



PERGAMON

Pattern Recognition 34 (2001) 903–922

**PATTERN
RECOGNITION**

THE JOURNAL OF THE PATTERN RECOGNITION SOCIETY

www.elsevier.com/locate/patcog

A Bayesian framework for 3D surface estimation

Mick Turner, Edwin R. Hancock*

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

Received 30 August 1999; received in revised form 2 January 2000; accepted 2 January 2000

Abstract

We develop an evidence-combining framework for extracting locally consistent differential structure from curved surfaces. Existing approaches are restricted by their sequential multi-stage philosophy, since important information concerning the salient features of surfaces may be discarded as necessarily condensed information is passed from stage to stage. Furthermore, since data representations are invariably unaccompanied by any index of evidential significance, the scope for subsequently refining them is limited. One way of attaching evidential support is to propagate covariances through the processing chain. However, severe problems arise in the presence of data non-linearities, such as outliers or discontinuities. If linear processing techniques are employed covariances may be readily computed, but will be unreliable. On the other hand, if more powerful non-linear processing techniques are applied, there are severe technical problems in computing the covariances themselves. We sidestep this dilemma by decoupling the identification of non-linearities in the data from the fitting process itself. If outliers and discontinuities are accurately identified and excluded, then simple, linear processing techniques are effective for the fit, and reliable covariance estimates can be readily obtained. Furthermore, decoupling permits non-linearity estimation to be cast within a powerful evidence combining framework in which both surface parameters and refined differential structure come to bear simultaneously. This effectively abandons the multi-stage processing philosophy. Our investigation is firmly grounded as a global MAP estimate within a Bayesian framework. Our ideas are applicable to volumetric data. For simplicity, we choose to demonstrate their effectiveness on range data in this paper. © 2001 Pattern Recognition Society. Published by Elsevier Science Ltd. All rights reserved.

Keywords: Surface reconstruction curvature estimation; Bayes theorem; MAP estimation; Relaxation labelling

1. Introduction

Since the advent of sensor technology capable of volumetric imaging, three-dimensional scene analysis has become a critical area of investigation in computer vision. The main types of data under study are range images, which consist of an array of sensed depth values, and density slice data generated by techniques such as magnetic resonance imaging, or X-Ray computed tomography. Conventionally, researchers have adopted a four-stage strategy to scene interpretation [1–6]. Firstly, surface points are estimated. Whilst unoccluded sur-

faces are detected directly in range imagery, in the case of slice data intensity features must be tracked over a number of adjacent slices in order to establish their persistence [7,8]. In practice, the features are usually 3D edges which present themselves as intensity discontinuities [3]. As with 2D edges, hysteresis linking is applied to improve the connectivity of the detected features. For example, Monga et al. [1] have developed an extension of the Canny–Deriche edge detector [9,10] to 3D slice data processing. Once feature detection has been performed the next stage involves fitting a parametric model of local surface to the putative surface points. Commonly, a conventional least-means-squares approach is adopted [6,11], although Monga et al. [1] use a Kalman filter in a weighted least-squares approach. In the third-stage local surface estimates obtained by the fitting process are defined to improve their consistency. Sander and Zucker have recently introduced a relaxation process designed to

* Corresponding author. Tel.: + 44-1904-43-3374; fax: + 44-1904-43-2767.

E-mail address: erh@cs.york.ac.uk (E.R. Hancock).

improve the consistency of differential structure of neighbouring points [6]. In practice, the algorithm employed updates charts via a least-means-squares estimate. Finally, with refined differential structure to hand, identification of global surfaces based upon curvature information can proceed [6,11–13].

It is this sequential multi-stage philosophy, and its band limitation of information, that restricts the effectiveness of most existing approaches to 3D scene interpretation. Viewed as a system of finite information bandwidth, as necessarily condensed data representations are passed from stage to stage, important information concerning the salient features of surfaces may be discarded. This is particularly critical for surfaces containing fine curvature detail. Moreover, since the data representations are effectively static, and invariably lack any measure of evidential significance, the scope for subsequently adapting them is limited. As a concrete example, and as we will demonstrate in this paper, given a poor initial segmentation of volumetric feature points it may prove highly advantageous to disregard outliers or surface discontinuities in the estimation of differential structure. Critical to this approach is a means of assessing the evidential impact, or lack thereof, of each feature point in the overall assessment of a consistent curvature description.

One way to address this need and to establish a metric of statistical significance at each stage is to propagate covariance through the processing chain [1,4,14]. Whilst this may be a correct procedure, there are severe difficulties associated with its practical application. Whilst covariance computation and propagation is indeed feasible when the processing stages involve linear operation, in many practical situations the covariance estimates are unreliable due to the inherent non-linearity of the information being processed. On the other hand, if more powerful non-linear processing techniques are employed there are severe technical problems in the computation of the covariances themselves. Conventional attempts at resolving this dilemma must choose between linear or non-linear algorithms, based upon some compromise position between algorithm effectiveness and the reliability of the computed covariances.

In contrast, we effectively sidestep this dilemma by estimating and excluding non-linearities from the fit processes. While this approach is identical to the stated aim of robust statistics [15–18], and relates to regularisation technique [19], there is a key aspect to our approach which leads to major advantages over other methods. Specifically, we decouple non-linearity estimation from the fitting process. Then, provided non-linearities can be reliably identified and excluded, linear processing techniques will be appropriate and effective for the fit, and reliable covariance estimates can be readily computed and propagated. This contrasts with regularisation and robust fitting techniques in which non-linearity estima-

tion and data fitting are performed simultaneously. As a consequence, non-linear algorithms are required, and computation of covariance is often intractable. Furthermore, our decoupling facilitates more accurate identification of non-linearities in the data since evidence from multiple sources can be utilised, without worrying about introducing additional complexity into the fitting algorithm. Robust and regularisation methods are generally restricted to using fit-residuals alone, since incorporation of higher-order features soon makes analysis unmanageable. However, higher-order features are necessary to properly identify non-linearities [20].

In the work reported here we apply our ideas to surface parameter estimation and surface refinement, both of which may be considered as fitting problems. In relation to surfaces, non-linearities may arise from gross data outliers, or points which lie across genuine surface boundaries or discontinuities. Our goal is to identify such points (which we term irregularities) so that they may be excluded prior to the fit, thereby facilitating the effective use of linear processing operations. Of course, before reliable estimates of irregularities can be made, a reasonable estimate of the underlying surface is required. Evidently, such information may not be available, since it is the underlying surface that we are trying to determine in the first place. We solve this dilemma by basing estimates of irregularity upon the most recently available surface estimates. Since surface estimates will themselves be influenced by which data is excluded from the fits, the implication is that surface estimation should be a recursive process which is iterated until convergence. The idea is that improved estimates of irregularities will in turn lead to better surface estimates, and vice versa. Furthermore, by casting the estimation of irregularities in terms of a unified evidence combining framework in which both surface parameters and refined differential structure information come to bear simultaneously, we effectively abandon the conventional sequential, multi-stage philosophy. Specifically, information from one level can be utilised to improve estimates at the other level — i.e., the parameter estimation process and the refinement process become coupled.

We formulate the recovery of consistent differential structure as a global MAP estimate within a Bayesian framework. The assumption of uniform measurement probability density functions for irregularities leads to a weighted least-mean-squares technique for surface parameter estimation and surface refinement, respectively. These linear algorithms facilitate the characterisation of surface parameters and charts in terms of well-defined probability density functions related to uncertainties in the data, as required by the Bayesian principles that underpin the research. The weights embody evidence of irregularity, and effectively exclude those points identified as non-linearities in the data from the linear fits. Thus, as improved surface estimates lead in turn to better

irregularity estimates, so the linear processes become increasingly effective, and covariance estimates ever more reliable.

In this paper we develop our problem formulation, modelling strategy and algorithms. We identify the important advances the resultant technique offers over the methods introduced by Sander and Zucker [6]. In particular, we emphasise four key features not present in [6]: (1) the propagation of uncertainty throughout processing levels, (2) a mechanism for estimating and weighting against irregularities, (3) the ability to re-adjust surface estimates during processing, and (4) the coupling of parameter estimation and refinements processes.

There are a number of important features which differentiate our refinement technique from the method introduced in Ref. [6]. Firstly, the technique incorporates uncertainty. The influence of a neighbourhood chart on the update equations is weighted by the uncertainty in its differential structure. Secondly, the algorithm is robust. Neighbourhood charts estimated to be irregular with respect to the underlying differential structure at i are weighted against in the fit. Thirdly, the covariance of each chart is iteratively refined based upon the propagation of uncertainties from its neighbourhood. Fourthly, our refinement process is coupled with the parameter estimation process. This arises since estimates of irregular data points depend upon estimates of consistency at the level of refinement. Thus, changes in differential structure through contextual refinement lead to re-estimation of data irregularity, which in turn leads to a re-estimation of surface parameters. An illustration of the different processes and their relationships in the complete system is given in Fig. 1.

The structure of this paper is as follows. In Section 2 we describe the standard representation of differential structure which we employ. In Section 3 we introduce and develop our global MAP approach. The recursive process for surface parameter estimation, its assumed local surface model, and the resulting robust algorithm are described in Section 4. Before consistence can be imposed, differential structures and their associated covariances must be derived from surface parameters and covariances. We show how to do this in Section 5. The relaxation of refinement process, its model components and algorithm are then outlined in Section 6. Results of the applications of the technique to range data are illustrated in Section 7, before some concluding remarks about the work and the direction of future research.

2. Local representation of differential structure

Following [6] we choose as our local representation of differential structure the augmented Darboux frame. The frame contains the set of elements comprising the basic

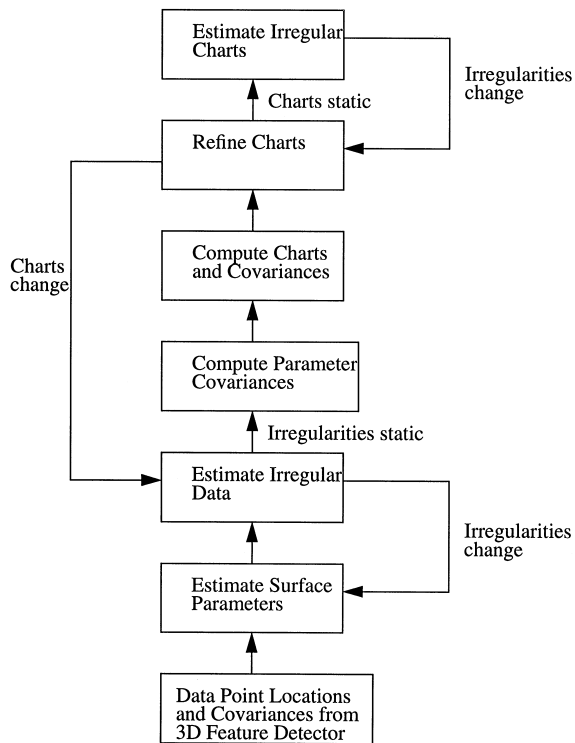


Fig. 1. Illustration of the surface estimation processes and their relationships.




Fig. 2. Representation of the differential structure of a local surface.

local differential geometry of a surface. These elements can be described with the aid of Fig. 2 which shows the local neighbourhood of a point labeled i , at location X_i , on a surface S .

Suppose T_i is the plane tangent to the surface at the point i , and D_i is the plane orthogonal to T_i containing the unit normal vector to i , N_i . As D_i is rotated about N_i it intersects S in a contour called the normal section. The curvature of the normal section at i is called the

normal curvature, K_i . The two special directions in T_i for which K_i takes on maximum and minimum values, K_{M_i} and K_{m_i} are referred to as the principal directions, M_i and m_i . Thus, N_i , M_i and m_i form an orthogonal coordinate system established using the local differential structure. The scalar quantities K_{M_i} and K_{m_i} are referred to as the principal curvatures at i . Following the convention of Ref. [6], the augmented Darboux frame at i is given by the set

$$\Lambda_i = (X_i, N_i, M_i, m_i, K_{M_i}, K_{m_i}). \quad (1)$$

Throughout this paper we refer interchangeably to the set of differential structure components for a point, as given in Eq. (1), as its chart.

3. Global MAP estimation

The basis of our approach is to formulate the recovery of consistent differential structure as a global MAP estimate. In this section we introduce our MAP estimation scheme and show how it leads to separate processes for surface parameter estimation, and chart refinement. Further, we show that optimisation can be realised by the iterative re-assignment of single charts. This has the desirable effect of permitting a manifestly global process to be realised in terms of local computations.

Our starting assumption is that we are given the estimated locations and associated covariances of putative surface points, as returned by some 3D feature detection process. Let the estimated locations and covariances, defined in the global coordinate system be denoted by the sets $X^g = \{X_1^g, \dots, X_N^g\}$, and $W^g = \{W_1^g, \dots, W_N^g\}$ respectively, where N is the number of putative surface points. Further let the estimated local differential structures for the points be denoted by the set $\Lambda = \{\Lambda_1, \dots, \Lambda_N\}$. Our goal is to find the best set of self-consistent charts which is compatible with the estimated location. We formulate this task as finding the chart assignments over the entire set of points which maximises the a posteriori probability of the joint chart assignment $P(\Lambda|X^g)$, i.e., the MAP estimate. The formulation follows along the lines described in detail in Ref. [21] in application to discrete relaxation.

The optimisation procedure over the space of possible global chart assignments, Ω , may be written as the decision rule:

make the global assignment $\Lambda = \omega$ if

$$P(\Lambda = \omega|X^g) > P(\Lambda = \hat{\omega}|X^g) \quad \forall \hat{\omega} \in \Omega, \hat{\omega} \neq \omega. \quad (2)$$

Now, we have chosen a standard representation of differential structure. This is necessary in order to take advantage of the relaxation refinement scheme introduced in Ref. [6]. However, there is no direct link between this representation of differential structure and the esti-

mated locations of putative surface points. Rather, we have to link the two indirectly by introducing surface parameterisations into the MAP estimate at each point. Using the theorem of probability we expand $P(\Lambda = w|X^g)$ over the space of possible global parameter value assignments, Φ

$$\begin{aligned} P(\Lambda = w|X^g) &= \int_{\phi \in \Phi} P(\Lambda = w, \phi|X^g) d\phi \\ &= \int_{\phi \in \Phi} P(\Lambda = w|\phi, X^g)P(\phi|X^g) d\phi. \end{aligned} \quad (3)$$

Our basic assumption is that the integral in Eq. (3) can be approximated by the contribution from the best-fit global parameter values, ϕ^b . Specifically, we assume

$$\begin{aligned} P(\Lambda = w|X^g) &= \int_{\phi \in \Phi} P(\Lambda = w|\phi, X^g)P(\phi|X^g) d\phi \\ &\approx P(\Lambda = w|\phi^b)P(\phi^b|X^g), \end{aligned} \quad (4)$$

where $\phi^b = \{\phi_1^b, \dots, \phi_N^b\}$ is such that

$$P(\phi^b|X^g) > P(\phi|X^g) \quad \forall \phi \in \Phi, \phi \neq \phi^b. \quad (5)$$

From the Bayes rule, the quantity of interest in the MAP estimation process becomes

$$\begin{aligned} P(\Lambda = w|X^g) &\approx P(\Lambda = w|\phi^b)P(\phi^b|X^g) \\ &= \frac{p(\phi^b|\Lambda = w)P(\Lambda = w)P(\phi^b|X^g)}{p(\phi^b)}. \end{aligned} \quad (6)$$

Ignoring the joint density of the denominator, $p(\phi^b)$, and, for the moment, the conditional probability $P(\phi^b|X^g)$, since they are both independent of the charts, there are two quantities which influence the decision process stated in Eq. (2). Firstly, there is the joint conditional density, $p(\phi^b|\Lambda)$, which models the process by which parameter values have been obtained from local differential structures. Secondly, there is the a priori probability of the chart assignments, $P(\Lambda = w)$, which models the consistency of the differential structures.

To develop our iterative procedure consider two realisations of the global assignments in which the chart at point i takes on the values of α and β , respectively. The ratio of the a posteriori probabilities of the two assignments is given by

$$\frac{P(\Lambda_i = \alpha, \Lambda'|X^g)}{P(\Lambda_i = \beta, \Lambda'|X^g)} = \frac{p(\phi^b|\Lambda_i = \alpha, \Lambda')P(\Lambda_i = \alpha, \Lambda')}{p(\phi^b|\Lambda_i = \beta, \Lambda')P(\Lambda_i = \beta, \Lambda')}, \quad (7)$$

where Λ' is the set of $N - 1$ charts, $\{\Lambda_1, \dots, \Lambda_{i-1}, \Lambda_{i+1}, \dots, \Lambda_N\}$. If we assume that the observational process for parameter values is conditionally independent, then we can rewrite the ratio in Eq. (7) as

$$\frac{P(\Lambda_i = \alpha, \Lambda'|X^g)}{P(\Lambda_i = \beta, \Lambda'|X^g)} = \frac{p(\phi_i^b|\Lambda_i = \alpha)P(\Lambda_i = \alpha, \Lambda')}{p(\phi_i^b|\Lambda_i = \beta)P(\Lambda_i = \beta, \Lambda')}, \quad (8)$$

This formula is the basis of our MAP estimate. Eq. (8) suggests a two-stage realisation of the MAP estimation. The stages may be thought of as corresponding to surface parameter estimation and surface refinement, respectively. Firstly, the best-fit parameter values, $\phi^b = \{\phi_1^b, \dots, \phi_N^b\}$, should be computed for the N putative surface points returned by the 3D feature detector. We develop a robust, recursive scheme to do this in the next section. Once this has been achieved the algorithm suggested by Eq. (8) is to make the assignment to the i th chart, $\Lambda_i = \Lambda_i^b$, which maximises the quantity

$$p(\phi_i^b | \Lambda_i) P(\Lambda_i, \Lambda). \quad (9)$$

By systematically repeating the procedure at all N points the MAP estimate will monotonically increase to its global optimum.

Before we can commence with the relaxation refinement process there are two main concerns to address. Firstly, we need to be able to compute best-fit parameter values (and their covariances). We show how to do this in the next section. The second prerequisite is to characterise the relationship between our two surface representations, i.e., the parameters and the charts, in terms of the conditional probability density functions in Eq. (9). We establish the link between the two levels in Section 5. We will return to the global MAP estimate, and the relaxation refinement process, in Section 6.

4. Surface parameter estimation

In this section we commence development of our weighted, least-squares technique for estimating best-fit parameter values for a patch model of a local surface. By using separate probability density functions to characterise the measurement process of points which obey and violate smoothness assumptions, we obtain a robust technique. Further, by basing estimates of violation on previous estimates of best-fit parameters and differential structure, we arrive at a recursive fit process which may be coupled to the surface refinement stage.

In fitting a parametric model at the i th point we seek to find the parameter values $\phi_i = \phi_i^b$ which maximises $P(\phi_i | X^g)$. By application of Bayes rule, and assuming that the data are conditionally independent of one another, then we have

$$P(\phi_i | X^g) = \frac{\prod_j p(X_j^g | \phi_i) P(\phi_i)}{p(X^g)}. \quad (10)$$

Ignoring the joint measure density in the denominator, since this is a static property, there are two terms of interest in Eq. (10). Firstly, there is the a priori probability of surface parameters, $P(\phi_i)$. A number of approaches are possible for modelling this quantity. For example, high curvatures could be penalised. A more common

approach is to assume that all parameter values are equally likely (the maximum entropy assumption). This is the approach we adopt. Secondly, and of more concern to us here, is the conditional density, $p(X_j^g | \phi_i)$, which models the affinity of the estimated location of the j th point to the realisation of the local parametric model for the i th point.

At this point it is worth stressing the novelty of our approach by underlining the modelling demands that it imposes. Much research in 3D scene interpretation has focused on developing strategies for extracting curvature information from smooth surfaces [1–6,11,22]. Our goal is different since we would like to be able to cope with situations where smoothness assumptions break down (e.g. in the presence of outliers, or across ridge and ravine lines, say). The least-means-squares approach of Refs. [6,11], or the precise Kalman filter applications in Ref. [1], are unsuitable for our purposes. We require a robust parameter estimation technique, which will neither over-smooth across boundaries, nor be oversensitive to outliers. The basis for accommodating this prerequisite is to distinguish between the conditional measurement densities for points which obey our smoothness assumptions (we refer to such points as regular points), and points which violate smoothness (these are the non-linearities in the data which we refer to as irregular points). It is the irregularities which we wish to exclude from the fit. If this can be reliably achieved then simple, essentially linear, fitting techniques will be effective, and the covariances of the estimated parameters can be readily and reliably computed and propagated. We will return to this issue later in this section. For the moment we have three modelling requirements: the conditional measurement densities of regular and irregular points, respectively, and an estimate of the probability of smoothness violation. We now consider each in turn.

4.1. Measurement probability density for regular points

We now introduce the parametric model for local smooth surfaces and the appropriate measurements to take for parameter estimation. We characterise these measurements by a probability density function. We consider the local surface at just a single point, but the process is identical for all points returned by a 3D detector. The development proceeds in a local coordinate system. We map data locations and covariances back into global coordinates at the end of the subsection. In the local coordinate system we denote data point coordinates by the set $X^1 = (X_1^1, X_2^1, \dots, X_M^1)$, where M is the number of data points involved in the parametric fit, and $X_i^1 = (x_i^1, y_i^1, z_i^1)^T$, is the local 3D coordinate of an individual point.

Following Refs. [1,6], we choose as our local parametric model of smooth surface the parabolic quadric patch since this is the simplest form which is appropriate for the

computation of curvatures. We assume that the estimated locations of regular points are noisy estimates from a smooth parametric surface which can be expressed in terms of the position variables, u and v ,

$$h(u, v) = (u, v, w(u, v)), \quad (11)$$

where $w(u, v)$ is the mapping from a u - v plane of the local coordinate system onto the surface. We can approximate $w(u, v)$ in the vicinity of the local origin using a Taylor expansion. Since our ultimate goal is the computation of curvature we take the simplest expansion appropriate for this purpose, i.e., we terminate the expansion at the second-order terms,

$$w(u, v) = w|_{0,0} + \left. \frac{\partial w}{\partial u} \right|_{0,0} u + \left. \frac{\partial w}{\partial v} \right|_{0,0} v + \frac{1}{2} \left. \frac{\partial^2 w}{\partial^2 u} \right|_{0,0} u^2 + \left. \frac{\partial^2 w}{\partial u \partial v} \right|_{0,0} uv + \frac{1}{2} \left. \frac{\partial^2 w}{\partial^2 v} \right|_{0,0} v^2. \quad (12)$$

We can now be specific about the physical meaning of our parameters by re-writing Eq. (12) as

$$w(u, v) = A(u, v)\phi, \quad (13)$$

where

$$A(u, v) = \left(\frac{1}{2}u^2, uv, \frac{1}{2}v^2, u, v, 1 \right)^T \quad (14)$$

and ϕ is a vector of surface derivatives,

$$\phi = \left(\left. \frac{\partial^2 w}{\partial^2 u} \right|_{0,0}, \left. \frac{\partial^2 w}{\partial u \partial v} \right|_{0,0}, \left. \frac{\partial^2 w}{\partial^2 v} \right|_{0,0}, \left. \frac{\partial w}{\partial u} \right|_{0,0}, \left. \frac{\partial w}{\partial v} \right|_{0,0} \right)^T = (a, b, c, d, e, f)^T. \quad (15)$$

Parameters a , b and c relate to the curvature of the surface (at $u = v = 0$), parameters d and e relate to the surface normal, and the parameter f relates to surface position. In line with the conventional approach adopted elsewhere, we assume that the appropriate measurements to take in estimating parameters are the residuals $r = r_1, \dots, r_M$, given by

$$r_j = A(x_j^l, y_j^l)\phi - z_j^l \quad (16)$$

and that a Gaussian distribution, with zero mean is appropriate for the residuals. That is,

$$p(r_j | j \in S_i^f) = \frac{1}{\sqrt{2\pi\sigma_{r_j}^2}} \exp\left(-\frac{r_j^2}{2\sigma_{r_j}^2}\right). \quad (17)$$

In a manner similar to that employed in Ref. [1], we relate the uncertainty of the measurement for point j to the covariance of its location, W_j^l , defined in the local coordinate system. This is achieved via the first order approximation

$$\sigma_{r_j}^2 = E[\Delta r_j \Delta r_j^T] = J_j W_j^l J_j^T, \quad (18)$$

here J_j is the 1×3 Jacobian matrix for the transformation, given by

$$J_j = \left(\phi \frac{\partial}{\partial x_j^l} A(x_j^l, y_j^l), \phi \frac{\partial}{\partial y_j^l} A(x_j^l, y_j^l), -1 \right) = (ax_j^l + by_j^l + d, bx_j^l + cy_j^l + e, -1). \quad (19)$$

We now show how to propagate data locations and covariances back into the global coordinate system. If a data point is at a location $(x_j^g, y_j^g, z_j^g)^T$ in global coordinates then to express it on the local coordinate system (P, Q, N) at location O we compute

$$\begin{pmatrix} x_j^l \\ y_j^l \\ z_j^l \end{pmatrix} = R \begin{pmatrix} x_j^g - O_x \\ y_j^g - O_y \\ z_j^g - O_z \end{pmatrix}, \quad (20)$$

where

$$R = \begin{pmatrix} P_x & P_y & P_z \\ Q_x & Q_y & Q_z \\ N_x & N_y & N_z \end{pmatrix}. \quad (21)$$

Assuming that the covariance in the location of a point as returned by the 3D feature detector is also given in the global coordinate system, then the corresponding 3×3 matrix, W_j^g , in the local coordinate system (P, Q, N) at O , is given by

$$W_j^l = R W_j^g R^T. \quad (22)$$

We have now described all the elements required by the measurement probability density function for regular points. We have assumed that regular points are noisy estimates from a smooth surface which can be approximated locally by parabolic quadric patch. We further assumed that the appropriate measurements to take were fit-residuals, and that they follow a Gaussian distribution with zero mean. We then related the standard deviation of the residual measurement for a point to the covariance of its location. We now turn our attention to characterising irregular points.

4.2. Measurement probability density for irregular points

We adopt a very different modelling strategy for irregular points, i.e., outliers or points across surface boundaries. Specifically, we assume a uniform probability density for their measurements, i.e.,

$$p(r_j | j \notin S_i^f) = q_1. \quad (23)$$

This approach has the dual advantages of being simple, and facilitating the estimation of parameter

covariances (we will justify this later in Section 4.5). The consequence of assuming a uniform distribution is that estimated irregularities are ignored by the surface fit algorithm (see Section 4.4). This is in line with the philosophy adopted in the application of standard robust estimation technique [15,17,18].

4.3. Estimation of irregular points

We have now specified measurement density functions for both regular and irregular points. However, as yet we have no mechanism for determining which density function is appropriate. Therefore, there are two distinct estimates of interest: $P(j \in O_i^f)$, the probability that joint j is an outlier with respect to the local surface at i , and $P(j \in B_i^f)$, the probability that j lies across a boundary from i . Under the independence assumption the probability of irregularity is

$$P(j \notin S_i^f) = P(j \in O_i^f \text{ or } j \in B_i^f) \\ = P(j \in O_i^f) + P(j \in B_i^f) - P(j \in O_i^f)P(j \in B_i^f). \quad (24)$$

Of course, it is only possible to make reliable estimates of irregular points if a good estimate of the underlying surface itself is available, and vice versa. However, the underlying surface may not be available since it is the very thing that we are attempting to compute. One approach is to base the estimation of an irregularity at some time step, n say, upon the most recently available (stable) surface estimates. This is the approach taken here. Since estimates are made prior to each re-estimation of the surface parameters we are decoupling non-linearity estimation from the fit process itself. This has important consequences. First of all, it permits the use of powerful techniques for estimating irregularity, without worrying about introducing additional complexity into the fitting algorithm. In turn, the power of the irregularity estimator means that relatively simple fitting techniques are effective. As a consequence, after fitting has terminated, reliable covariances for surface parameters may be readily computed and propagated to the surface refinement level. We now consider the outlier process and the boundary process in turn.

4.3.1. Outlier process

We commence by specifying more precisely what we mean by an outlier, and describe how it can be identified. Conventionally, robust estimation techniques have based outlier detection solely upon fit-residual measurements. However, Parent and Zucker [20] have shown in relation to 2D images that high-order features up to and including curvature variation are necessary to properly discriminate between the influence of nearby curves. An analogous situation arises for surface in 3D, and provides the motivation for including high-order features into the outlier model. One approach might be to extract orienta-

tions and curvatures from surface-fit parameters. Unfortunately, we take the view that high-order derivatives are highly susceptible to noise [20], and we prefer to extract them from the refined charts later on in processing, since they provide for more stable and consistent estimates.

We consider a point j to be an outlier with respect to the underlying local surface at point i if it lies off the surface, or if its associated (refined) chart is inconsistent with the true differential structure at i . We state more precisely what is meant by consistency of differential structure in Section 6. The basic idea that we wish to exploit is that, by extrapolating outwards from the local surface patch point i , it is possible to obtain a prediction for the appearance of the differential structure at point j , giving rise to a measure of compatibility or consistency. This prediction can then be compared with the actual chart estimate at point j . Since the representation of differential structure we employ incorporates information up to and including curvature, so derivatives up to curvature variation are implicitly included in the consistency measure by virtue of the extrapolation process. The consistency measure is formulated within the Bayesian framework which underpins our approach.

Under the assumptions outlined above, the probability of an outlier is

$$P^{(n)}(j \in O_i^f) = P(j \in O_i^f | r_j^{(n-1)}, \Lambda_{ij}^{e(n-1)}, \Lambda_j^{(n-1)}). \quad (25)$$

Applying Bayes rule then

$$P^{(n)}(j \in O_i^f) = \frac{p(r_j | j \in O_i^f)P(\Lambda_j, \Lambda_{ij}^e, j \in O_i^f)}{p(r_j | j \in O_i^f)P(\Lambda_j, \Lambda_{ij}^e, j \in O_i^f) + p(r_j | j \notin O_i^f)P(\Lambda_j, \Lambda_{ij}^e, j \notin O_i^f)}, \quad (26)$$

where Λ_{ij}^e is the result of extrapolating the chart at the point i to the j th point (see Section 6), and we have temporarily dropped the $(n-1)$ suffixes on the right-hand side of Eq. (26). The conditional measurement densities in Eq. (26) are available from Sections 4.1 and 4.2. The joint chart probabilities relating to consistency of differential structure are given in Section 6.

4.3.2. Boundary process

The approach we have taken to modelling surface boundaries is simple, computationally attractive, and works well for the range data applications reported in Section 7. We estimate points lying across depth discontinuities alone, based upon the difference between the depths returned by the 3D feature detector, and the depths predicted by the surface parameters. Specifically, we take

$$P^{(n)}(j \in B_i^f) = \exp\left(-\frac{(z_j - z_{ij}^p)^2}{2\sigma_z^2}\right), \quad (27)$$

where σ_z^2 is a weighted estimate of the variances in surface position, given by

$$\sigma_z^2 = \frac{\sum_k P^{(n-1)}(k \in S_i^f) (z_j - z_{ij}^p)^2}{\sum_k P^{(n-1)}(k \in S_i^f)}, \quad (28)$$

where z_j is the z -value of the j th point, and z_{ij}^p is the z -value predicted by the i th patch.

We recognise the limitations of this approach. In particular, only depth discontinuities are estimated, no structural information is taken into account, and measurements are viewpoint-dependent. Work currently in progress addresses these issues. For example, we are considering an orientation discontinuity process, and the incorporation of boundary models based on spatial coherence.

4.4. Algorithm

We have now met our three modelling requirements as specified at the start of this section, namely, appropriate measurement probability density functions for regular and irregular points, respectively, and a method for estimating the probability that data is regular. We are now in a position to specify the fit algorithm itself. Restating Eq. (4.1), the requirement is to find $\phi_i = \phi_i^b$ which maximises

$$P(\phi_i|X) = \frac{\prod_j P(X_j|\phi_i)P(\phi_i)}{p(X)}, \quad (29)$$

where, from Sections 4.1 and 4.2,

$$p(X_j|\phi_i) = \begin{cases} \frac{k}{\sigma_{r_j}} \exp\left(-\frac{r_j^2}{2\sigma_{r_j}^2}\right) & \text{if } j \in S_i^f, \\ q_1 & \text{if } j \notin S_i^f \end{cases} \quad (30)$$

and k is a constant. Given that the joint probability density in the denominator is independent of the parameters, and assuming that the prior probability of parameter values, $P(\phi_i)$, is constant over all possible values, then the logarithm of Eq. (30) can be written as

$$E_i^{(n)}(\phi_i|X) = -\sum_j w_{ij}^{(n)} \left(\frac{r_j^2}{2\sigma_{r_j}^2} - \ln\left(\frac{k}{\sigma_{r_j}}\right) \right) + (1 - w_{ij}^{(n)}) \ln q_1 + c_1, \quad (31)$$

where c_1 is a constant, and we have made the substitution

$$w_{ij}^{(n)} = P^{(n)}(j \in S_i^f) \quad (32)$$

which is available from Section 4.3. Now, in line with our approach of decoupling irregularity estimation from the fitting process, then $w_{ij}^{(n)}$ is constant throughout optimisation of Eq. (30). As a consequent, the quantity to maximise takes on a relative simple form. It is effectively

a weighted least-means estimate, with a correction term since measurement uncertainty depends upon patch parameters:

$$E_i^{(n)}(\phi_i|X) = -\sum_j w_{ij}^{(n)} \left(\frac{r_j^2}{2\sigma_{r_j}^2} - \log\left(\frac{k}{\sigma_{r_j}}\right) \right) + c_2. \quad (33)$$

Maximisation of Eq. (33) may be carried out by gradient descent. The procedure at each iteration is to update parameter values,

$$\phi \rightarrow \phi + \Delta\phi, \quad (34)$$

where

$$\Delta\phi = -\eta \nabla_p E_i^{(n)} = \eta \sum_j w_{ij}^{(n)} \left(\left(\nabla_p r_j \frac{r_j}{\sigma_{r_j}^2} \right) + \left(\nabla_p \sigma_{r_j} \left(-\frac{1}{\sigma_{r_j}} - \frac{r_j^2}{\sigma_{r_j}^3} \right) \right) \right). \quad (35)$$

Unfortunately, instabilities can occur in the gradient descent procedure where the surface gradient is large [23]. To alleviate this problem, and in line with the predominant approach in the literature [1,6,11,12], we break optimisation into two phases. The phases correspond to (1) surface position and surface normal estimation, and (2) surface curvature estimation. To be more specific, in the first phase we fix those parameters relating to second-order surface derivatives (refer to Eq. (15)) to zero, i.e., $a = b = c = 0$, and concentrate solely upon estimating the zeroth- and first-order parameters, d, e and f . Gradient descent takes place in the global coordinate system. Once complete, estimates of surface position and surface normal can then be extracted from the parameters (see Section 5.1) and used to establish a local coordinate system such that, by definition, the gradient at the origin with respect to the parameters is zero, and the problem of instability is mitigated. This is achieved simply by aligning the w -axis with the surface normal, and setting the origin to lie on the surface. We can then proceed to estimate the parameters relating to curvature in relative safety, provided first-order parameters are fixed at zero in the local frame during this second phase.

At this point it is worth stressing four key features of the novel fitting algorithm which differentiate it from the most commonly used paradigm for surface parameters estimation [1,6,11]. Firstly, the algorithm incorporates measurement uncertainty. Clearly, from Eq. (35), points whose measurement uncertainty, σ_{r_j} , is relatively high will have less bearing on the estimation process. Monga et al. [18] have demonstrated that improved patch parameters can be obtained by including measurement uncertainties into a Kalman filter technique. Within our formation measurement uncertainties are incorporated into a gradient descent paradigm. Secondly, our algorithm is robust. We have a mechanism for estimating and

excluding data outliers or data points which lie across the surface discontinuities. This contrasts both with the least-means-squares approach of [6,11] and the Kalman filter application of [18]. Thirdly, the fit process is recursive. The idea is that improved estimates of irregularities will in turn lead to better surface patch estimates, and vice versa. This process should be contrasted with the technique of Sander and Zucker [6] wherein surface patches are extracted in a one-off process, and remain static thereafter. Fourthly, estimation and refinement processes become coupled. Therefore, improved chart estimates also facilitate improved surface estimates, and vice versa.

4.5. Summary of fit procedure

A summary of the main processing steps in the surface fitting process for a single data point is given here. The procedure is identical for all points, and may be carried out in parallel.

(1) Estimate surface position and surface normal by fitting a plane to data points in the global coordinate system. (2) Determine the local coordinate system whose w-axis is aligned with the surface normal and whose origin lies on the surface. (3) Estimate surface curvature by fitting a parabolic quadric patch to data points in this local coordinate system. (4) Make re-estimates of data irregularities based upon the extracted parameters and (refined) charts, if available. (5) If the new estimates are different from the old estimates go to step 1, else stop the fitting procedure.

Once fitting has been terminated covariances may be computed. We now show how to do this.

4.6. Computing the covariance for best-fit parameters

We now derive a relationship between the covariance of estimated parameters and the uncertainties in the putative surface point locations. To proceed, we assume that all data points are regular, i.e., noisy estimates from an underlying smooth surface. Of course, in general, this will not be the case, and we will weight against irregular points at the end of the section. We assume that application of Eq. (35) solves the set of M simultaneous equations,

$$\tilde{A}(x, y)\phi - z = 0, \quad (36)$$

where the $M \times 6$ matrix

$$\tilde{A}(x, y) = (A(x_1, y_1), \dots, A(x_i, y_i), \dots, A(x_M, y_M))^T \quad (37)$$

and

$$\begin{aligned} x &= (x_1, \dots, x_i, \dots, x_M)^T, \\ y &= (y_1, \dots, y_i, \dots, y_M)^T, \\ z &= (z_1, \dots, z_i, \dots, z_M)^T. \end{aligned} \quad (38)$$

We further assume that the estimated locations of regular points are related to their true, underlying values by

$$\begin{aligned} x &= x^t + \Delta x, \\ y &= y^t + \Delta y, \\ z &= z^t + \Delta z, \end{aligned} \quad (39)$$

where

$$\begin{aligned} \Delta x &= (\Delta x_1, \dots, \Delta x_i, \dots, \Delta x_M)^T, \\ \Delta y &= (\Delta y_1, \dots, \Delta y_i, \dots, \Delta y_M)^T, \\ \Delta z &= (\Delta z_1, \dots, \Delta z_i, \dots, \Delta z_M)^T. \end{aligned} \quad (40)$$

Then, by substituting Eq. (39) into Eq. (36) we have

$$\tilde{A}(x^t + \Delta x, y^t + \Delta y)(\phi + \Delta\phi) = z^t + \Delta z. \quad (41)$$

Taking a Taylor expansion, and ignoring terms which are second order and above in the subsequent expression, we obtain

$$\tilde{A}(x^t, y^t)\Delta\phi + (\phi \nabla_x \tilde{A}(x^t, y^t))\Delta x + (\phi \nabla_y \tilde{A}(x^t, y^t))\Delta y = \Delta z. \quad (42)$$

Rearranging terms we have

$$\Delta\phi = \tilde{A}^{-1}(x^t, y^t)B(\phi, x^t, y^t, \Delta x, \Delta y, \Delta z), \quad (43)$$

where B is the $M \times 1$ column,

$$\begin{aligned} B(\phi, x^t, y^t, \Delta x, \Delta y, \Delta z) &= (b_1(\phi, x^t, y^t, \Delta x, \Delta y, \Delta z), \dots, b_M(\phi, x^t, y^t, \Delta x, \Delta y, \Delta z))^T \end{aligned} \quad (44)$$

whose i th entry, b_i , is given by

$$\begin{aligned} b_i(\phi, x^t, y^t, \Delta x, \Delta y, \Delta z) &= \left(\phi \frac{\partial}{\partial x_i^t} A(x_i^t, y_i^t), \phi \frac{\partial}{\partial y_i^t} A(x_i^t, y_i^t), -1 \right) (\Delta x_i, \Delta y_i, \Delta z_i)^T \end{aligned} \quad (45)$$

From Eq. (43) the covariance of ϕ , W_ϕ is then

$$W_\phi = E[\Delta\phi\Delta\phi^T] = \tilde{A}^{-1}E[BB^T](\tilde{A}^{-1})^T, \quad (46)$$

where, for brevity, we have written

$$\begin{aligned} \tilde{A} &\equiv \tilde{A}(x, y), \\ B &\equiv B(\phi, x, y, \Delta x, \Delta y, \Delta z). \end{aligned} \quad (47)$$

Now, assuming that the uncertainties in data points locations are independent of one another, i.e.

$$E[(\Delta x_i, \Delta y_i, \Delta z_i)(\Delta x_i, \Delta y_i, \Delta z_i)^T] = \begin{cases} W_j^i & \text{if } i = j, \\ 0_{3 \times 3} & \text{if } i \neq j, \end{cases} \quad (48)$$

then, from Eqs. (48) and (19), the entries, C_{ij} , of the $M \times M$ matrix $C = BB^T$, have expected values,

$$E[C_{ij}] = J_i E[(\Delta x_i, \Delta y_i, \Delta z_i)(\Delta x_i, \Delta y_i, \Delta z_i)^T] J_i^T \\ = \begin{cases} \sigma_{r_i}^2 & \text{if } i = j, \\ 0 & \text{if } i \neq j \end{cases} \quad (49)$$

and the expression for the covariance of parameters takes on a relatively simple form,

$$W_\phi = W[\Delta\phi\Delta\phi^T] = A^{-1}D(A^{-1})^T = (A^T D^{-1} A)^{-1}, \quad (50)$$

where D is an $M \times M$ diagonal matrix whose i th diagonal entry, $d_i = \sigma_{r_i}^2$. It is easily shown that W_ϕ^{-1} can be expressed as

$$W_\phi^{-1} = \begin{pmatrix} \sum e_i x_i^4 & \sum e_i x_i^3 y_i & \sum e_i x_i^2 y_i^2 & \sum e_i x_i^3 & \sum e_i x_i^2 y_i & \sum e_i x_i^2 \\ \sum e_i x_i^3 y_i & \sum e_i x_i^2 y_i^2 & \sum e_i x_i y_i^3 & \sum e_i x_i^2 y_i & \sum e_i x_i y_i^2 & \sum e_i x_i y_i \\ \sum e_i x_i^2 y_i^2 & \sum e_i x_i y_i^3 & \sum e_i y_i^4 & \sum e_i x_i y_i^2 & \sum e_i y_i^3 & \sum e_i y_i^2 \\ \sum e_i x_i^3 & \sum e_i x_i^2 y_i & \sum e_i x_i y_i^2 & \sum e_i x_i^2 & \sum e_i x_i y_i & \sum e_i x_i \\ \sum e_i x_i^2 y_i & \sum e_i x_i y_i^2 & \sum e_i y_i^3 & \sum e_i x_i y_i & \sum e_i y_i^2 & \sum e_i y_i \\ \sum e_i x_i^2 & \sum e_i x_i y_i & \sum e_i y_i^2 & \sum e_i x_i & \sum e_i y_i & \sum e_i \end{pmatrix}, \quad (51)$$

where all summations are over the data points in the fit, labelled $i = 1, \dots, M$, and,

$$e_i = \frac{1}{\sigma_{r_i}^2}. \quad (52)$$

We have thus far assumed that all data points are regular. To make the estimate of covariance robust we simply re-define the weights, e_i , as,

$$e_i = \frac{P(i \in S^f)}{\sigma_{r_i}^2}. \quad (53)$$

We have now given the necessary details to compute the surface parameters' covariance matrix. The covariance has been linked to the uncertainties in the fit-residual measurements. The residual uncertainties were themselves related to the covariances of data point locations in Section 4.1. We have excluded estimated irregular points from the estimate, in line with our robust approach.

5. Calculation of charts and their covariances

In the previous section we were concerned with a parametric representation of local surface. We showed how to compute best-fit surface parameters and their covariances. In this section we move on to our second level of surface representation, namely local differential structures or charts. Such a presentation is necessary given our stated aim of taking advantage of the relaxational refinement scheme introduced in Ref. [6], which imposes consistency at the level of differential structure. However, in contrast to the work reported in Ref. [6], we are interested in incorporating uncertainty (and model robustness, see Section 6) into the refinement stage. As such, not only charts but also their covariances are required. In this section we show how they can be obtained from the parameter values and covariances returned by the fitting process.

5.1. Chart elements in terms of patch parameters

In this section we state how to compute elements of basic differential structure from surface parameters. The process is identical to that described in Refs. [6,11]. We include it here for the sake of completeness only, and recognise that the reader may wish to move directly on to Section 5.2. To begin, we detail the process when parameters are defined in some arbitrary local coordinate system (with axes labelled u, v and w , respectively). We then show how the expressions for chart elements are successively simplified if the local coordinate system is transformed such that (1) its origin lies on the (estimate) surface and its w -axis is aligned with the surface normal, and (2) its u - and v -axis are rotated to be in alignment with the principal directions of the surface. The first transformation is utilised during the fitting process. The second transformation is required on termination of the fit, but prior to parameters covariance estimation. It leads to a particular simple expression for chart elements, and importantly, greatly simplifies the computation of chart covariances.

According to Section 4.1, the local parameterisation of a surface, S , can be expressed as

$$h(u, v) = (u, v, w(u, v)), \quad (54)$$

where taking terms up to and including second order in the Taylor expansion for $w(u, v)$, we obtain the parabolic quadric patch.

$$w(u, v) = au^2 + buv + cv^2 + du + ev + f. \quad (55)$$

We wish to compute the elements of basic differential structure as defined in Section 2, i.e. the set (X, N, K_M, K_m, M, m) . We consider each element in turn. The location of the estimated surface is given by

$$X = (0, 0, w(0, 0)) = (0, 0, f) \quad (56)$$

and its normal to the surface, N , is given by

$$N = \frac{(-w_u, -w_v, 1)}{\sqrt{1 + w_u^2 + w_v^2}} \tag{57}$$

Further, it can be shown [24] that,

$$dN \begin{pmatrix} du \\ dv \end{pmatrix} = A \begin{pmatrix} du \\ dv \end{pmatrix} \tag{58}$$

with

$$A = \begin{pmatrix} a_{11} & a_{12} \\ a_{21} & a_{22} \end{pmatrix} = \frac{1}{PR - Q^2} \begin{pmatrix} qQ - pR & rQ - qR \\ pQ - qP & qQ - rP \end{pmatrix} \tag{59}$$

and

$$p = \frac{w_{uu}}{\sqrt{1 + w_u^2 + w_v^2}},$$

$$P = 1 + w_u^2,$$

$$q = \frac{w_{uv}}{\sqrt{1 + w_u^2 + w_v^2}},$$

$$Q = w_u w_v,$$

$$r = \frac{w_{vv}}{\sqrt{1 + w_u^2 + w_v^2}},$$

$$R = 1 + w_v^2, \tag{60}$$

where the eigenvalues of A are the principal curvature, K_M and K_m , and the eigenvectors of A are the principal directions M and m . Explicitly, the principal curvature and principal directions are given by

$$K_M = -\frac{1}{2}(a_{11} + a_{22} - s),$$

$$K_m = -\frac{1}{2}(a_{11} + a_{22} + s),$$

$$M = \begin{cases} (a_{12}, -\frac{1}{2}(a_{11} - a_{22} + s)), & a_{11} \geq a_{22}, \\ (\frac{1}{2}(a_{11} - a_{22} - s), a_{21}), & a_{11} < a_{22}, \end{cases}$$

$$m = \begin{cases} (\frac{1}{2}(a_{11} - a_{22} + s), a_{21}), & a_{11} \geq a_{22}, \\ (-a_{12}, \frac{1}{2}(a_{11} - a_{22} - s)), & a_{11} < a_{22}, \end{cases} \tag{61}$$

where s is defined by

$$s = \sqrt{(a_{11} - a_{22})^2 + 4a_{21}a_{12}}. \tag{62}$$

The expressions in Eq. (61) give the elements of differential structure in a general local coordinate system.

Now consider the particular case of determining the differential structure for a local parameterisation $(u, v, w(u, v))$ in which the w -axis is aligned with the surface normal, N , and the origin of the local frame lies just on the surface. In such a case by definition,

$$(w_u|_{0,0}, w_v|_{0,0}, w|_{0,0}) = (d, e, f) = (0, 0, 0) \tag{63}$$

and the elements of differential structure are given by

$$X = (0, 0, 0),$$

$$N = (0, 0, 1),$$

$$K_M = a + c + \sqrt{(a - c)^2 + b^2},$$

$$K_m = a + c - \sqrt{(a - c)^2 + b^2}, \tag{64}$$

$$M = \begin{cases} (a - c + \sqrt{(a - c)^2 + b^2}, b), & a \geq c, \\ (b, a - c - \sqrt{(a - c)^2 + b^2}), & a < c, \end{cases}$$

$$m = \begin{cases} (-b, a - c + \sqrt{(a - c)^2 + b^2}), & a \geq c, \\ (a - c - \sqrt{(a - c)^2 + b^2}, b), & a < c. \end{cases} \tag{65}$$

Suppose that the local tangent plane is now rotated about the w -axis so as to coincide with the principal directions, M and m , as given in Eq. (65) this implies that $b = 0$ and the elements of differential structure have a particularly simple form

$$X = (0, 0, 0),$$

$$N = (0, 0, 1),$$

$$K_M = a,$$

$$K_m = c,$$

$$M = \begin{cases} (1, 0), & a \geq c, \\ (0, 1), & a < c, \end{cases}$$

$$m = \begin{cases} (0, 1), & a \geq c, \\ (1, 0), & a < c. \end{cases} \tag{66}$$

We now have a way of manipulating the local coordinate system so that the resultant parameter values lead to simple expressions for the elements of differential structure. The transformations described in this section are utilised both during and after the fitting process. Refer back to the fitting procedure outlined in Section 4.5. The first phase in the process is the fitting of a plane to data points in the global coordinate system. The resultant estimated parameter value will be of the form $\phi = (0, 0, 0, d, e, f)^T$, since the second-order parameters are fixed to zero. The second phase concerns estimation of curvature. Unfortunately, as stated previously, estimation is unstable at high gradients. To alleviate this

problem we transform the local coordinate system so that its w -axis is aligned with the surface normal, and its origin lies just on the surface. Then, by definition, the estimated surface gradient at the origin will be zero. The appropriate transformation can be determined using Eq. (61). In this local coordinate system the set of parameter values will be $\phi = (0, 0, 0, 0, 0, 0)^T$. After estimating curvatures the parameters will be of the form $\phi = (a, b, c, 0, 0, 0)^T$. By computing the principal directions from Eq. (65), and subsequently rotating the u - and v -axis about the w -axis so that they align with these directions, the parameter values will be transformed (66). Furthermore, by computing parameter covariances in this particular coordinate system, calculation of chart covariances will be greatly simplified, as we illustrate below.

5.2. Covariance of chart elements in terms of covariance of parameters

We make the simplifying assumption that all the elements of differential structure are independent. While this assumption is naive and frequently violated, it makes computation feasible. This is the approach taken in Refs. [6,11]. (In fact, in Refs. [6,11] uncertainties are ignored all together). So, our reduced task is to relate the covariance of each element of differential structure to the covariance of surface parameters. We achieve this through first-order approximations.

Under our independence assumption, the 14×14 covariance matrix for differential structures, W_Λ , is

$$W_\Lambda = \begin{pmatrix} \begin{pmatrix} W_X & 0_{3 \times 3} \\ 0_{3 \times 3} & W_N \end{pmatrix} & 0_{6 \times 6} & 0_{2 \times 6} \\ 0_{6 \times 6} & \begin{pmatrix} W_M & 0_{3 \times 3} \\ 0_{3 \times 3} & W_m \end{pmatrix} & 0_{6 \times 6} \\ 0_{2 \times 2} & 0_{6 \times 2} & \begin{pmatrix} W_{K_M} & 0 \\ 0 & W_{K_m} \end{pmatrix} \end{pmatrix}, \quad (67)$$

where W_X , W_N , W_M and W_m are 3×3 covariance matrices associated with estimates of surface position, surface normal and surface principal directions, respectively, and W_{K_M} and W_{K_m} are the variances on principal curvatures. We will not go through the process of calculating all the elements of W_Λ here. The general procedure for each element, $T = F_T(\phi)$, say, is to make the first-order approximation,

$$\Delta T \approx (\nabla_\phi F_T(\phi)) \Delta \phi. \quad (68)$$

The covariance of T , W_T is then given by

$$\begin{aligned} W_T &= E[(\nabla_\phi F_T(\phi)) \Delta \phi][(\nabla_\phi F_T(\phi)) \Delta \phi]^T \\ &= \nabla_\phi F_T(\phi) W_\phi (\nabla_\phi F_T(\phi))^T. \end{aligned} \quad (69)$$

As an illustrative example, the covariance for surface normal is computed by noting that

$$N = F_N(\phi) = \frac{(-w_u, -w_v, 1)}{\sqrt{1 + w_u^2 + w_v^2}} = \frac{(-d, -e, 1)}{\sqrt{1 + d^2 + e^2}}. \quad (70)$$

Given that the parameter values are of the form $(a, 0, c, 0, 0, 0)$ then

$$\begin{aligned} \nabla F_N(\phi) &= \begin{pmatrix} \frac{\partial N_x}{\partial a} & \frac{\partial N_x}{\partial b} & \frac{\partial N_x}{\partial c} & \frac{\partial N_x}{\partial d} & \frac{\partial N_x}{\partial e} & \frac{\partial N_x}{\partial f} \\ \frac{\partial N_y}{\partial a} & \frac{\partial N_y}{\partial b} & \frac{\partial N_y}{\partial c} & \frac{\partial N_y}{\partial d} & \frac{\partial N_y}{\partial e} & \frac{\partial N_y}{\partial f} \\ \frac{\partial N_z}{\partial a} & \frac{\partial N_z}{\partial b} & \frac{\partial N_z}{\partial c} & \frac{\partial N_z}{\partial d} & \frac{\partial N_z}{\partial e} & \frac{\partial N_z}{\partial f} \end{pmatrix} \\ &= \begin{pmatrix} 0 & 0 & 0 & -1 & 0 & 0 \\ 0 & 0 & 0 & 0 & -1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \end{pmatrix} \end{aligned} \quad (71)$$

which is in a particularly simple form. W_N follows directly from Eq. (69).

We now have described how to instantiate chart elements and their covariances. We now return to our global MAP estimation scheme, and to the surface refinement process for improving the consistency of instantiated charts.

6. Surface refinement

Recently, Sander and Zucker have introduced an iterative minimisation procedure for local surface refinement [6]. The basis for their algorithm is to optimise the local consistency of differential structure. The basic idea is that, since the surface at each point is represented by a parabolic quadric patch, it is possible to extrapolate outward from a point j , say, to its neighbour i , in order to get an idea of what the differential surface at i looks like according to j . If the differential structure extrapolated to i is close to the actual differential structure there, then chart i is said to be consistent with the chart at j (note that i consistent with j , does not necessarily imply j consistent with i). Now, in imposing consistency, Sander and Zucker perform the extrapolation procedure for a set of neighbours to each point and update each chart via a least-means-squares estimate. Chart consistency forms the basis for our refinement strategy and we adopt their extrapolation method. Details of the method are beyond the scope of this paper but may be found in Ref. [6]. However, by contrast, we wish to avoid the ad hoc least-squares approach to minimisation adopted in Refs. [6,11] since it is well known that it oversmooths at genuine surface discontinuities and boundaries and it is over sensitive to outliers. We require a robust technique,

which also incorporates uncertainty. Our approach follows that employed at the level of surface parameter estimation, and similarly, we arrive at a weighted-least-squares fit. The linearity facilitates the straightforward computation and propagation of covariances, while the weights embody increasingly reliable estimates of non-linearity (i.e., irregularity). Consequently, the linear fitting techniques become ever more appropriate and effective as irregularities are more reliably estimated and excluded from the fit.

We now return to our global MAP estimate which was introduced in Section 2. The appropriate ratio in the MAP estimate is

$$\frac{p(\phi_i^b|\Lambda_i = \alpha)P(\Lambda_i = \alpha, \Lambda')}{p(\phi_i^b|\Lambda_i = \beta)P(\Lambda_i = \beta, \Lambda')} \quad (72)$$

There are two quantities of interest here. Firstly, there is the conditional probability density function, $p(\phi_i^b|\Lambda_i)$, which models the process by which the surface parameters for a point have been obtained from its chart. We cannot immediately link charts and parameters, since they comprise different surface representations. However, in Section 5, we showed how to compute a chart and its covariance from the set of best-fit parameter values and its covariance. This allows us to express the conditional density function as the Gaussian distribution:

$$p(\phi_i^b|\Lambda_i) = p(\Lambda_i^\phi|\Lambda_i) = \frac{1}{(2\pi)^7 \sqrt{|W_i^\phi|}} \times \exp\left(-\frac{1}{2}(\Lambda_i^\phi - \Lambda_i)(W_i^\phi)^{-1}(\Lambda_i^\phi - \Lambda_i)^T\right), \quad (73)$$

where Λ_i^ϕ is the chart computed from the best-fit parameters values ϕ_i^b and W_i^ϕ is its corresponding covariance.

The second quantity of interest in Eq. (72) is the joint probability, $P(\Lambda_i, \Lambda')$, which models the consistency of differential structure. Following a standard approach [6,22], we assume pairwise interactions for charts. The ratio of joint probabilities relevant to the MAP estimate can then be written as

$$\frac{P(\Lambda_i = \alpha, \Lambda')}{P(\Lambda_i = \beta, \Lambda')} = \frac{\prod_{j \in N_i} P(\Lambda_i = \alpha, \Lambda_j)}{\prod_{j \in N_i} P(\Lambda_i = \beta, \Lambda_j)}, \quad (74)$$

where N_i is some neighbourhood set for point i . We model the joint probability that they are pairwise consistent, i.e. i is consistent with j and j is consistent with i . That is,

$$P(\Lambda_i, \Lambda_j) = P(\Lambda_i|\Lambda_{ji}^e)P(\Lambda_j|\Lambda_{ij}^e), \quad (75)$$

where Λ_{ji}^e is the extrapolation of the chart at point j to point i . Substituting Eq. (75) into Eq. (74) gives

$$\frac{P(\Lambda_i = \alpha, \Lambda')}{P(\Lambda_i = \beta, \Lambda')} = \frac{\prod_{j \in N_i} P(\Lambda_i = \alpha|\Lambda_{ji}^e)P(\Lambda_j|\Lambda_{ij}^e = \alpha_j^e)}{\prod_{j \in N_i} P(\Lambda_i = \beta|\Lambda_{ji}^e)P(\Lambda_j|\Lambda_{ij}^e = \beta_j^e)}. \quad (76)$$

Unfortunately, incorporation of the conditional probability $P(\Lambda_j|\Lambda_{ij}^e)$ into Eq. (76) leads to an intractable non-linear update equation of Λ_i . This arises because the extrapolation method itself is non-linear, drawing on second derivative approximations of local structure. We are currently investigating the use of a linear extrapolation scheme to overcome this problem. However, in the work reported here we have taken the extreme step of ignoring this term completely. This is equivalent to assuming that the consistency of a chart depends solely upon extrapolations from, and not to, its neighbours. This assumption it also made by Sander and Zucker in their relaxational refinement technique. (In fact, asymmetry of compatibility is frequently admitted in relaxation schemes.) It ensures that the refinement algorithm is linear, hence covariances can be easily updated.

Using Bayes rule, and assuming, a priori, that all single-chart assignments are equiprobable, then we can rewrite Eq. (76) as

$$\frac{p(\phi_i^b|\Lambda_i = \alpha) \prod_{j \in N_i} p(\Lambda_{ji}^e|\Lambda_i = \alpha)}{p(\phi_i^b|\Lambda_i = \beta) \prod_{j \in N_i} p(\Lambda_{ji}^e|\Lambda_i = \beta)}. \quad (77)$$

Our development of a robust technique at the level of surface refinement proceeds along similar lines to those adopted at the surface fit level. We wish to estimate which charts violate the local smooth surface model so that they do not bias the fit. The approach taken is to distinguish between charts which obey smoothness assumptions (which we term regular charts), and those which violate smoothness (which we term irregular charts), and to treat each differently. We have three modelling components: the probability density functions of extrapolations from regular and irregular charts, respectively, and an estimate of smoothness violation itself. A brief outline of each is now given.

6.1. Extrapolation density function of regular charts

For regular charts we assume a Gaussian distribution for their extrapolations

$$p(\Lambda_{ji}^e|\Lambda_i, j \in S_i^r) = \frac{1}{(2\pi)^7 \sqrt{|W_{ji}^e|}} \times \exp\left(-\frac{1}{2}(\Lambda_{ji}^e - \Lambda_i)(W_{ji}^e)^{-1}(\Lambda_{ji}^e - \Lambda_i)^T\right), \quad (78)$$

where here, and throughout the rest of this section, we dropped the suffix labelling the chart covariance.

6.2. Extrapolation density function of irregular charts

For irregular charts we assume a uniform probability distribution

$$p(\Lambda_{ji}^e|\Lambda_i, j \notin S_i^r) = q_2. \quad (79)$$

6.3. Estimation of irregular charts

The chart at point j is deemed to be irregular with respect to the underlying surface at i , if its differential structure is inconsistent with the structure of the true surface. Obviously, the true differential structure is not available — it is what we are trying to compute. Taking a similar approach to that employed at the level of surface fit in estimating irregular data, we base chart irregularity upon the most recent surface estimates available. As a consequence, we are decoupling the estimation of chart irregularities from the refinement itself. Estimation takes place prior to refinement, and estimates remain fixed during a refinement iteration. This means we need not worry about introducing additional complexity into the refinement algorithm. In particular, we can safely extrapolate from the i th chart to its neighbourhood, without introducing non-linearity into the algorithm. We have

$$P^{(n)}(j \in O_i^r) = P(j \in O_i^r | \Lambda_{ij}^{e(n-1)}, \Lambda_j^{(n-1)}). \quad (80)$$

Application of the Bayes rule gives

$$P^{(n)}(j \in O_i^r) = \frac{p(\Lambda_j | \Lambda_{ij}^e, j \in O_i^r) P(\Lambda_{ij}^e, j \in O_i^r)}{p(\Lambda_j | \Lambda_{ij}^e, j \in O_i^r) P(\Lambda_{ij}^e, j \in O_i^r) + p(\Lambda_j | \Lambda_{ij}^e, j \notin O_i^r) P(\Lambda_{ij}^e, j \notin O_i^r)}, \quad (81)$$

where we have dropped the suffixes on the right-hand side of Eq. (81).

We have not incorporated a boundary process in the refinement stage in the work reported here and we take

$$P^{(n)}(j \in S_i^r) \equiv 1 - P^{(n)}(j \in O_i^r). \quad (82)$$

However, we are currently investigating the introduction of shape-index labels [25,26] to characterise boundary structures, using a dictionary-based relaxation labelling technique [21,27].

6.4. Algorithm

We have now described the three modelling components requested in Section 6.1, and are in a position to specify the algorithm for relaxational refinement of consistency. Referring back to our global MAP estimate in Eq. (77), the aim is to find $\Lambda_i = \Lambda_i^b$ which maximises

$$p(\Lambda_i^\phi | \Lambda_i) \prod_{j \in N_i} p(\Lambda_{ji}^e | \Lambda_i). \quad (83)$$

Maximising Eq. (83) is equivalent to solving, for Λ_i , the equation

$$(\Lambda_i^\phi - \Lambda_i)(W_i^\phi)^{-1}(\Lambda_i^\phi - \Lambda_i)^T + \sum_{j \in N_i} w_{ji}^{(n)}((\Lambda_{ji}^e - \Lambda_i)(W_{ji}^e)^{-1}(\Lambda_{ji}^e - \Lambda_i)^T) = 0, \quad (84)$$

where we have made the substitution

$$w_{ji}^{(n)} = P^{(n)}(j \in S_i^r). \quad (85)$$

This equation can be rewritten, to emphasise its weighted least-means-squares nature

$$\sum_k c_k (\Lambda_i - \Lambda_k^\circ) W_k^{\circ-1} (\Lambda_i - \Lambda_k^\circ)^T = 0, \quad (86)$$

where

$$c = (1, w_{1i}^{(n)}, \dots, w_{Mi}^{(n)}),$$

$$\Lambda^\circ = (\Lambda_i^\phi, \Lambda_{1i}^e, \dots, \Lambda_{Mi}^e),$$

$$W^\circ = (W_i^\phi, W_{1i}^e, \dots, W_{Mi}^e) \quad (87)$$

and M is the number of points within the neighbourhood set N_i . This has the solution

$$\Lambda_i = \left(\sum_k c_k (W_k^\circ)^{-1} \right)^{-1} \sum_k c_k (W_k^\circ)^{-1} \Lambda_k^\circ \quad (88)$$

with covariance

$$W_i = \left(\sum_k c_k (W_k^\circ)^{-1} \right)^{-1}. \quad (89)$$

We have now elucidated our algorithm for iterative surface refinement to improve the consistency of differential structures based upon the propagation of contextual information. By systematically repeating the update algorithms (88) and (89) at all N putative surface points, the MAP estimate will increase to its global optimum, as required.

7. Experimentation

Our research goal is to develop a framework for interpreting the differential structure of 3D scenes represented in terms of density slice data, rather than range data. This is due to its greater 3D coverage and wealth of information concerning the sensed scene. However, in the experiments reported in this paper we have taken range data as our input. The main reasons for this are the saving in computation over slice data, and the ease with which range data can be visualised. In fact, data for experimentation originated from the estimated locations of 3D feature points returned by a multi-frame feature tracking technique [28] applied to cranial MRI slice data. To convert the data into range form we projected a ray from each point (i, j) in a discrete 2D grid onto the set of 3D feature points, and measuring the distance, $d(i, j)$, to the first ‘hit’. Fig. 3 shows the range data obtained by projecting onto the front and back views of the 3D data points.




Fig. 3. Range data generated for front and back views of head.

Our technique requires as input not only the estimated locations of putative surface points, but also estimates of their associated covariances. In experiments we took the covariance of range data points (in the global coordinate system with the z -axis aligned with the depth direction) to be of the form

$$W_i = \begin{pmatrix} 0 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & \sigma^2 \end{pmatrix}. \quad (90)$$

Techniques for estimating σ are available, but in all the experiments reported here we took $\sigma = 0.1$. We took neighbourhood regions for both parameter estimation and surface refinement to be 7×7 pixels, and the step size used in the gradient descent procedure in Section 4 was 0.1 throughout. Two further parameters must be chosen before experimentation can proceed. These are q_1 and q_2 , the height and of the uniform distributions assumed for outliers at the levels of fit and refinement, respectively. In estimating q_1 we choose that value which gives a suitably low probability of a point being classed as an outlier if its fit residual is zero. Referring to Eq. (26), if the residual is zero and assuming no prior information is available, then we have

$$P(j \in O^f) = \frac{q_1}{2\pi\sigma_j^2 + q_1}. \quad (91)$$

In experiments we have taken the probability of an outlier for zero fit residual to be 0.001 and hence $q_1 = (2\pi\sigma^2)/999$. The parameter q_2 is determined similarly.

We have applied both least-squares estimators and our robust techniques to the height data in Fig. 3. We use two

images in our experiments. The top view is of the front of a head and includes features such as eye-sockets and ears. The bottom view is of the back of the head. The results of our experiments are shown in Figs. 4–9. In each case we present our results in four panels. The top panel is the reconstructed height data. The second panel shows the reconstructed surface orientation. In the third panel we show the extracted mean and Gaussian curvature labels. Finally, the bottom panel shows the principal curvature direction. Elements of differential structures estimated at the level of surface parameter estimation and surface refinement are not given.

Fig. 4 shows the estimated differential structures after application of the standard one-pass least-squares estimator to the front-of-head data. Fig. 5 shows the corresponding results obtained after five time steps (i.e. four re-estimations of irregular data) by our robust surface parameter estimation technique. It is clear from the figures that the robust approach avoids many of the problems of oversmoothing and oversensitivity to noise associated with the least-squares estimator. For example, compare the surface normals in the two figures. The improved performance of our technique is particularly evident where smoothness assumptions break down, for example, at the sides of the head and around the ears and eyes. Notice how the data outliers above the right eye are well handled in our method but that the conventional approach is oversensitive to them. The results for surface refinement are illustrated in Figs. 6–9. Figs. 6 and 8 display the refined differential structures for front and back of head, respectively, after 5 iterations of the least-squares approach to imposing consistency. The corresponding results for our technique are given in Figs. 7 and 9. Again, it is clear that the robust method handles genuine surface boundaries and the effects of noise rather well, whereas the conventional method oversmooths.




Fig. 4. Differential structure using least-means estimator for surface fit (front of head).




Fig. 5. Differential structure using robust estimator for surface fit (front of head).




Fig. 6. Differential structure using least-means estimator for surface refinement (front of head).




Fig. 7. Differential structure using robust estimator for surface refinement (front of head).




Fig. 8. Differential structure using least-means estimator for surface refinement (back of head).




Fig. 9. Differential structure using robust estimator for surface refinement (back of head).

8. Discussion

In this paper we have formulated a technique for extracting locally consistent differential structure for surfaces, as a step towards our goal in establishing a framework for 3D scene interpretation. Our approach incorporated standard models of local (smooth) surfaces at the levels of surface fit and surface refinement, namely the parabolic quadric patch and consistency of differential structure, respectively. Its novelty stemmed from (1) its formulation as a global MAP estimate within a Bayesian framework, and (2) the decoupling of the irregularity estimation process from the fits themselves. This led to iterative, robust techniques wherein evidential support for data representations was readily computed and readjustable. As processing proceeded, more reliable evidential support led to better surface estimates, and vice versa. By application to range data we demonstrated that the resultant technique can overcome some of the problems of oversmoothing and oversensitivity to noise associated with least-squares estimators.

In our view, there are two main shortcomings associated with the work reported here. Firstly, there is no representation of boundary structures in our formulation [17]. Secondly, in moving from surface parameter to chart representations of surface (as required by the extrapolation method adopted) some information is lost (in particular we are forced to assume that chart elements are independent). We are currently addressing both these issues. In particular, we are (1) developing and incorporating a boundary-labelling process [21,27] based on the spatial coherence of shape-index labels [25,26], and (2) introducing an extrapolation scheme which acts on the surface parameter representation.

A further concern is to apply our method of 3D slice data. Of interest is the possibility of developing a unified evidence-combining scheme which can be applied directly to the raw slice images. Evidential support for data representations at any one of the feature detection, parameter estimation or surface refinement levels would be determined by information from all the other levels simultaneously. This would increase the scope for their effective adaptation. In this way, the band limitation of the conventional sequential multi-stage philosophy to 3D scene interpretation would be mitigated.

References

- [1] O. Monga, N. Ayache, P. Sander, Using uncertainty to link 3D edge detection and local surface modelling, Proceedings of the Seventh Scandinavian Conference on Image Analysis, 1991, pp. 273–284.
- [2] O. Monga, N. Ayache, From voxel to curvature, IEEE Computer Vision and Pattern Recognition Conference, 1991, pp. 644–649.
- [3] O. Monga, R. Deriche, G. Malandain, Recursive filtering and edge closing: two primary tools for 3D edge detection, *Image Vision Comput.* 9 (1991) 203–214.
- [4] O. Monga, N. Ayache, P. Sander, Modelling uncertainty for estimating local surface geometry, Proceedings of the Seventh Scandinavian Conference on Image Analysis, Vol. 1, 1992, pp. 379–382.
- [5] P.T. Sander, S.W. Zucker, Singularities of principal direction field from 3D images, *IEEE Pattern Anal. Mach. Intell. PAMI* 14 (1992) 309–317.
- [6] P.T. Sander, S.W. Zucker, Inferring surface structure and differential structure from 3D images, *IEEE Pattern Anal. Mach. Intell. PAMI* 12 (1990) 833–854.
- [7] O. Monga, S. Beayoun, Using differential geometry in to extract typical features in 3D density images, 11th International Conference on Pattern Recognition, Vol. 1, 1992, pp. 379–382.
- [8] A. Sha'ashua, S. Ullman, Structural saliency: the detection of globally salient structures using a locally connected network, Proceedings of the Second International Conference on Computer Vision, 1988, pp. 321–327.
- [9] J.F. Canny, A computational approach to edge detection, *IEEE Pattern Anal. Mach. Intell. PAMI* 8 (1986) 679–698.
- [10] R. Deriche, Optimal edges detection using recursive filtering, *Int. J. Comput. Vision* (1987) 167–187.
- [11] F.P. Ferrie, J. Lagarde, P. Whiate, Darboux frames, snakes and super-quadratics: geometry from the bottom up, *IEEE Pattern Anal. Mach. Intell. PAMI* 15 (1993) 771–784.
- [12] P.J. Besl, R.C. Jain, Segmentation through variable order surface fitting, *IEEE Pattern Anal. Mach. Intell. PAMI* 10 (1988) 167–192.
- [13] P.J. Flynn, A.K. Jain, On reliable curvature estimation, *IEEE Computer Vision and Pattern Recognition Conference*, 1989, pp. 110–116.
- [14] R.M. Haralick, Propagating covariance in computer vision, Proceedings of the 12th ICPR, 1994, pp. 493–498.
- [15] M.J. Black, P. Anandan, A framework for robust estimation of optical flow, Proceedings of the Fourth International Conference on Computer Vision, 1993, pp. 231–237.
- [16] A. Blake, A. Zissermann, *Visual Reconstruction*, MIT Press, Boston, 1987.
- [17] F.R. Hampel, E.M. Ronchetti, P.J. Rousseeuw, W.A. Stahel, *Robust Statistics: The Approach Based on Influence Functions*, Wiley, New York, 1986.
- [18] P. Meer, D. Mintz, A. Rosenfeld, Robust regression methods for computer vision: a review, *Int. J. Comput. Vision* 6 (1991) 59–70.
- [19] T. Poggio, V. Torre, C. Koch, Computational vision and regularisation theory, *Nature* 317 (1985) 314–319.
- [20] P. Parent, S.W. Zucker, Trace inference, curvature consistency and curve detection, *IEEE Pattern Anal. Mach. Intell. PAMI* 11 (1989) 823–839.
- [21] E.R. Hancock, J. Kittler, Discrete relaxation, *Pattern Recognition* 23 (1990) 711–733.
- [22] M. Kass, A. Witken, M. Terzopoulos, Constraints on deformable models: Recovering 3D Shape from non-rigid motion, *Artif. Intell.* 36 (1988) 91–123.
- [23] R. Johnson, D. Wichern, *Applied Multivariate Statistical Analysis*, Prentice-Hall, Englewood Cliffs, NJ, 1982.

- [24] M.P. do Carmo, *Differential Geometry of Curves and Surface*, Prentice-Hall, Englewood Cliffs, NJ, 1976.
- [25] J.J. Koenderink, *Solid Shape*, MIT Press, Boston, 1990.
- [26] J.J. Koenderink, Surface shape and curvature scales, *Image Vision Comput.* 10 (1992) 557–565.
- [27] E.R. Hancock, J. Kittler, Adaptive estimation of hysteresis thresholds, *Proceedings IEEE Computer Vision and Pattern Recognition Conference*, 1991, pp. 196–201.
- [28] N. Sharp, E.R. Hancock, Multi-frame feature tracking by probabilistic relaxation. *Proceedings of the Fifth British Machine Vision Conference*, 1994, pp. 407–418.

About the Author—MICK TURNER gained a first-class B.Sc. degree in physics and a Ph.D. for research in speech recognition from the University of Manchester. From 1990 to 1998 he worked as a research associate in the Department of Computer Science at the University of York. During this period he undertook research on projects concerned with chromosome recognition, 3D surface analysis and molecular matching. In 1998 he co-founded a company, “Pattern Computing”, to provide consultancy in the application of pattern-recognition techniques to problems from industry. His research interests are in statistical pattern recognition, neural networks and image analysis.

About the Author—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. He is currently Professor of Computer Vision in the Department of Computer Science at the University of York where he leads a group of some 15 researchers in the areas of computer vision and pattern recognition. Professor Hancock has published about 200 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 outstanding paper award in 1997. Professor Hancock serves as an Associate Editor of the journals *IEEE Transactions on Pattern Analysis and Machine Intelligence*, and, *Pattern Recognition*. He 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.