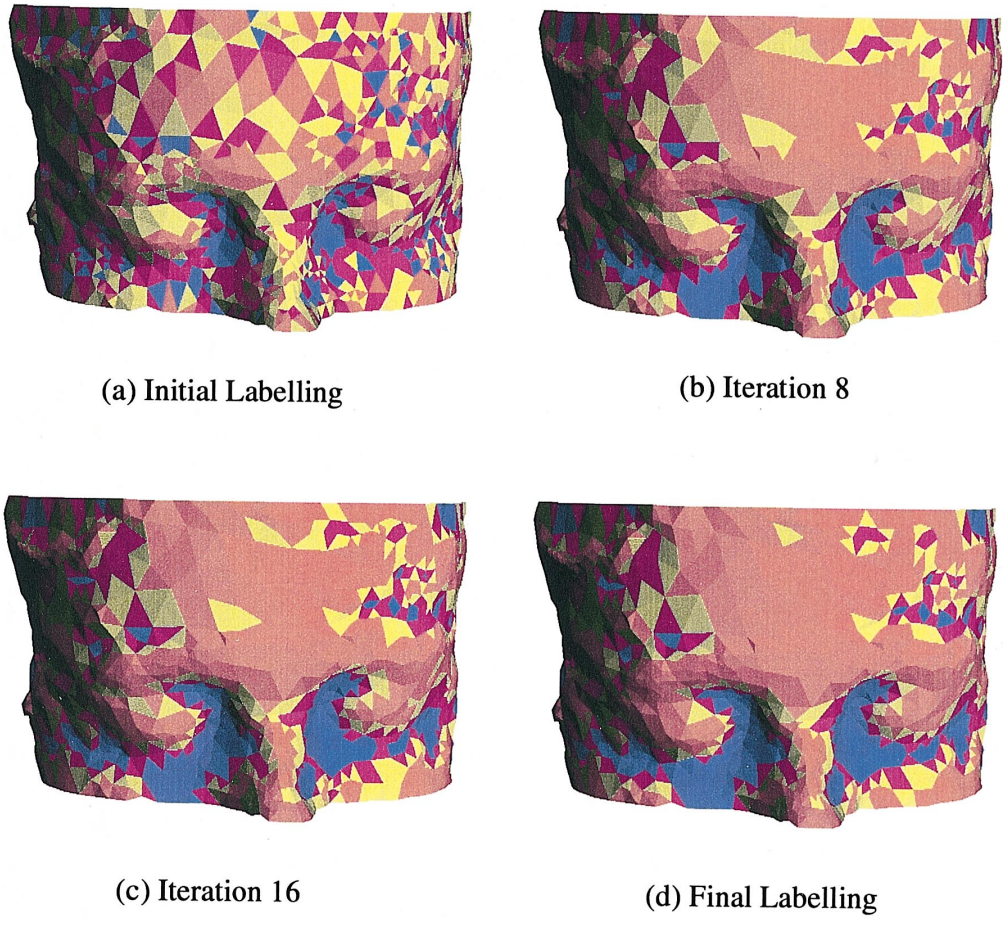

Fig. 5. Iterative sequence showing successive topographic labels for 3D slice data; Red labels represent domes, blue labels are cups, yellow labels are saddle ridges and purple labels are saddle valleys. (a) Initial labelling, (b) iteration 8, (c) iteration 16, (d) final labelling.

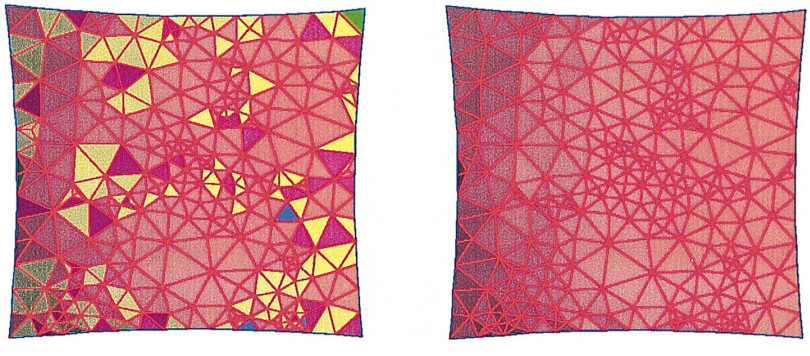

Fig. 6. Initial and final labelling of a noisy spherical surface.

Figs. 10 and 11 we investigate the labelling performance as a function of noise-variance in the raw data-point heights. The two plots respectively show the results for the spherical and cosine surfaces. These figures underline in a more quantitative way some of the qualitative observations made earlier. The two curves in the plots show the initial and final fractions of correct labels when a winner-take-all decision is made. In both cases the final

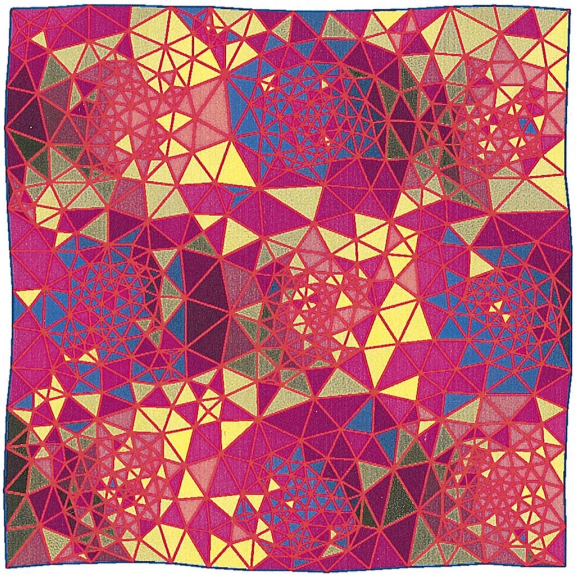

Fig. 7. Initial labelling of cosine surface.

Fig. 8. Final labelling of cosine surface.

labelling represents a considerable improvement in accuracy over the initial non-contextual assignment of topographic labels. It is interesting to note that this is the case for both the spherical surface, which is of uniform topographic composition, and the cosine surface, which is more structured. However, in the case of the cosine surface, the improvement over the initial labelling is not as marked for large values of the height noise-variance.

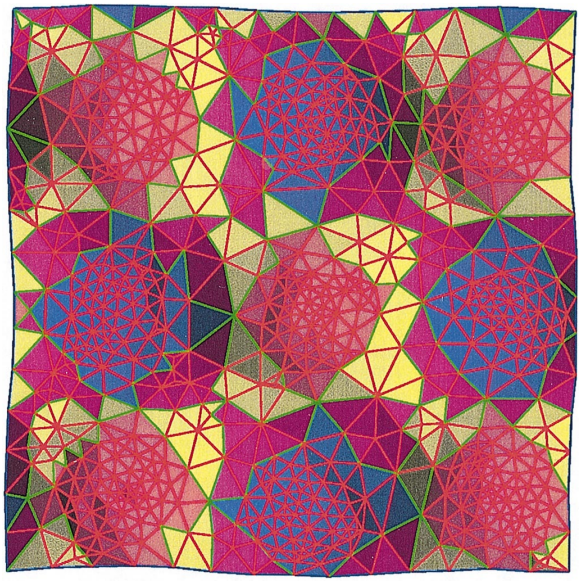

Fig. 9. True labelling of cosine surface.

7. Conclusions

Our main contributions in this paper are twofold. In the first instance, we have demonstrate how topographic surface labelling can be realised using a constrained dictionary of feasible surface-label configurations. These configurations observe certain constraints on the adjacency of different topographic classes together with their region continuity or line contiguity. The second contribution is to develop a statistical model which allows the topographic surface-labels to be assigned probabilities based on the measured values of the mean and Gaussian curvatures. These probabilities model the assignment of topographic classes using surface normal information. We have also shown how these two elements can be integrated using a relaxation scheme. Results have shown that the resulting labelling is a considerable improvement over the raw topographic labels.

There are a number of ways in which the ideas developed in this paper can be extended. In the first instance, we intend to construct a more realistic model of the distribution of curvature error. At the moment there is little justification for our Gaussian distribution model. Secondly, we intend to explore alternative curvature representations. These include the perceptually motivated angular shape-index of Koenderinck [20]. This will have a number of implications for our surface model. In the first-instance, the dictionary will have a more natural structure. Studies directed at these issues are in progress and will be reported in due course.

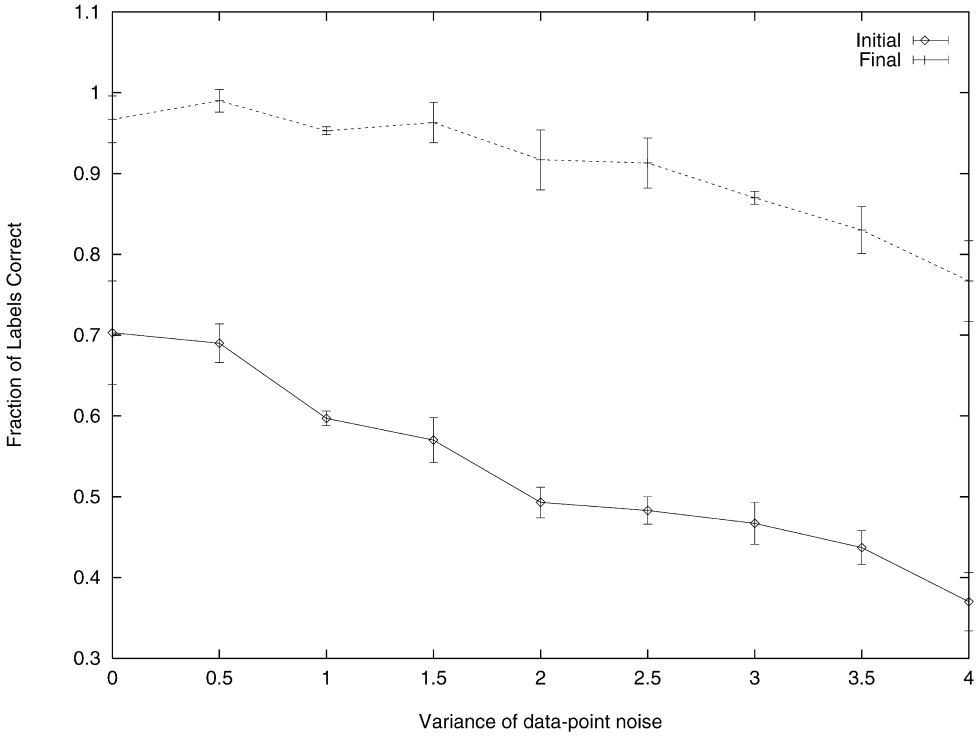

Fig. 10. Graph of the performance of the surface-labelling scheme on a noisy sphere surface.

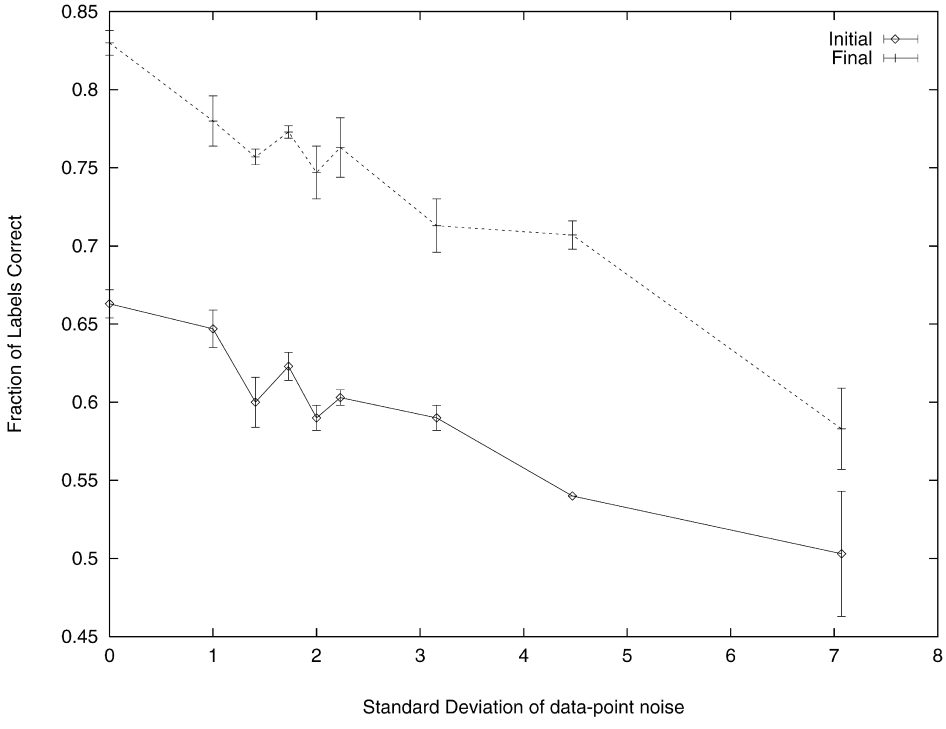

Fig. 11. Graph of the performance of the surface-labelling scheme on a noisy cosine surface.

References

[1] P.J. Besl, R.C. Jain, Segmentation through variable order surface fitting, IEEE Trans. Pattern Anal. Machine Intell. PAMI 10 (1988) 167–192.

[2] Y.S. Lim, T.I. Cho, K.H. Park, Range image segmentation based on 2D quadratic function approximation, Pattern Recognition Lett. 11 (1990) 699–708.

[3] E. Trucco, R.B. Fisher, Experiments in curvature-based segmentation of range-data, IEEE Trans. Pattern Anal. Machine Intell. PAMI 17 (1995) 177–182.

[4] P.T. Sander, S.W. Zucker, Inferring surface structure and differential structure from 3D images, IEEE Trans. Pattern Anal. Machine Intell. PAMI 12 (1990) 833–854.

[5] P.T. Sander, S. Zucker, Singularities in the principal direction field from 3D images, IEEE Trans. Pattern Anal. Machine Intell. PAMI 14 (1992) 309–317.

[6] O. Monga, N. Ayache, P. Sander, Using uncertainty to link 3D edge detection and local surface modelling, Lecture Notes in Computer Science, vol. 511, Springer, Berlin, 1991, pp. 273–284.

[7] O. Monga, R. Deriche, G. Malandain, Recursive filtering and edge closing; two primary tools for 3D edge detection, Image Vision Comput. 9 (1991) 403–417.

[8] O. Monga, S. Benayoun, Using differential geometry in R4 to extract typical features in 3D density images, 11th Internat. Conf. on Pattern Recognition, Vol. 1, 1992, pp. 379–382.

[9] P.J. Flynn, A.K. Jain, On reliable curvature estimation, IEEE Computer Vision and Pattern Recognition Conf., 1989, pp. 110–116.

[10] R.M. Haralick, L. Watson, A facet model for image data, Comput. Graph. Image Process. 15 (1981) 115–129.

[11] R.M. Bolle, D.B. Cooper, Bayesian recognition of of local 3D shape by approximating image intensity functions with quadric polynomials, IEEE Trans. Pattern Anal. Machine Intell. PAMI 6 (1984) 418–429.

[12] N.N. Abdelmalek, Algebraic error analysis for surface curvatures and segmentation of 3-D range images, Pattern Recognition 23 (8) (1990) 807–817.

[13] A. Hilton, J. Illingworth, T. Windeatt, Statistics of surface curvature estimates, Pattern Recognition 28 (1995) 1201–1221.

[14] E.R. Hancock, J. Kittler, Edge labelling using dictionary based relaxation, IEEE Trans. Pattern Anal. Machine Intell. PAMI 12 (1990)165–181.

[15] E.R. Hancock, J. Kittler, Discrete relaxation, Pattern Recognition 23 (1991) 711–733.

[16] N.G. Sharp, E.R. Hancock, Density propagation for surface tracking, Pattern Recognition Letters 19 (1998) 177–188.

[17] N.G. Sharp, E.R. Hancock, Feature tracking by multi-frame relaxation, Image Vision Comput. 13 (1995) 637–644.

[18] R.C. Wilson, E.R. Hancock, A minimum variance adaptive surface mesh, IEEE Computer Vision and Pattern Recognition Conf., 1997, pp. 634–639.

[19] R.C. Wilson, E.R. Hancock, Variance-bias tradeoff for adaptive surface meshes, Computer Vision ECCV 98, Lecture Notes in Computer Science, Vol. 1407, Springer, Berlin, 1998, pp. 449–465.

[20] J.J. Koenderinck, A. Van Doorn, Surface shape and curvature scales, Image Vision Comput. 10 (1992) 557–564.

[21] F.P. Ferrie, J. Lagarde, P. Whaite, Darboux frames, snakes, and super-quadratics: geometry from the bottom up, IEEE Trans. Pattern Anal. Machine Intell. PAMI 15 (1993) 771–784.

[22] E.M. Stokely, Surface parameterisation and curvature measurement of arbitrary 3-D objects – 5 practical methods, IEEE Trans. Pattern Anal. Machine Intell. PAMI 14 (1992) 833–840.

About the Author—DR WILSON was awarded an open scholarship to read physics at St John's College, University of Oxford, graduating with first class honours in 1992. Between 1992 and 1995 he undertook research at the University of York on the topic of relational graph matching for which he was awarded the DPhil degree. He is currently an EPSRC Advanced Research Fellow in the Department of Computer Science at the University of York. He has published some 45 papers in journals, edited books and refereed conferences. He recieved an honorable mention in the 1997 Pattern Recognition Society awards. His research interests are in high-level vision, scene understanding, volumetric image analysis and structural pattern recognition.

About the Author—PROFESSOR EDWIN HANCOCK gained his B.Sc. in physics in 1977 and Ph.D. in high energy nuclear physics in 1981, both from the University of Durham, UK. After a period of postdoctoral research working on charm-photo-production experiments at the Stanford Linear Accelerator Centre, he moved into the fields of computer vision and pattern recognition in 1985. Between 1981 and 1991, he held posts at the Rutherford-Appleton Laboratory, the Open University and the University of Surrey. Professor Hancock is currently Professor of Computer Vision at the University of York. He leads a group of some 15 researchers in the areas of computer vision and pattern recognition. He has published about 180 refereed papers in the fields of high energy nuclear physics, computer vision, image processing and pattern recognition. He was awarded the 1990 Pattern Recognition Society Medal and received an honorable mention in 1997. Professor Hancock serves as an Associate Editor of the journal Pattern Recognition and has been a guest editor for the Image and Vision Computing Journal. He is currently guest-editing a special edition of the Pattern Recognition journal devoted to energy minimisation methods in computer vision and pattern recognition. He chaired the 1994 British Machine Vision Conference and has been a programme committee member for several national and international conferences.