

A Bayesian Approach to 3D Surface Fitting and Refinement.

M. Turner and E.R. Hancock

Department of Computer Science, University of York

York, YO1 5DD, UK

email: mick@minster.york.ac.uk

Abstract

This paper presents a novel methodology for extracting differential structure from noisy and discontinuous surfaces. We formulate the problem as a global MAP estimate which we develop under Bayesian principles. Development leads to two distinct MAP estimates for initial surface patch fitting and for subsequent surface refinement. Each estimate is realised through an iterative scheme which gauges the effects on consistency of updating single local surfaces. In turn, each update is itself realised by iterating between two modes, for surface smoothing and outlier identification, respectively. Smoothing reduces to a simple weighted least-squares technique which delivers robust surface estimates and covariances. Outlier estimation is achieved by a powerful non-linear method which incorporates feedback. As a consequence, the two processing levels become strongly coupled since curvature information from the refinement level can influence patch estimation. The net effect is a robust, evidence-combining scheme in which early fit and segmentation errors can be rectified. We demonstrate the utility of the scheme on range data.

1 Introduction

Conventionally, researchers have adopted a 4-level, hierarchical approach to volumetric scene interpretation in which increasingly condensed, salient surface descriptions are built up in a sequential fashion [2, 4, 9, 10, 11]. This approach is exemplified by the work of Sander and Zucker [11]. Initially, a 3D feature extraction technique is employed to detect the locations of putative surface points. In the next stage surface parameterisations are estimated by fitting parabolic quadric patches to these estimated locations in a conventional least-means approach. At the third level in the hierarchy the differential structure of local surfaces is extracted from the surface parameters and subsequently refined. The refinement aims to improve the curvature consistency of neighbouring surfaces through a relaxation scheme. This is achieved through the repeated optimisation of a least-squares error metric. Finally, with consistent differential structure to hand, the extraction of global surfaces and curves may proceed based upon mean and Gaussian curvature labels [11].

We identify three limitations associated with this approach. Firstly, the surface descriptions lack any measure of uncertainty. Specifically, there are no covariances associated with data delivered to each processing stage. Secondly, the descriptions do not accommodate outliers. This can give rise to gross fitting errors which then propagate within and between levels. Thirdly, the bottom-up nature of processing means that there is no mechanism for using high-level information to rectify errors. This is a severe limitation when the extraction of fine curvature detail is requested since early processing tends to be unreliable.

Individual techniques exist with the potential for separately addressing each (but not all) of the shortcomings mentioned above. Monga, Ayache and Sander [9] have recognised the need to attach a measure of uncertainty to surface estimates. Accordingly, they propagate the covariances of putative surface points to the surface parameterisation stage. However, covariances are under-utilised in the iterative refinement itself. Moreover, both the surface estimates and the covariances are unreliable in the presence of outliers.

The need to identify and exclude outliers from fitting (or smoothing) has long been accepted. The simplest approaches attempt identification (or labeling) in a one-off pre-processing stage [7]. However, reliable identification can only take place if reliable estimates of the surface are to hand, which is obviously not the case a priori. Regularisation [13] and robust estimation [3, 7] techniques attempt simultaneous labeling and smoothing, but are restricted since only relatively simple outlier models can be employed without introducing unmanageable complexity into the algorithms. Yuille [15] extends the Graduated Non-Convexity idea of Blake and Zisserman [3] using mean-field theory to average out the line (or discontinuity) process. This has the effect of smoothing over outlier structure, and as such, important information is lost. The EM algorithm [5] adopts a different strategy in the fitting of incomplete data. Processing is split into two modes, for parameter estimation and likelihood maximisation respectively. Optimal surface descriptions may be located by iterating between these two modes [1]. Estimated probability densities naturally weight against outliers in the robust fitting of surfaces, without introducing undue additional complexity. However, the approach has traditionally been applied at a single level of representation [1]. For surfaces, multiple levels of representation appear necessary, and both intra-level and inter-level constraints are important.

In this paper we address the problem of extracting differential surface structure from noisy and discontinuous 3D data. Our primary concern is to address the issues of uncertainty, robustness and feedback mentioned above. We begin by formulating the problem as a global MAP estimate, and we develop the estimate under Bayesian principles. Our development leads to two distinct MAP estimates, for initial surface patch estimation (surface fitting), and subsequent surface refinement. Each estimate is realised through an iterative scheme which gauges the effects of updating single local surfaces on intra-level and inter-level consistency. In gauging consistency we wish to differentiate between outliers and genuine surface structure in order to make the scheme robust. Since outliers are not known a priori we adopt an EM-type approach. Specifically, each single update is realised by iterating between two modes, for smooth surface estimation and outlier identification, respectively. Smoothing reduces to a simple least-squares technique

which delivers both robust surface estimates and their robust covariances. Outlier identification takes place within a non-linear, probabilistic labeling scheme which incorporates feedback. As such, the fitting and refinement processing levels become strongly coupled since, not only do patch estimates affect the charts, but curvature information at the refinement level can influence patch estimation itself. The net effect is a robust, evidence-combining scheme in which early fit and segmentation errors can be rectified.

The structure of this paper is as follows. In section 2 we introduce and develop the global MAP estimate. Models of inter-level and intra-level surface consistency are introduced in section 3, along with a brief outline of our approach to modeling outliers. The resulting algorithms are given in section 4. We illustrate the application of the scheme to range data in section 5.

2 Process

We begin by assuming that we have N nodes which are spatially organised in some way. We further assume that we are given the estimated location and associated covariance of the putative surface point at each node, 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 $w^0 = \{w_0^0, \dots, w_N^0\}$, and $W^0 = \{W_0^0, \dots, W_N^0\}$ respectively. Further, let the local differential structures, or charts, for the nodes be denoted by the set $s^2 = \{s_0^2, \dots, s_N^2\}$. (The notation will become clearer shortly). We formulate the task as finding the global chart assignment, $s^2 = \omega^2$, over the space of possible global assignments, S^2 , which maximises the a posteriori probability $P(s^2 | w^0)$, i.e., the MAP estimate. In other words we require:

$$\omega^2 =: \arg \max_{\hat{s}^2 \in S^2} P(s^2 = \hat{s}^2 | w^0) \quad (1)$$

Now, we choose a standard representation of differential structure, i.e., the augmented Darboux frame [11]. (This describes a local surface's position, normal, principal directions and principle curvatures). Its choice is necessary in order to take advantage of the relaxational refinement scheme introduced in [11]. However, there is no direct link between this representation and the estimated locations of putative surface points. Rather, we have to link the two indirectly by introducing surface parameterisations into the MAP estimate at each point. Specifically, we introduce the parabolic quadric patch parameterisation. Using the theorem of total probability we expand $P(s^2 = \hat{s}^2 | w^0)$ over the space of possible global patch assignments, S^1 , i.e.,

$$P(s^2 = \hat{s}^2 | w^0) = \int_{s^1 \in S^1} P(s^2 = \hat{s}^2 | s^1) P(s^1 | w^0) ds^1 \quad (2)$$

Our basic assumption is that the integral in (2) can be approximated by the contribution from the best-fit global patch assignments. That is, we assume

$$P(s^2 = \hat{s}^2 | w^0) \approx P(s^2 = \hat{s}^2 | \omega^1) P(\omega^1 | w^0) \quad (3)$$

where ω^1 is given by,

$$\omega^1 = \arg \max_{\hat{s}^1 \in S^1} P(s^1 = \hat{s}^1 | w^0) \quad (4)$$

Further, since the optimal patches are static with respect to the charts, then equation (1) may be re-written as

$$\omega^2 = \arg \max_{\hat{s}^2 \in S^2} P(s^2 = \hat{s}^2 | w^1) \quad (5)$$

Thus, our task is now reduced to a two-stage process, which is in line with conventional approaches [11]. Firstly, there is the optimisation of patch parameters given the (static) putative surface point locations. Secondly, there is the optimisation of charts given the (static) optimal patches.

2.1 Developing MAP Estimates

The two quantities of interest in the MAP estimates in (4) and (5) are of the form $P(s^k | w^{k-1})$, where $k = 1, 2$ indexes the level of surface representation. Here, w^{k-1} may be viewed as static data delivered by the $k - 1$ level of surface representation, and s^k may be viewed as surface parameters to be estimated. (For $k = 1$ the data are the putative surface point locations, and the parameters relate to patches, while for $k = 2$ the optimal patches serve as data, and parameters relate to charts). Our approach to developing the MAP estimates follows along the lines taken in [8] in relation to discrete relaxation labeling. In [8] the aim is to optimise label assignments on nodes in such a way as to achieve a MAP estimate which is globally consistent. This is achieved through the iterative replacement of single label assignments. Effectively, the scheme gauges the relative effects on structural consistency of different label assignments and selects the best. Our development differs in that we are dealing with continuous parameter assignments rather than discrete labels. Further, we need to gauge the effects of an update on intra-level and inter-level smooth surface consistency.

Consider two realisations of the assignments at level k , in which the parameters for node i takes on the values α and β respectively. In particular, consider the ratio

$$\frac{P(s_i^k = \alpha, \tilde{\omega}^k | \omega^{k-1})}{P(s_i^k = \beta, \tilde{\omega}^k | \omega^{k-1})} \quad (6)$$

where $\tilde{\omega}^k = \omega_j^k, \forall j \in N, j \neq i$. If we can systematically find, for each node i , that assignment α for which the ratio in (6) is $> 1, \forall \beta \neq \alpha$ then the MAP estimate will monotonically increase to its global optimum [8].

We develop this ratio by applying Bayes' rule to the numerator and denominator. Assuming that the data are conditionally independent of one another, and further, that all the information from the data relating to the surface estimate is conveyed by nodes within some local spatial neighbourhood, N_i^{k-1} , then the probability density functions in (6) may be expanded to give

$$\frac{P(s_i^k = \alpha, \tilde{\omega}^k | \omega^{k-1})}{P(s_i^k = \beta, \tilde{\omega}^k | \omega^{k-1})} = \frac{\prod_{j \in N_i^{k-1}} p(\omega_j^{k-1} | s_i^k = \alpha) P(s_i^k = \alpha, \tilde{\omega}^k)}{\prod_{j \in N_i^{k-1}} p(\omega_j^{k-1} | s_i^k = \beta) P(s_i^k = \beta, \tilde{\omega}^k)} \quad (7)$$

for $k = 1$ or $k = 2$. There are two quantities of interest here. Firstly, there is $p(\omega_j^{k-1} | s_i^k)$ which models the inter-level affinity of the two local surfaces. Secondly, there is the quantity $P(s_i^k = \alpha, \tilde{\omega}^k)$ which models intra-level surface consistency. We now introduce our models for these quantities. Our approach requires a third modeling ingredient, for outliers, and we shall outline this as well.

3 Models

3.1 Inter-level Surface Consistency

We are interested in modeling the inter-level affinity of two local (smooth) surfaces. We assume that the affinity may be characterised via a d -dimensional measurement vector, r_{ji}^{k-1} , taken between the two surfaces, which is of the form $r_{ji}^{k-1} = A_{ji}^{k-1} s_i^k - B_{ji}^{k-1}$ where $B_{ji}^{k-1} \equiv B_{ji}^{k-1}(s_j^{k-1})$ is some (possibly non-linear) transformation of the ‘datum’, s_j^{k-1} , which makes it suitable for comparison with the surface estimate, s_i^k . We further assume that if both the datum and the surface estimate belong to the true local surface for node i , then the measurement follows a Gaussian distribution with zero mean. In contrast, if either the datum or the local surface estimate is an outlier with respect to the true local surface then a uniform distribution is appropriate. Formally, we have

$$p(s_j^{k-1} | s_i^k) = \begin{cases} G_d(r_{ji}^{k-1}, W_{ji}^{k-1}) & \text{if } l_{ii}^k = S \text{ and } l_{ji}^{k-1} = S \\ q & \text{otherwise} \end{cases} \quad (8)$$

where l_{ji}^k labels the membership of the j th local surface estimate, at representational level k , to the underlying (i.e., true) local surface at point node i , S denotes the true (local) surface class, and G_d is the d -dimensional Gaussian distribution

$$G_d(r, W) = \frac{1}{(2\pi)^{n/2} |W|^{1/2}} \exp\left(-\frac{1}{2} r^T W^{-1} r\right) \quad (9)$$

In this paper we take standard measurements at both levels. For $k = 1$ the measurement of interest is simply the scalar difference in height between patch and putative surface point location. Details may be found in [14]. For $k = 2$ the measurement is the vector difference between s_i^2 and the chart extracted from s_j^1 . Note that the extracted chart is obtained in a two-stage process involving (a) initial extraction from the patch parameters and (b) transformation to make the extracted chart suitable for comparison with the surface at node i . We adopt the extrapolation scheme of Sander and Zucker for this purpose [11]. We also require the covariance on the extrapolated chart in order to estimate W_{ji}^{k-1} , and we obtain this through first-order approximation.

3.2 Intra-Level Surface Consistency

At the level of patch estimation we assume that the local surface at each point is independent of all other local surfaces. The idea is that the role of the intermediate surface representation is purely to deliver data for chart estimation. We adopt a pairwise model of local (curvature) consistency for charts. This follows the approach taken in [11]. Formally,

$$P(s_j^2, \forall j \in N_i^2) = \prod_{j \in N_i^2} P(s_i^2 | s_j^2) \quad (10)$$

where N_i^2 is some non-inclusive local spatial neighbourhood for node i . Our approach to modeling $P(s_j^2 | s_i^2)$ is analogous to that used in connection with inter-level parameter affinity, as detailed in the previous subsection. Specifically, we take

$$P(s_j^2 | s_i^2) = \begin{cases} G_d(r_{ji}^2, W_{ji}^2) & \text{if } l_{ii}^2 = S \text{ and } l_{ji}^2 = S \\ q & \text{otherwise} \end{cases} \quad (11)$$

where the measurement is of the form $r_{ji}^2 = A_{ji}^2 s_i^2 - B_{ji}^2$. Again, we use the extrapolation scheme of Sander and Zucker [11] to obtain $B_{ji}^2 \equiv B_{ji}^2(s_j^2)$. Taking $A_{ji}^2 = I$ the measurement is simply the vector difference between the chart extrapolated to node i , and the estimated chart there.

3.3 Outliers

The models of pairwise surface affinity introduced in the previous subsections necessitate a mechanism for estimating outlying surfaces. Specifically, we request a means of computing the probability that a local (smooth) surface estimate belongs to underlying local surfaces. Of course, accurate outlier identification can only take place if reliable surface estimates are to hand, and this is obviously not the case a priori. To overcome this dilemma we adopt an EM style approach [5]. The idea is that the updating of each local surface description as indicated in (9) is split into two modes, for surface smoothing and outlier identification, respectively. Each single update is then realised by iterating between the two modes, with an improved smooth surface leading to better estimation of data outliers, and vice versa, until convergence. Under this methodology, outlier estimates remain static during smoothing, and likewise, surface estimates remain static during outlier identification. Within our probabilistic framework, the quantity of interest in modeling outliers is the *a posteriori* label probability $P(l_{ji}^k | s)$ for $k = 0, 1, 2$, and where $s = \{s^0, s^1, s^2\}$ is the complete set of current (static) surface estimates. A full discussion of the development of this quantity is beyond the scope of this paper. The basic idea is to expand the *a posteriori* label probability using Bayesian principles, incorporating both the models of smooth surface consistency in (8) and (11), and a model of outlier structure. We are currently incorporating an outlier model which contains both intra-level and inter-level constraints on spatial configurations. However, in the work reported here we have adopted a much simpler model which ignores spatial context. Details of this model are not given here, but

may be found in [14]. Of primary concern is that the model incorporates feedback. In particular, outlier estimates at the chart representational level constrain the labeling probabilities for both patches and putative surface points. If inconsistencies arise between levels then the implication is that early segmentation was erroneous, and that patches should be re-estimated. Thus fitting and refinement levels become strongly coupled since, not only do the patches affect the charts, but curvature information extracted at the refinement level may be employed to improve the patch estimates themselves. This abandons the conventional sequential processing philosophy. We now have three levels of iteration in processing: between patch and chart levels of representation, at a single level to realise its MAP estimate, and between smoothing and outlier modes.

4 Algorithms

4.1 Patch Estimation

Substituting in our model for intra-level consistency in subsection 3.2, and assuming that all surface parameter assignments are equiprobable a priori, then the ratio of interest for patch estimation is

$$\frac{P(s_i^1 = \alpha, \tilde{\omega}^1 | \omega^0)}{P(s_i^1 = \beta, \tilde{\omega}^1 | \omega^0)} = \frac{\prod_{j \in N_i^0} p(\omega_j^0 | s_i^1 = \alpha)}{\prod_{j \in N_i^0} p(\omega_j^0 | s_i^1 = \beta)} \quad (12)$$

According to (12), a single update at a single node requires an assignment to s_i^1 which maximises the product $\prod_{j \in N_i^0} p(\omega_j^0 | s_i^1)$. By substituting in the expression for measurement probability density functions in (8), it is straightforward to show that maximisation becomes equivalent to finding that assignment to s_i^1 which minimises

$$\sum_{j \in N_i^0} r_{ji}^0 (e_{ji}^{01}) r_{ji}^{0T} \quad (13)$$

where e_{ji}^{01} , the inter-level evidential support matrix for datum ω_j^0 which couples smoothing and outlier modes, is given by

$$e_{ji}^{01} = -\log(2\pi^{d/2} |W_{ji}^0|^{1/2}) (W_{ji}^0)^{-1} P(l_{ii}^1 = S) P(l_{ji}^0 = S) \quad (14)$$

There are two components to the evidential support. Firstly, there are the probabilities of datum or patch belonging to the true local surface. If the datum is an outlier then its measurement is excluded from the summation in (13). If the patch estimate is an outlier, then all measurements taken from it are invalid. This makes estimation robust. (Recall that, under our approach to processing label probabilities are treated as static quantities during parameter optimisation). Secondly, there is a contribution from measurement uncertainty. In general, this contribution is non-linearly related to surface parameters, and this can make optimisation unstable, unreliable or intractable. For these reasons we treat measurement uncertainties as independent of changes in parameters during optimisation

of (13) itself. Rather, they are pre-determined, and static, based upon the parameter values determined at last update. This assumption simplifies the resulting fit algorithm. It implies that evidential support is static during optimisation. In such a circumstance we have a weighted least-squares problem with solution

$$s_i^1 = \left(\sum_{j \in N_i^0} A_{ji}^0 e_{ji}^{01} (A_{ji}^0)^T \right)^{-1} \sum_{j \in N_i^0} A_{ji}^0 e_{ji}^{01} B_{ji}^0 \quad (15)$$

and the associated robust covariance is readily computed [14]. Each update of the single-level MAP estimate in (4) involves iteration between this least-squares estimator and the outlier estimator at each node. At convergence, the next update may commence. In turn, at convergence of the MAP estimate chart estimation may proceed.

4.2 Chart Estimation

In a similar vein, if we incorporate our model of chart consistency, and assume that chart assignments are equiprobable a priori, then the MAP estimate in (5) requires an assignment to the chart s_i^2 which minimises the quantity

$$\sum_{p=1}^{p=2} \sum_{j \in N_i^p} r_{ji}^p (e_{ji}^{h(p)}) r_{ji}^{pT} \quad (16)$$

where $h(1) = \{1, 2\}$ and $h(2) = \{2, 2\}$. By once again interpreting evidential support (both inter-level and intra-level) as static during optimisation, we arrive at a weighted least-squares problem whose solution is

$$s_i^2 = \sum_{p=1}^{p=2} \left(\sum_{j \in N_i^p} A_{ji}^p e_{ji}^{h(p)} (A_{ji}^p)^T \right)^{-1} \sum_{j \in N_i^p} A_{ji}^p e_{ji}^{h(p)} B_{ji}^p \quad (17)$$

Again, it is a straightforward matter to estimate the covariance of chart estimates [14]. Each update of the MAP estimate in (5) involves iteration between the least-squares estimator in (17) and outlier estimation. However, at convergence we cannot immediately begin the next MAP update. Rather, we must first check to see if the update has changed the evidential support for the data in patch estimation. This may arise through the feedback of outlier information. If this is the case, the implication is that errors have been made in segmentation, and so patches should be re-estimated. Note a distinction between processing levels. For $k = 1$ local surfaces are treated as independent, and data is static. This is not the case at $k = 2$ since, under the imposition of intra-level consistency, extrapolations from each chart serve as data for other nodes. The MAP estimate can only be realised through repeated iteration in which the updated charts after one iteration serve as data for the next iteration (i.e., $w^2 \leftarrow s^2$). The MAP estimate process is a relaxational refinement scheme for the robust propagation of local smoothness constraints. The scheme of Sander and Zucker [11] may be viewed as a special case of the above with: (1) evidential support matrices for intra-level measurements given by $e_{ji}^2 = I, \forall i, j$, at all iterations, and (2) evidential support matrices




Figure 1: Least-means squares (left) and robust techniques applied to front of head.




Figure 2: Least-means squares (left) and robust techniques applied to back of head.

e_{ji}^1 set to the identity matrix at the first update, but to zero thereafter. That is, measurements from the patch level are used to initialise the refinement process but are subsequently ignored. This allows local surfaces to move away from the data. Furthermore, since covariances are ignored the scheme is oversensitive to uncertain data. In addition, all data is implicitly assumed to belong to true local surfaces. This can lead to gross oversmoothing across boundaries.

5 Experiments and Conclusions

Although our primary interest is in interpreting the differential structure of 3D scenes represented in terms of density slice data, in the experiments reported here we have taken range data as our input for reasons of computational and visualisation simplicity. In fact, data for experimentation originated from the estimated locations of 3D feature points returned by a multi-frame feature tracking technique [12] applied to cranial MRI slice data. To convert the data into range form we projected a ray from each point (x,y) on a discrete 2D grid onto the set of 3D feature points, and measured the distance, $z(x,y)$, to the first ‘hit’.

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 raw covariance of range data points to be solely in the sensing direction. All local neighbourhoods were taken over 3×3 nodes (pixels).

We have compared our coupled fitting technique with both a conventional least squares algorithm for surface fitting, and, Sander and Zucker's relaxation scheme [11]. The lefthand pair of images in both Figure 1 and Figure 2 show the results obtained by conventional means. The upper image shows reconstructed surface normals obtained by least squares fitting while the lower image shows the mean and Gaussian curvature labels obtained using the Sander and Zucker algorithm. For comparison, the righthand images show the results obtained using our coupled fitting process. In the case of the conventional methods fine curvature detail is lost and surface structure is blurred. By contrast, in the case of our coupled fitting method, fine surface detail and curvature structure is preserved. This is particularly noticeable in the vicinity of the ears and the bridge of the nose.

Based on these experimental observations, our coupled fitting process appears to provide a powerful framework for extracting consistent differential surface structure from noisy input data without oversmoothing fine detail. The method not only outperforms its conventional counterparts, it also provides an objective framework which integrates the different stages of surface segmentation and interpretation.

References

- [1] Becker S. and Hinton G.E., "Learning Mixture Models of Spatial Coherence", *Neural Computation*, **5**, pp 267-277, 1992.
- [2] Besl P.J. and Jain R.C., "Segmentation through variable order surface fitting", *IEEE PAMI*, **10**, pp 167-192, 1988.
- [3] Blake A. and Zissermann A., "Visual Reconstruction", MIT press, Boston, 1987.
- [4] Ferrie F.P., Lagarde J. and Whaite P., "Darboux frames, snakes, and super-quadratics: geometry from the bottom up", *IEEE PAMI*, **15**, pp 771-784, 1993
- [5] Dempster N.M. and Rubin D.B., "Maximum-likelihood from incomplete data via the EM algorithm", *J. Royal Statistical Soc. Ser. B (methodological)*, **39**, pp 1-38, 1977.
- [6] Flynn P.J., and Jain A.K., "On reliable curvature estimation", *IEEE Computer Vision and Pattern Recognition Conference*, pp. 110-116, 1989.
- [7] Hampel F.R., Ronchetti E.M., Rousseeuw P.J. and Stahel W.A., "Robust Statistics: the approach based on influence functions", J. Wiley and sons, NY, 1986.
- [8] Hancock E.R. and Kittler J., "Discrete relaxation", *Pattern Recognition*, **23**, pp.711-733, 1990.
- [9] Monga O., Ayache N. and Sander P., "From Voxel to Intrinsic Surface Features", *Image and Vision Computing*, **10**, pp 403-413, 1992.
- [10] Monga O., Deriche R., and Malandain G., "Recursive Filtering and edge closing; two primary tools for 3D edge detection", *Image and Vision Computing*, **9**, pp. 203-214, 1991.
- [11] Sander P.T. and Zucker S.W., "Inferring surface structure and differential structure from 3D images", *IEEE PAMI*, **12**, pp 833-854, 1990.
- [12] Sharp N. and Hancock E.R., "Multi-frame feature tracking by probabilistic relaxation", *Proceedings of the Fifth British Machine Vision Conference*, pp 407-418, 1994.
- [13] Terzopoulos D., "Regularisation of inverse problems involving discontinuities", *IEEE PAMI*, **8**, pp 129-139, 1986.
- [14] Turner M. and Hancock E.R., "Bayesian extraction of differential surface structure", to appear in CAIP 95.
- [15] Yuille A. L., "Generalised deformable models, statistical physics and matching problems", *Neural Computation*, **2**, pp 1-24, 1990.