

 Open access • Journal Article • DOI:10.1109/34.57680

Inferring surface trace and differential structure from 3-D images — [Source link](#)

P.T. Sander, Steven W. Zucker

Institutions: McGill University

Published on: 01 Sep 1990 - IEEE Transactions on Pattern Analysis and Machine Intelligence (IEEE Computer Society)

Topics: Principal curvature, Real image, Curvature and Three-dimensional space

Related papers:

- [Differential geometry of curves and surfaces](#)
- [Darboux frames, snakes, and super-quadrics: geometry from the bottom up](#)
- [Snakes : Active Contour Models](#)
- [Describing surfaces](#)
- [Dynamic 3D models with local and global deformations: deformable superquadrics](#)

Share this paper:    

View more about this paper here: <https://typeset.io/papers/inferring-surface-trace-and-differential-structure-from-3-d-3hypsifc7q>

INRIA

UNITE DE RECHERCHE
INRIA-SOPHIA ANTIPOLIS

Institut National
de Recherche
en Informatique
et en Automatique

Domaine de Voluceau
Roddencourt
BP 105
78153 Le Chesnay Cedex
France
Tel. (1) 39 63 55 11

Rapports de Recherche

N° 1117

Programme 6
Robotique, Image et Vision

INFERRING SURFACE TRACE AND DIFFERENTIAL STRUCTURE FROM 3-D IMAGES

Peter T. SANDER
Steven W. ZUCKER

Novembre 1989



Inferring Surface Trace and Differential Structure from 3-D Images

Peter T. Sander
INRIA
2004, Route des Lucioles
06561 Valbonne Cedex, France

Steven W. Zucker*
McRCIM
McGill University, 3480 University St.
Montréal QC, H3A 2A7, Canada

Abstract

Early image understanding seeks to derive analytic representations from image intensities. This paper presents steps towards this goal by considering the inference of surfaces from three-dimensional images. We limit ourselves to smooth surfaces, and focus on the coupled problems of (i) inferring the trace points (the points through which the surface passes), and (ii) estimating the associated differential structure given by the principal curvature and direction fields over the estimated smooth surfaces. Computation of these fields is based on the determination of an atlas of local charts or parametrizations at estimated surface points. Algorithm robustness and the stability of results are essential for analysis of real images; to this end, we present a functional minimization algorithm utilizing overlapping local charts to refine surface points and curvature estimates, and develop an implementation as an iterative constraint satisfaction procedure based on local surface smoothness properties. Examples of the recovery of local structure are presented for synthetic images degraded by noise and for clinical magnetic resonance images.

Détermination de la Structure Différentielle des Surfaces à partir d'une Image Tridimensionnelle

Résumé

Une des premières étapes de l'interprétation d'images numériques est la recherche des représentations analytiques. Cet article présente un algorithme d'estimation des surfaces à partir d'images tridimensionnelles. Nous nous limitons au cas de surfaces C^∞ , et nous privilégions les deux problèmes conjugués de l'estimation de la trace des surfaces (les points par lesquels passent les surfaces) et de leurs propriétés différentielles fournies par les champs de courbure principale et les directions principales. Le calcul de ces champs principaux est effectué en déterminant une carte locale en chaque point de la surface estimée. L'obtention de résultats stables nécessite des techniques robustes. A cet effet, nous avons conçu un algorithme de minimisation qui utilise des cartes superposées pour affiner l'estimation des courbures principales et des directions principales des surfaces. L'algorithme procède par itérations devant satisfaire les contraintes fournies par les propriétés différentielles des surfaces estimées. Des exemples de reconstitution de structures locales sont présentés pour des images synthétiques dégradées par du bruit simulé et pour des images tomographiques obtenues par résonance magnétique nucléaire.

Support for this research was partially provided by NSERC Grant A4470 and MRC Grant MA 6125.

*Fellow, Canadian Institute for Advanced Research

1 Introduction

There is a dilemma inherent in all approaches to inferring the structure of surfaces from three-dimensional images: if the set of points that lie on the surface — the *trace* of the surface — were known, then an analytic mapping defining the surface could be determined. However, since these trace points are not given directly in intensity images, a segmentation problem arises: how to distinguish surface trace points from other image points? In other words, how can the trace of the surface be traced out from the image? This is an important question because non-trace points, or points not on the surface, corrupt estimates of surface properties. The question is analogous to removing ‘outliers’ in curve fitting, and is the question that we shall deal with in this paper for 3-D images. Specifically, our goal is to infer the trace points and to estimate differential geometric properties. We believe that this will provide the proper foundation for recognizing objects and determining shape, and in this paper concentrate on *reliably* estimating principal curvature and direction fields on smooth surfaces. The analysis is taken a step further in [SZ] with the estimation of direction field singularities.

To gain a perspective on 3-D trace segmentation, observe that an analogous problem arises in 2-D images: before an edge or curve can be traced out, the edge points must somehow be segmented from the other, non-edge, image points [PZ89]. The analogy with 2-D image understanding is tight. The boundary of a two-dimensional object is a curve separating inside from outside. In a sampled image, an estimate of the curve might be represented by a sequence of pixels each with an associated unit vector, an estimate of the local tangent direction of the curve. The boundaries of 3-D objects are surfaces estimated locally by ‘voxels’ with associated unit surface normal vectors. Initial estimates of boundary points can be obtained by convolving the image with a digital gradient operator. However, as in 2-D edge detection, it is difficult to directly infer global boundaries from these raw local measurements. Intermediate processing seems to be required, and it is this intermediate stage with which we shall be concerned.

A brief remark concerning terminology: we prefer to describe our goal as ‘surface trace point inference’ rather than the more common phrase (in the computer vision community) ‘edge point detection’ for the following reason. Edge point detection generally means running some gradient (or differential) operator and then thresholding the response. But this is just the first step in determining trace points of underlying object surfaces. Trace points (the image points lying on the surface) are rooted in the mathematical idea of explicit surfaces and they carry the implication of relations between themselves, something which is markedly absent from the independent gradient operator measurements. These relations provide the foundation for our intermediate stage of analysis.

Our approach to inferring surface (trace) points in 3-D images is based on differential topology and geometry. Limiting ourselves to smooth surfaces, the local structural information relevant to shape is embodied in the principal curvature, normal, and principal direction fields over a surface. Theoretically, we pose the determination of these fields as the covering of the putative surface by local parametrized surface patches, or charts (the requisite technical background and terminology is provided in §4). But translating the differential theory into an implementation is complicated by the ‘chicken and egg’ nature of the problem: the trace of the surface is necessary to determine structural properties of the surface, but the trace is not directly available and must itself be determined from the emerging surface structure. In short, differential geometry assumes the existence of the surface, but the surface trace exists only implicitly in the image. We break the ‘chicken and egg’ cycle by computing the trace and differential structure simultaneously, using each to ‘refine’ the other, as opposed to taking the thresholded gradient operator results as sacrosanct.

Concern with stability and robustness issues leads naturally to variational methods. Thus, we introduce a functional at each trace point, whose minimization determines an estimate of the surface differential structure at that point. We implement the minimization as an iterative constraint satisfaction procedure, with constraints drawn from consistency relations between overlapping charts at putative surface trace points. While this framework is similar to ‘standard’ relaxation labeling procedures [HZ83], there are significant differences. In particular, the combinatorial load of dealing with 3-D data obliges keeping the labels implicit and in the continuous domain.

Shape characterization and decomposition [HR84,KvD82] provide the overall context in which to place this paper. We believe that the sorts of descriptions that will underlie successful theories of shape will be inspired by notions from differential topology, i.e., they will be stable in that small changes in the image (whether due to noise in the image or small changes in the underlying physical surfaces themselves) should lead to small changes in the surface descriptions. The direction fields that we compute are stable in this sense, and furthermore they permit the successful calculation of generic singularities such as umbilic points [SZ]. The principal curvature and direction fields thus provide a qualitative foundation of precisely the sort on which ‘third paradigm’ computational vision stands [Zuc87].

We consider 3-D imagery for two reasons, one theoretical and the other practical. First, theoretically, true three-dimensional images¹ avoid the ambiguity introduced by projection from a 3-D world to a 2-D image. However, many of the fundamental problems of image understanding remain (boundary point detection, organization of points into object boundaries, object determination and identification [RK81,BB82,Mar82,Lev85]), and are complicated by the additional dimension. In this paper, we concentrate on the earliest stages of recovering object surfaces from 3-D images, specifically, the estimation of trace points. Such estimation procedures indicate essential characteristics for generic part decomposition algorithms. The second reason for considering 3-D imagery is that scanners for non-invasive imaging are becoming practical and widespread. In biomedicine, for example, true 3-D images are produced by magnetic resonance imaging (MRI) [Pyk82], computed tomography [GHJ75], and positron emission tomography [TRS80]. But, for diagnostic purposes, they are treated as a stack of 2-D images (see Fig. 1), and it is left to clinicians to infer in their “mind’s eye” the spatial layout of internal structure from a sequence of images viewed temporally. This is unnatural for people, however, making many diagnoses inefficient or, in some cases, almost impossible. Representations of 3-D surfaces could alleviate these problems by permitting the use of computer graphics display techniques. And, perhaps more importantly, if the 3-D surface structure were available, information from the different types of scanners could be integrated into a single presentation. It is now widely held that such presentations would be a significant advance over the individual 2-D slice sequences.

We should also point out that, in this paper, we are only interested in the geometric properties of 3-D images. Indeed, there is information in the absolute values of the image intensities which may help in the particular image understanding task, but it is the *relative* values on which the geometry is based. Thus we do not use any knowledge specific to the domain of MR images, e.g., that certain tissue types correspond to certain intensity ranges or that certain structures are likely to be found at specific image locations. As well, we only deal with estimating the trace points and their local structure, and using this information for a coarse segmentation of the surfaces. Developing such a segmentation into analytic surface patches which can support a theory of parts is an immediate research direction suggested by this work.

¹Where the image value at each point is a function of three spatial coordinates. We generalize ‘intensity’ to mean any function $I(x, y, z)$.




Figure 1: A sequence of two-dimensional slices from a magnetic resonance image (MRI) of the head. Slice (a) is at mid-nose level, and the slices proceed toward the top of the head. The darker grey structures visible in the centre of slices (f) through (i) are the lateral ventricles. The images are 256×256 pixels in the xy plane, with a resolution of 0.2 cm. in the z direction. Here we show 9 images from a sequence of 41. (Courtesy Dr. Terry Peters of the Montreal Neurological Institute.)

In the following section, we outline other research related to that presented in this paper, and §3 is an overview of our methods. The necessary mathematical background is in §4. Sections 5 and 6 are the core of the paper, where we present our minimization methods for finding the principal curvatures and directions, and give details of our implementation as iterative constraint

satisfaction: §5 applies the differential geometry to estimate initial local properties of possible surface points; and in §6 the local information is refined by making it more consistent (according to certain smoothness properties) with that in its local context. We present examples in §7 of the application of our methods to synthetic imagery and to clinical magnetic resonance images, which demonstrate clearly the necessity and the power of the refinement methods for real imagery. Section 8 is optional to the understanding of the paper; it briefly introduces notions from more advanced differential geometry which, in fact, inspired our approach.

2 Related work

Related work falls into two broad categories: (i) processing of true 3-D biomedical images; and (ii) processing of rangefinder images. Researchers working on (i) begin, of course, with the same type of data as we do, but, as we explain below, our methods are much closer in spirit to those of (ii). For example, Herman [HL79] has concentrated on the 3-D display of organ surfaces rather than on their analysis. Such implementations often require a user to interactively indicate a starting point on an object surface, after which neighbouring points are grouped by connecting voxels with similar intensity levels. Their goal is to allow a physician to perform interactive 'digital surgery' on the image (see also [VMW83,RGA83,Tri86]).

Much of the work on 3-D images treats them as a stack of 2-D images and can be summarized as follows: first, do edge detection on each of the 2-D slices, then interpolate between the contours of each slice to capture the 3-D structure. This approach can be found in, among others, [FKU77, AFH81,Boi84,BTTU84]. We believe that it is important to treat a 3-D problem in 3-D — that there is useful information in the constraints *between* the slices (see [Cap87] for a comparison of two- and three-dimensional methods; see also [LC87] for a gradient based approach which does use inter-slice connectivity).

Bajcsy and her collaborators are concerned with matching clinical data into an idealized model image in order to identify internal structures. They pose the problem as that of elastically deforming a 'rubber sheet' image until it matches pictures in the model image atlas [BLR83]. This method, essentially based on the correlation between intensity levels of the voxels in the image and the image atlas, is dependent on a highly constrained domain. The method does not deal with domain independent segmentation of the image itself into possible objects.

As far as we are aware, this is the only research which analyzes 3-D biomedical images by making explicit the domain independent differential geometric properties.

It is worth stating explicitly the distinction between our 3-D images and idealized '3-D' rangefinder images. A true 3-D image is a function $I : \mathbb{R}^3 \rightarrow \mathbb{R}$ (appropriately sampled and quantized), whereas rangefinders give $I' : \mathbb{R}^2 \rightarrow \mathbb{R}$ which can be considered as $I' : \mathbb{R}^3 \rightarrow \{0,1\}$ (indicating the absence/presence respectively of a surface trace point). There is no 'surface detection problem' as such in rangefinder images. Techniques for analyzing rangefinder (or depth) images are well represented in the literature, e.g., [BJ85,BPYA85,FH86]. Our work and the recent work on range images share common roots in differential geometry, but this similarity can be somewhat misleading. Range images should really be considered '2 1/2-D' images since the information as to which points in space are points of some surface is explicitly available (in the ideal case). The trace points we seek are, for them, given *a priori*: it is only once these points are estimated that we are able to bring the differential geometric machinery to bear on the problem. Thus, for example, when [HW81,PB85] fit local surfaces of varying degree, they presume to know which points of \mathbb{R}^3 to use as data points; we must first segment the data points from non-surface image points. It follows, of course, that for real rangefinder images (as opposed to idealized ones) techniques like ours can be used to segment

the ‘noise’ points from the data points, and hence should be useful in these domains as well; indeed, preliminary research by Ferrie *et al* [FLW89] extending our techniques to real rangefinder configurations is very promising.

Once local analytic surface models are instantiated, derivative information is readily available by differentiation. Mathematicians naturally classify points on a surface by the signature of the Gaussian and mean curvatures into elliptic (concave or convex), hyperbolic (saddle), parabolic, and planar type; this has become a paradigm of researchers in range image analysis as well, see [HWL83,BJ86,YL87], cf., [HR84]. However, issues of stability have not been properly appreciated. For example, parabolic lines separating elliptic and hyperbolic regions are theoretically stable [BGM82], but observe that, in practice, small errors in local models could lead to large shifts in the classification, resulting in large changes in the parabolic lines and little stability. As well, the theoretical importance of the extrema of the normal curvatures, that is, the principal curvatures and their associated directions, has been recognized [BJ85,BPYA85,FH86], but their computation has been considered to be unreliable due to instability in the presence of noise. We now show how these quantities — principal curvatures, and principal directions — can be reliably computed.


3 Overview

Our specific aim is to infer from 3-D images sets of points which may be reliably considered as the points of a smooth surface, and to make the differential geometric properties of the surface explicit at those points. An initial estimate of the surface is built up by constructing local surface models based on the results of a 3-D gradient operator. Subsequently, we use certain technical smoothness properties of the local surfaces to refine the estimated surface points. The need for this refinement is illustrated by the example of Fig. 2, in which geometric properties computed from initially-estimated local surfaces are compared with refined properties. As can be seen, these initial local estimates are generally unsatisfactory. The subsequent non-local processing is posed as a minimization problem subjecting the initial estimates to compatibility constraints imposed by the estimated information at neighbouring points.


The compatibility constraints are derived as follows. Let Γ_P be an estimate of the differential structure at P (termed the *chart* at P ; all the necessary technical background is supplied in §4), consisting of the principal curvatures and directions there, and let Γ_α , $\alpha = 1, \dots, k$, be the charts at neighbouring surface points Q_α . A functional r_P is defined for the charts at Q_α , which measures how compatible they are with the chart at P , given a putative (local) surface model, and this leads to a global functional $\Phi = \sum_P r_P$, the natural measure of compliance over all estimated trace points. The functional Φ is minimized by iteratively replacing each estimated chart by a ‘best fit’ (in the least squares sense) of those in its neighbourhood. The key to making it work is the manner in which charts at neighbouring points are compared by using overlapping parametrizations. The process is iterative, with each step providing a better, in the sense of locally more consistent, chart at each surface point.

Finally, at convergence, the computed local information admits two complementary forms of grouping: (i) integration of the principal directions leads to lines of curvature; and (ii) grouping of points according to Gaussian and mean curvature type leads to patches of points with qualitatively similar curvature, see Fig. 3(d). The latter can be used to effect a coarse partition of the surfaces.


The initialization step is special. Initial surface points are proposed where a 3-D gradient operator [ZH81] convolved over the intensities of the input image responds strongly, i.e., suggests the existence of a surface tangent plane. Since it doesn’t figure directly in the constraint satisfaction itself, details of the gradient operator are left to App. A. Figures 3(a), (b), (c) show the local tangent




(a)



(b)



(c)



(d)

Figure 2: The frontal view of the face region of the MR image. (a) The orientation of the tangent planes is shown by considering a light source from the right. (b) Normal directions of the tangent planes. (c) Initial estimates of the principal direction field corresponding to the smaller principal curvature, based on the local information from the gradient operator using the methods of §5. (d) Refined estimate of the direction field after five iterations, illustrating contextual constraints improve the initial estimates. The algorithm for refining the initial estimates is in §6. Bear in mind that although the fields are shown projected onto a viewing plane, they exist on a curved surface in space.

planes computed by the gradient operator for synthetic and MR images. The initialization also

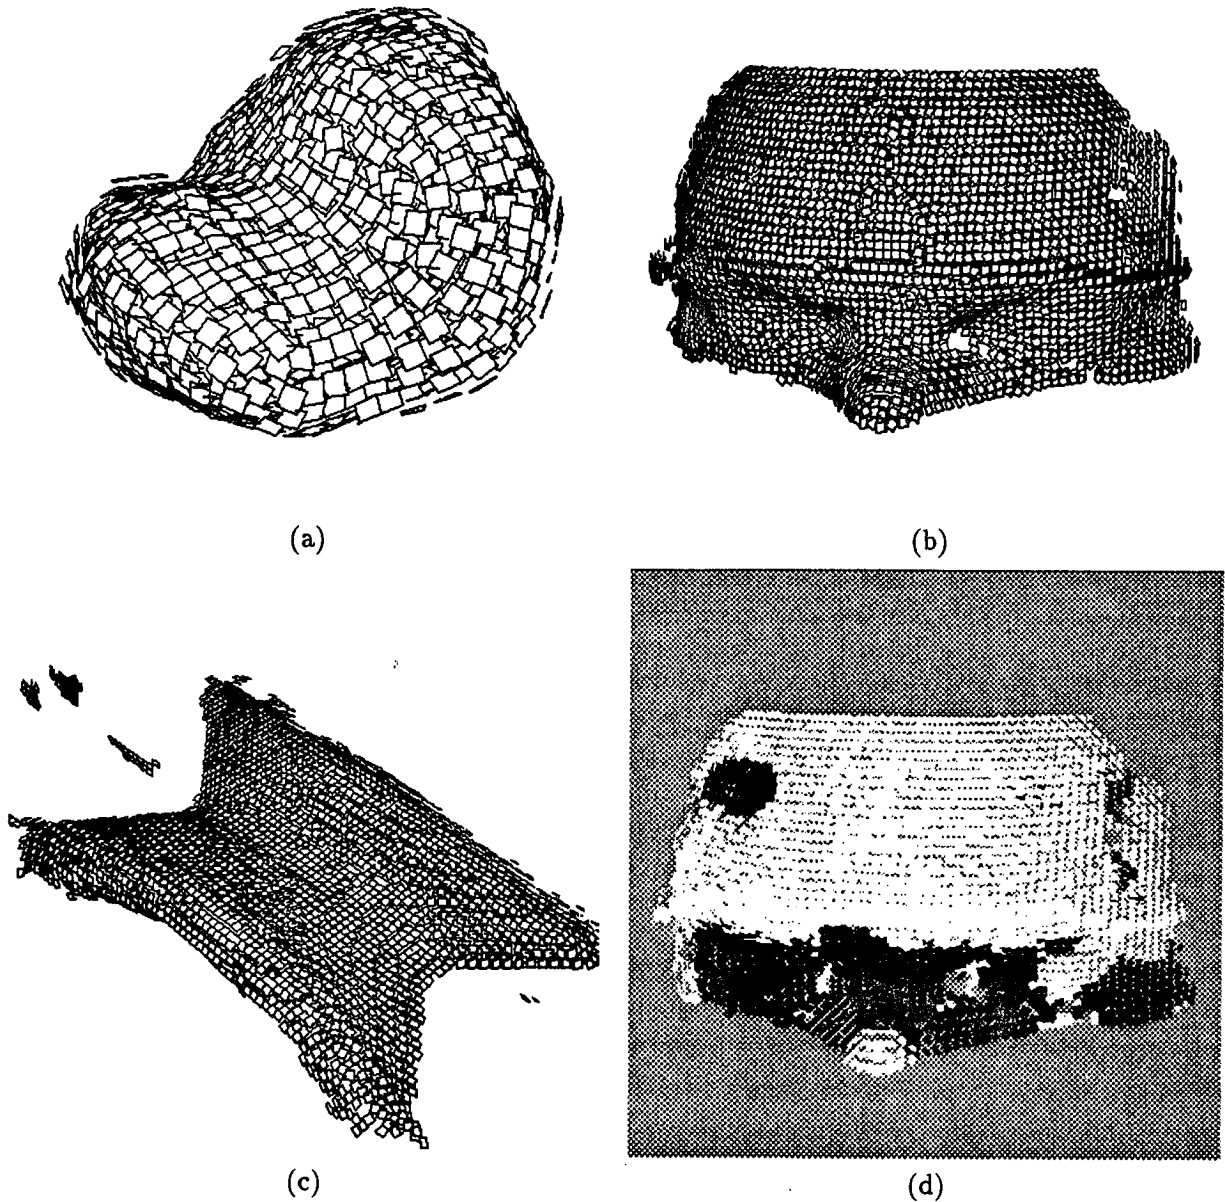


Figure 3: Local tangent planes of (a) a synthetic 'bloblike' 3-D image, (b) the face view of the magnetic resonance image, and (c) the top view of the lateral ventricles. (d) Segmentation of the face image into regions of elliptic (bright) and hyperbolic (dark) points. Weak convolution operator responses have been eliminated. The synthetic image is $32 \times 32 \times 32$, the face image is $90 \times 48 \times 41$, and the ventricle image is $70 \times 96 \times 27$.

includes the estimation of local surfaces at the proposed surface points. We pose the estimation as the instantiation of a quadric surface patch at each candidate surface point by means of a nonlinear 'best fit' (in a sense made precise in §5) over neighbourhoods of surface points.

Singularities of the principal direction fields are handled quite naturally, although some special care must be taken in the application. We focus on these in [SZ]; for now one should simply note that (i) the presence of singularities affects our iteration, but (ii) the resultant direction fields are sufficiently stable to serve as a basis for calculating them.

4 Mathematical background and terminology

Let $S \subset \mathbb{R}^3$ be a given set of points with an open subset $\mathcal{O} \subset S$. A *(local) parametrization* is determined by $\phi : U \rightarrow \mathcal{O}$, where ϕ is a diffeomorphism (recall that a diffeomorphism is a smooth bijective map with a smooth inverse) and $U \subset \mathbb{R}^2$ is an open set. The inverse mapping ϕ^{-1} is a *(local) coordinate system* or *chart*. An *atlas* is a collection of local charts which cover S subject to the condition that overlapping charts be smoothly related, i.e., when $P \in \mathcal{O}_1 \cap \mathcal{O}_2$ then the *change of coordinate* mapping

$$\phi_2^{-1} \circ \phi_1 : \phi_1^{-1}(\mathcal{O}_1 \cap \mathcal{O}_2) \rightarrow \phi_2^{-1}(\mathcal{O}_1 \cap \mathcal{O}_2)$$

is smooth with a smooth inverse. The set S equipped with the additional structure provided by such an atlas is our working definition of a *surface*. Our refinement methods of §6 are based in an essential way on the change of coordinate mappings at overlapping parametrizations.

We will generally speak interchangeably of parametrizations or charts (since they are smoothly invertible). We depart slightly from the usual language of differential topology [GP74] by explicitly including a set of right-handed orthonormal axes in the notion of local chart, thus writing (ϕ, U, ξ) , where $\xi = (\vec{P}, \vec{Q}, \vec{N})$, $\vec{P}, \vec{Q}, \vec{N} \in \mathbb{R}^3$, and $U \subset \text{span}(\vec{P}, \vec{Q})$. The vector \vec{N} is always the normal of the surface at P , and (\vec{P}, \vec{Q}) span the tangent plane; thus ξ is a *Darboux trihedron* or *frame* [dC76] of S at P , see Fig. 4. Note that, as defined, the tangent plane basis (\vec{P}, \vec{Q}) is arbitrary up to orthonormality constraints; when \vec{P}, \vec{Q} are the principal directions of the surface, then ξ is the *principal direction frame* at P . As we shall see below in §5, working in the *principal chart* allows us to choose a particularly simple form for the mapping ϕ .

We assume that the reader is familiar with the differential geometric notions of principal curvatures and directions of a surface, and with Gaussian and mean curvature (do Carmo [dC76] is an excellent reference). The Gaussian curvature K and the mean curvature H naturally partition the set of surface points into four basic types: *elliptic* ($K > 0$), *hyperbolic* ($K < 0$), *parabolic* ($K = 0$ but $H \neq 0$), and *planar* ($K = H = 0$), see Fig. 5. In effect, they provide a qualitative basis for characterizing smooth surface patches: regions composed entirely of elliptic points are necessarily separated from hyperbolic regions by a curve of parabolic points. Parabolic lines are stable since, under small perturbations of the surface, elliptic and hyperbolic regions maintain their respective characters [BGM82].

5 Initial local estimates

In this section, we compute an initial estimate of the atlas of the candidate surface S . The local parametrizations comprising the atlas are modeled by *parabolic quadric* surface patches, and are instantiated at the points of S via surface fitting by least squares error minimization. Principal curvatures and directions are computed from the estimated local surface parameters.

5.1 Local surface parametrization model

To build an atlas of S , we need to specify at each $P \in S$ a parametrization (or chart) (ϕ, U, ξ) , where $\phi : U \rightarrow \mathcal{O}$, $P \in \mathcal{O} \subset S$, and ξ is the coordinate frame.




Figure 4: The Darboux frame $(\vec{P}, \vec{Q}, \vec{N})$ at P on any oriented curve C along the oriented surface S . \vec{P} is tangent to C , \vec{N} is normal to S , and \vec{Q} is orthonormal to both \vec{P} and \vec{N} , with $\det(\vec{P}, \vec{Q}, \vec{N}) = +1$.

We are interested in differential surface structure up to and including curvature, and the simplest surface sufficient for the job is a *parabolic (or non-central) quadric* surface [Eis60]. It is the Taylor expansion of bivariate function f truncated after the second order term, with graph

$$\phi(p, q) = (p, q, n(p, q)); \quad n(p, q) = \frac{1}{2} (ap^2 + 2bpq + cq^2), \quad (1)$$

where $a = \frac{\partial^2 f}{\partial p^2}$, $b = \frac{\partial^2 f}{\partial p \partial q}$, $c = \frac{\partial^2 f}{\partial q^2}$. In theory, a local coordinate system can always be found in which this form is appropriate. In our application, the surface normal \vec{N} produced at P by the gradient operator allows us to build just such a parametrization, where $P = \phi(0, 0) = (0, 0, 0)$ is a critical value of ϕ and $U = \phi^{-1}(\mathcal{O})$ for $\mathcal{O} = V \cap S$, $V \subset \mathbb{R}^3$ an open-ball neighbourhood of P . Let $\vec{P}, \vec{Q} \in \mathbb{R}^3$ be any two vectors such that $\xi = (\vec{P}, \vec{Q}, \vec{N})$ forms a frame at P , i.e., (\vec{P}, \vec{Q}) forms a basis for the tangent plane of S at P (the arbitrary choice of the tangent plane basis vectors \vec{P}, \vec{Q} in frame ξ has no effect on the curvature determination, as we will see below). We will consider Eq. (1) in this coordinate system, i.e., p is along \vec{P} , q is along \vec{Q} , and n is along \vec{N} , to be the local surface model for the parametrizations of the atlas.




Figure 5: Surface point P is classified as (a) elliptic, (b) hyperbolic, or (c) parabolic according to its Gaussian and mean curvatures.

5.2 Instantiating the local model

Instantiation of the local parabolic quadric model at candidate trace point P means determining the parameters a, b, c of Eq. (1) appropriate to the neighbourhood \mathcal{O} and the frame ξ . This could be done, in principle, by directly computing second partial derivatives over the image intensities. However, the sensitivity of differentiation to noise (increasing for higher orders) suggests a more robust surface fitting approach.

We take the candidate surface points given by the thresholded gradient operator as samples of the surface \mathcal{S} . Thus the local parametrization Eq. (1) is to be inferred from the estimated surface points $Q_i = (p_i, q_i, n_i) \in \mathcal{O}$, $i = 1, \dots, k$ of the neighbourhood of P . We assume for the moment that the Q_i are in the (ϕ, U, ξ) coordinate system local to P ; we will come back to this shortly. A straightforward method of determining the quadric parameters $\mathbf{x} = (a, b, c)$ is by a linear least-squares error minimization of the simultaneous linear equations $A\mathbf{x} = \mathbf{b}$

$$\frac{1}{2} \begin{pmatrix} p_1^2 & 2p_1q_1 & q_1^2 \\ \vdots & \vdots & \vdots \\ p_k^2 & 2p_kq_k & q_k^2 \end{pmatrix} \begin{pmatrix} a \\ b \\ c \end{pmatrix} = \begin{pmatrix} n_1 \\ \vdots \\ n_k \end{pmatrix},$$

since by Eq. (1) each data point $Q_i = (p_i, q_i, n_i)$ should ideally satisfy

$$E_1(a, b, c, p_i, q_i, n_i) = \frac{1}{2} (ap_i^2 + 2bp_iq_i + cq_i^2) - n_i = 0. \quad (2)$$

This method of fitting a surface at P to the data points has the desirable property of supplying information as to the quality of the fit (see [SZ86] for use of the residual and condition number). However, it does not make use of all the available information since only the positions of the points are used—the associated surface normal information is ignored. Unfortunately, the linear least-squares formulation does not lend itself directly to the incorporation of the surface normals for the following reason. Only the *directions* of the ‘model’ normal $\vec{N}(p, q)$ (derived from the surface model ϕ in Eq. (1) — see below) and the ‘data’ normal (N_{pi}, N_{qi}, N_{ni}) at Q_i (derived by the gradient operator from the contrast across image intensity values) can be meaningfully related in any equation—the magnitudes arise from different processes and must be normalized away since they are incommensurable. The model normal is determined as

$$\vec{N}(p, q) = \frac{\partial \phi}{\partial p} \times \frac{\partial \phi}{\partial q} = (-ap - bq, -bp - cq, 1),$$

and thus normal information adds the three equations

$$\begin{aligned} E_2(a, b, c, p_i, q_i, n_i) &= \frac{ap_i + bq_i}{\|\vec{N}\|} + \frac{N_{pi}}{\|(N_{pi}, N_{qi}, N_{ni})\|} = 0, \\ E_3(a, b, c, p_i, q_i, n_i) &= \frac{bp_i + cq_i}{\|\vec{N}\|} + \frac{N_{qi}}{\|(N_{pi}, N_{qi}, N_{ni})\|} = 0, \\ E_4(a, b, c, p_i, q_i, n_i) &= \frac{1}{\|\vec{N}\|} - \frac{N_{ni}}{\|(N_{pi}, N_{qi}, N_{ni})\|} = 0 \end{aligned} \quad (3)$$

to Eq. (2), for each point Q_i . The resulting system is, unfortunately, nonlinear in the parameters a, b, c , since the normal $\vec{N}(p, q)$ has magnitude

$$\|\vec{N}(p, q)\| = \sqrt{(a^2 + b^2)p^2 + 2(ab + bc)pq + (b^2 + c^2)q^2 + 1}.$$

Now, for each Q_i the surface mapping ϕ is presumed to satisfy the four² Eqs. (2, 3). The parameters a, b, c of this overconstrained system of simultaneous nonlinear equations can be determined by minimizing some function of the E_j , for example,

$$\min_{a, b, c} \sum_{i=1}^k \left[v_i \sum_{j=1}^4 \omega_j (E_j(a, b, c, p_i, q_i, n_i))^2 \right],$$

where the v_i weight the data points and the ω_j are weighting factors balancing position (E_1) and normal ($E_2 - E_4$) information³. We employ a quasi-Newton method for the minimization⁴.

Returning to the question of coordinate systems, we see that the model Eq. (1) is set up in the tangent plane parametrization (ϕ, U, ξ) whereas the ‘raw’ data points Q_i are only directly available in \mathbb{R}^3 image coordinates. To carry out the fit of the model to the data, we determine

²The unit surface normal, as a point on the unit sphere by the Gauss map, is determined by two parameters, but all three coordinates of $\vec{N}(u, v)$ are needed to compute those two parameters. For computational simplicity, we use the three coordinate equations directly in specifying the normal, rather than determining these two parameters and using only two normal component equations.

³We use weights $v_i = 1, \omega_1 = 1/2, \omega_2 = \omega_3 = \omega_4 = 1/6$ throughout the examples.

⁴IMSL [Ric] routine ZXMIN.

the transformation mapping a point in the neighbourhood of P from P 's tangent plane coordinate system to \mathbb{R}^3 image coordinates. This is done by means of the commutative diagram

$$\begin{array}{ccc}
 & \xrightarrow{(p, q, n(p, q))} & \\
 U & & U \times \vec{N} \\
 \uparrow \phi^{-1} & & \downarrow T \circ R \\
 \mathcal{O} & \xrightarrow{(x, y, z)} & \mathbb{R}^3
 \end{array}$$

where $T \circ R$ is a transformation consisting of a rotation followed by a translation which align P 's frame ξ with the \mathbb{R}^3 image axes at the origin, and (x, y, z) just indicates assignment of the point's \mathbb{R}^3 coordinates in the image.

In coordinates, let P , the point at which the surface ϕ is being fit, have \mathbb{R}^3 image coordinates (x_P, y_P, z_P) , and neighbouring point $Q \in \mathcal{O}$ have image coordinates (x, y, z) . Then the transformation of Q to the tangent plane coordinates centered at P is given by

$$R^{-1} \circ T^{-1} \begin{pmatrix} x \\ y \\ z \end{pmatrix} = \begin{pmatrix} P_x & P_y & P_z \\ Q_x & Q_y & Q_z \\ N_x & N_y & N_z \end{pmatrix} \begin{pmatrix} x - x_P \\ y - y_P \\ z - z_P \end{pmatrix}, \quad (4)$$

where $\vec{P} = (P_x, P_y, P_z)$, $\vec{Q} = (Q_x, Q_y, Q_z)$, $\vec{N} = (N_x, N_y, N_z)$ are the image coordinates of the frame ξ .

Now, substituting the coordinates from Eq. (4) into Eqs. (2, 3) gives us the form of the equations for the minimization, with everything in image coordinates. Writing $\vec{W}(x_i, y_i, z_i) = (x_i - x_P, y_i - y_P, z_i - z_P)$ we obtain

$$\begin{aligned}
 E_1(a, b, c, x_i, y_i, z_i) &= \frac{1}{2} (a\langle \vec{P}, \vec{W} \rangle^2 + 2b\langle \vec{P}, \vec{W} \rangle \langle \vec{Q}, \vec{W} \rangle + c\langle \vec{Q}, \vec{W} \rangle^2) - \langle \vec{N}, \vec{W} \rangle = 0 \\
 E_2(a, b, c, x_i, y_i, z_i) &= \frac{a\langle \vec{P}, \vec{W} \rangle + b\langle \vec{Q}, \vec{W} \rangle}{\|\langle \vec{N}, \vec{W} \rangle\|} + \frac{N_{xi}}{\|(N_{xi}, N_{yi}, N_{zi})\|} = 0, \\
 E_3(a, b, c, x_i, y_i, z_i) &= \frac{b\langle \vec{P}, \vec{W} \rangle + c\langle \vec{Q}, \vec{W} \rangle}{\|\langle \vec{N}, \vec{W} \rangle\|} + \frac{N_{yi}}{\|(N_{xi}, N_{yi}, N_{zi})\|} = 0, \\
 E_4(a, b, c) &= \frac{1}{\|\langle \vec{N}, \vec{W} \rangle\|} - \frac{N_{zi}}{\|(N_{xi}, N_{yi}, N_{zi})\|} = 0
 \end{aligned}$$

($\langle \bullet, \bullet \rangle$ indicates inner product). The inner products, of course, need only be computed once at each Q_α , before the minimization routine.

5.3 Estimating curvature at $P = \phi(0, 0) = (0, 0, 0)$

Curvature information at $P \in \mathcal{S}$ can be computed from the local chart $(\phi, U; \xi)$ instantiated at P as in the previous section. The principal curvatures and directions are determined from the Hessian

matrix of the parametrization about P , which is

$$G|_{(0,0)} = \left(\begin{array}{cc} \frac{\partial^2 n}{\partial p^2} & \frac{\partial^2 n}{\partial p \partial q} \\ \frac{\partial^2 n}{\partial q \partial p} & \frac{\partial^2 n}{\partial q^2} \end{array} \right) \Big|_{(0,0)} = \begin{pmatrix} a & b \\ b & c \end{pmatrix}$$

in the tangent plane parametrization of Eq. (1). The eigenvalues $\kappa_1 \geq \kappa_2$ of G are the principal curvatures at P

$$\begin{aligned} \kappa_1 &= \frac{1}{2} \left(a + c + \sqrt{(a-c)^2 + 4b^2} \right), \\ \kappa_2 &= \frac{1}{2} \left(a + c - \sqrt{(a-c)^2 + 4b^2} \right), \end{aligned}$$

which determine the Gaussian and mean curvatures

$$\begin{aligned} K &= \kappa_1 \kappa_2 = ac - b^2, \\ H &= \frac{(\kappa_1 + \kappa_2)}{2} = \frac{(a+c)}{2} \end{aligned}$$

respectively. At non-umbilic P (where the principal curvatures are distinct) the corresponding eigenvectors

$$\begin{aligned} \vec{p}_1 &= \left(b, \frac{1}{2} \left(c - a + \sqrt{(a-c)^2 + 4b^2} \right) \right), \\ \vec{p}_2 &= \left(b, \frac{1}{2} \left(c - a - \sqrt{(a-c)^2 + 4b^2} \right) \right), \end{aligned}$$

after normalization, are the principal directions.

Note that the eigenvectors as given are in the so far arbitrarily chosen (\vec{P}, \vec{Q}) basis of the tangent plane of the frame ξ . Note also that the simple form of these curvature computations is due to the fact that the Gauss frame of the chart (ϕ, U, ξ) coincides with the (\vec{P}, \vec{Q}) tangent plane basis of ξ at the local origin $P = \phi(0, 0) = (0, 0, 0)$, i.e.,

$$\frac{\partial \phi}{\partial p} \Big|_{(0,0)} = (1, 0, 0), \quad \frac{\partial \phi}{\partial q} \Big|_{(0,0)} = (0, 1, 0).$$

Compare this with the curvature computations away from the origin, as given in App. B.

The only loose end left to clear up is the conversion of the principal directions

$$\vec{x}_j = \frac{\vec{p}_j}{\|\vec{p}_j\|}, \quad j = 1, 2$$

from the (\vec{P}, \vec{Q}) system into \mathbb{R}^3 image coordinates. This is straightforward by the rotation

$$R \vec{x}_j = \begin{pmatrix} P_x & P_y & P_z \\ Q_x & Q_y & Q_z \\ N_x & N_y & N_z \end{pmatrix} \begin{pmatrix} x_{j1} \\ x_{j2} \\ 0 \end{pmatrix}, \quad j = 1, 2.$$

Henceforth, we use the following labeling

$$\begin{aligned} \vec{M} &= R \vec{x}_1 \\ \vec{m} &= R \vec{x}_2 \end{aligned}$$

to refer to the principal directions corresponding to the greater and lesser principal curvatures, κ_M, κ_m respectively, in image coordinates. In this *principal direction coordinate system*, that is, with the arbitrarily chosen tangent plane spanning vectors (\vec{P}, \vec{Q}) replaced by the principal directions (\vec{M}, \vec{m}) , the expression of the local surface becomes particularly simple: $\phi(u, v) = \left(u, v, \frac{1}{2}(\kappa_M u^2 + \kappa_m v^2) \right)$. The Gaussian curvature is $K = \kappa_M \kappa_m$ at P and the mean curvature is $H = (\kappa_M + \kappa_m)/2$.

5.4 Implementation considerations

The residual and condition number of the fit can, of course, still be computed to give an idea of the quality and stability of the fit. Incorporation of the surface normal information permits a better estimated local surface fit over a given neighbourhood size. We have observed empirically that using this additional information has about the same effect for attenuating noise as using purely positional data over a larger neighbourhood (5 pixels diameter vs. 7 pixels), but without so much of the accompanying danger of over-smoothing the data. The fidelity of the local surface model to a presumed underlying smooth surface inherently decreases with larger neighbourhoods, see [PS78, p.44] for an interesting illustration. For example, in cases where poor estimates are obtained from a 5 pixel diameter neighbourhood using position only, better results are achieved in practice either by using the normals over the same neighbourhood or by using positional information over a diameter 7 neighbourhood. The results are not critically sensitive to tuning the neighbourhood size parameter, in fact all the examples in this paper were generated with a 7 pixel diameter spherical neighbourhood.

Now that the estimated local structure is explicit, there is enough information to permit a coarse segmentation of the putative trace points by their Gaussian and mean curvatures. Unfortunately as can be seen from the examples in §7, any such grouping at this stage will fare poorly since the information is not stable with respect to noise. The elliptic and hyperbolic regions are generally small, hence likely to change character enormously with small changes to the surface points. It is generally conceded that noise has a deleterious influence on these results: [FH86] mentions the difficulties of reliably computing principal directions in the presence of noise. We now show how to do just that.

6 Refining estimated local information

In the previous section, we made use of derivative computations over an estimated atlas of local surface charts, employing lower order estimates to determine higher order properties (i.e., tangency from position, curvature from tangency). Figure 2 and §7 present examples showing that these, by themselves, are incapable of producing satisfactory differential structure estimates in realistic situations. Here, we present a remedy in the form of a ‘contextual peer review’ whereby information of a given order is refined by making it more consistent with information of the *same* order over local neighbourhoods. This leads in a more formal manner to a functional minimization.

As the paradigm for our refinement methods, we take the relaxation labeling techniques [HZ83] used for extracting structure in 2-D images. However, the far greater number of combinatorial possibilities inherent in 3-D precludes using explicit discrete labels. Instead, we develop an iterative constraint satisfaction method in the continuous domain, within the differential geometric formalism. At each surface point at each iteration step, a new chart is computed which best matches the information in the neighbourhood of the point. The key is in the notion of support between surface points — we develop a measure of consistency based on smooth transformations between overlapping charts of a smooth surface and update estimates to improve this consistency. This updating of charts leads to the determination of a new and better estimated atlas for the surface.

As in relaxation labeling, the fundamental idea is that of minimizing an appropriate functional over the space of structures of interest. We will develop a functional $\Phi : \{\mathcal{A}\} \rightarrow \mathbb{R}$ over the space of atlases \mathcal{A} , and a rule $\Gamma^i \mapsto \Gamma^{i+1}$; $\Gamma^i, \Gamma^{i+1} \in \mathcal{A}$ ($\Gamma = (\phi, U, \xi)$) for updating the principal chart at each surface point at each iteration i , such that $0 \leq \Phi(\mathcal{A}^{i+1}) \leq \Phi(\mathcal{A}^i)$. We specify Φ by presenting in detail the smoothness constraints imposed on the (principal) chart at P by overlapping charts

at neighbouring points, and develop the $(i+1)^{\text{th}}$ iteration of the algorithm by explicating the local rule taking the chart Γ_P^i to Γ_P^{i+1} . In particular, we examine these issues:

1. which neighbouring points should contribute to the support of the chart at P , i.e., which points determine its *contextual neighbourhood* \mathcal{N}_P ;
2. how to determine the support of the information at point $Q_\alpha \in \mathcal{N}_P$ for the chart at P , i.e., formulation of a mapping $\Gamma_\alpha \mapsto \Gamma_{P_\alpha}$ taking the chart at Q_α to a new chart at P , which can then be compared with Γ_P ; and
3. given 1 and 2, how to update the information⁵ at P by mapping $\Gamma_P^i \mapsto \Gamma_P^{i+1}$.

Throughout this section we adhere to the convention that P^i refers to the particular point at which the information is being refined at iteration i , and Q^i (with or without subscript) refers to a neighbouring candidate surface point supporting the local information at P . When it is clear that we are considering iteration step i , we omit explicit superscripts.

6.1 Preview of the algorithm

Before entering into the minutæ of the algorithm, we present a preview of the steps involved.

```

do      /* iteration step i */
  forall  $P^i \in \mathcal{S}^i$ 
    determine its contextual neighbourhood  $\mathcal{N}_P^i$ 
    forall  $Q_\alpha^i \in \mathcal{N}_P^i$ 
      compute the mapping  $\Gamma_\alpha^i \mapsto \Gamma_{P_\alpha}^i$ 
    end
    compute the chart updating  $\Gamma_P^i \mapsto \Gamma_P^{i+1}$  from local support  $\{\Gamma_{P_\alpha}^i\}_\alpha$ 
  end
  evaluate functional  $\Phi(\mathcal{A}^{i+1})$  over the new atlas of charts
  determine surface points  $\mathcal{S}^{i+1}$ 
   $i := i + 1$ 
until  $\Phi(\mathcal{A}^i) - \Phi(\mathcal{A}^{i+1}) \approx 0$ 

```

6.2 Determining the contextual neighbourhood

From a previous iteration, or from the initialization, assume that an estimated atlas of charts exists for the surface. We first determine a contextual neighbourhood, i.e., those surface points contributing to the support of the chart at P .

Recall that the principal chart $(\phi_\alpha, U_\alpha, \xi_\alpha)$ determines the quadric surface patch $\mathcal{O}_\alpha = \phi_\alpha(u, v) = \left(u, v, \frac{1}{2}(\kappa_{M_\alpha}u^2 + \kappa_{m_\alpha}v^2)\right)$, $(u, v) \in U_\alpha \subset \text{span}(\vec{M}_\alpha, \vec{m}_\alpha)$, at Q_α . We take the contextual neighbourhood of P to be the set

$$\mathcal{N}_P = \{Q_\alpha \in \mathcal{S} \mid \|P - Q_\alpha\| \leq r; \bar{P} \in \mathcal{O}_\alpha\}, \quad (5)$$

⁵We also allow the coordinates of P , i.e., its location on \mathcal{S} , to undergo refinement. We don't dwell on this aspect of updating in this paper as, for the most part, we have observed qualitatively that there is no dramatic change in the surface trace.

where $\bar{\in}$ means ‘belongs to the thick trace of’, i.e.,

$$P \in \left\{ (u, v, w) \in U \times \mathbb{R} \mid \left| \frac{1}{2}(\kappa_{M\alpha}u^2 + \kappa_{m\alpha}v^2) - w \right| \leq \epsilon \right\} = \bar{\mathcal{O}}_\alpha.$$

Essentially, this says that Q_α is within radius r of P and P lies near the local surface \mathcal{O}_α fit at Q_α .⁶ Note that the relation between P and Q_α need not be symmetric. It makes little sense to require that Q_α lie on \mathcal{O}_P , since this would be making assumptions based on the local information at P — the very information being refined. In fact, local information at P , other than position, plays no part in determining the neighbourhood of P ; it is the information at surrounding Q_α which is of importance here. Thus, the support of P comes from those neighbouring Q_α whose local surface parametrizations contain P . This is the constraint that must be satisfied, and is somewhat analogous to a generalization of the notion of curvature consistency [PZ89].

6.3 Local support

Given that Q_α is in the contextual neighbourhood of P , we next determine the mapping $\Gamma_\alpha \mapsto \Gamma_{P\alpha}$ which measures how compatible the chart $\Gamma_P = (\phi_P, U_P, \xi_P)$ is with $\Gamma_\alpha = (\phi_\alpha, U_\alpha, \xi_\alpha)$. Since a chart is composed of an instantiated local surface model ϕ and a frame ξ , both these components must be included in the determination of support.

Support is developed with reference to the following commutative diagram. γ , δ are the principal curvature and direction computations respectively over the local surface, and R is the orthogonal rotation taking the principal directions from tangent plane coordinates into the image space

⁶The algorithm is not particularly sensitive to the choice of parameter r — all examples use a fixed neighbourhood diameter $2r = 7$. Parameter ϵ determines the ‘thickness’ of \mathcal{O}_α . In practice, we have found this to be a not particularly critical parameter, and fix it at $\epsilon = 1$ for all examples. More complex definitions are conceivable, e.g., using the shortest distance from the surface, but we have found this simple form to be adequate.

coordinate system, as in Eq. (4).

$$\begin{array}{ccc}
 & \mathcal{O}_\alpha \cap \mathcal{O}_P & \\
 \phi_\alpha^{-1} \swarrow & & \searrow \phi_P \\
 \phi_\alpha^{-1}(\mathcal{O}_\alpha \cap \mathcal{O}_P) & & \phi_P^{-1}(\mathcal{O}_\alpha \cap \mathcal{O}_P) \\
 \delta_\alpha \downarrow & \gamma_\alpha \searrow & \swarrow \gamma_P & \downarrow \delta_P \\
 & \mathbb{R} \times \mathbb{R} & \\
 R_\alpha \swarrow & & \searrow R_P \\
 \mathbb{R}^3 \times \mathbb{R}^3 \times (0,0,1) & & \mathbb{R}^3 \times \mathbb{R}^3 \times (0,0,1) \\
 & \searrow R_\alpha & \swarrow R_P \\
 & \mathbb{R}^3 \times \mathbb{R}^3 \times \mathbb{R}^3 &
 \end{array} \tag{6}$$

By the assumption that $Q_\alpha \in \mathcal{N}_P$ we have $P \in \mathcal{O}_\alpha \cap \mathcal{O}_P$, where coordinates in the chart Γ_P are assigned by $\phi_P^{-1}(\mathcal{O}_\alpha \cap \mathcal{O}_P)$. In this chart, we consider that $P \in \mathcal{O}_P$ and the principal curvatures and directions at P are computed along the right-hand side of the commutative diagram (6) as $\gamma_P \circ \phi_P^{-1}(P)$ and $R_P \circ \delta_P \circ \phi_P^{-1}(P)$ respectively (as in §5.3). We have direct access to the principal coordinates and directions here since these are updated from the previous iteration, and are available in P 's principal chart itself since $\phi_P(0,0) = P = (0,0,0)$.

On the other hand, $P \in \mathcal{O}_\alpha$ and hence its coordinates in the chart Γ_α are given by $\phi_\alpha^{-1}(P)$, with the corresponding principal curvatures and directions computed along the left-hand side of the diagram as $\gamma_\alpha \circ \phi_\alpha^{-1}(P)$ and $R_\alpha \circ \delta_\alpha \circ \phi_\alpha^{-1}(P)$ respectively. These are less direct to determine than along the right-hand side since, in general, $P \neq \phi_\alpha(0,0)$; details are found in App. B. In fact by Eq. (5), P will generally not be exactly on \mathcal{O}_α , so in practice we take the closest surface point for determining its coordinates.

Now we have everything needed: the principal curvatures and directions at P in its own principal chart along the right-hand side of (6), and also as a point of Q_α 's chart along the left-hand side as the mapping $\Gamma_\alpha \mapsto \Gamma_{P_\alpha}$. Since the only transformations which we effect in carrying out the surface fit are change of coordinates by rigid translation and rotation which leave principal curvature measures invariant, these are directly comparable as scalars (indicated in the diagram by mappings into the same $\mathbb{R} \times \mathbb{R}$ space). The principal direction frames $\xi_P = (\vec{M}_P, \vec{m}_P, \vec{N}_P)$ at P as a point of \mathcal{O}_P , and $\xi_{P_\alpha} = (\vec{M}_{P_\alpha}, \vec{m}_{P_\alpha}, \vec{N}_{P_\alpha})$ at P as a point of \mathcal{O}_α consist of vector components bound to P and thus permit the computation of inner products $\langle \vec{M}_P, \vec{M}_{P_\alpha} \rangle, \langle \vec{N}_P, \vec{N}_{P_\alpha} \rangle$. Note the crucial fact that, e.g., the inner product $\langle \vec{M}_P, \vec{M}_\alpha \rangle$ is meaningless since $\vec{M}_P, \vec{M}_\alpha$ are bound to different surface points, P and Q_α respectively (computing the inner product would assume that the surface was flat; see [SZ88] for a discussion of the underlying notion of parallel transport).

The principal curvatures and directions computed along the left-hand side of the commutative diagram (6) by the mapping $\Gamma_\alpha \mapsto \Gamma_{P\alpha}$ say what the local information at P should look like according to the observed data at Q_α ; along the right-hand side they say what the current estimated data is at P . The contextual neighbourhood of P consisting of k points, say, contributes k estimates of the local information at P . In the next section we use this support $\{\Gamma_{P\alpha}^i\}_\alpha$ to map $\Gamma_P^i \mapsto \Gamma_P^{i+1}$, producing a better local estimate of the principal direction chart at P which contributes to updating the atlas from \mathcal{A}^i to \mathcal{A}^{i+1} .

6.4 Updating local information

In the previous section we discussed how neighbouring points support the local information at P . We now consider the local update rule determining a new chart at P from the local information in its contextual neighbourhood, that is, how the new chart satisfies the support constraints from its context. As well, we show how the minimization functional Φ follows naturally from the update rule $\Gamma_P^i \mapsto \Gamma_P^{i+1}$. We argue that the iterations based on the update rule converge.

Each of the k charts computed above expresses what the principal chart at P ‘ought’ to be according to the local information at neighbouring point Q_α . This collective support $\{\Gamma_{P\alpha}^i\}_\alpha$ is used to determine a better estimate Γ_P^{i+1} of the local structure at P . Briefly, the new chart is the one that best fits what the supporting points indicate that the chart at P ought to be. The ‘best fit’ here is in the natural least mean square error sense, since throughout we take the simplest form of estimate sufficient for the task. Subtleties still arise, however, and formulation of the fit must be done carefully to take into account the possibility of umbilics. By definition, if P is an umbilic of the surface \mathcal{O}_α , no unique principal direction frame exists there. However, the normal direction remains valid for the computation of the best fit \vec{N} at P . Thus we decompose the fit into two steps:

1. all k surface normals are used to determine \vec{N} , the surface normal component of the new frame; then
2. the \hat{k} non-umbilic frames determine the new frame by selecting principal direction \vec{M} (\vec{m} follows by orthonormality of \vec{N} , \vec{M}).

Principal curvatures can be computed at all k points in step 1, since only the principal directions are singular at umbilics. Figure 6 shows the increasing support of a representative P from the MR image with its neighbouring support points over the course of several iterations. In contrast, Figure 7 shows that there is no improvement in the compatibilities of the frames at an umbilic point with its supporting neighbourhood — the incompatibility is inherent due to the singularity of the direction field and not to noise in the estimates⁷.

Here we present the details of step 2, the fitting of the principal direction \vec{M} . We will make the common assumption that the noise in the data is roughly zero mean Gaussian i.i.d. (independent and identically distributed), so that a linear least-squares estimator can be used [Lue69] (this assumption is supported by example below).

Let us assume that we are given $\Gamma_{P\alpha}$ at P , with $\vec{N}_{P\alpha} = (N_{P\alpha x}, N_{P\alpha y}, N_{P\alpha z}) \in \xi_{P\alpha}$ etc., from the k neighbouring points $Q_\alpha \in \mathcal{N}_P$, $\alpha = 1, \dots, k$ of P . Step 1 above determines their ‘best fit’ unit normal as

$$\vec{N} = \left(\frac{\sum_{\alpha=1}^k N_{P\alpha x}}{\sqrt{d}}, \frac{\sum_{\alpha=1}^k N_{P\alpha y}}{\sqrt{d}}, \frac{\sum_{\alpha=1}^k N_{P\alpha z}}{\sqrt{d}} \right)$$

⁷A statistical consideration of the distribution of the support could form a possible method for the localization of umbilic points. We do not pursue this line of investigation in this paper; rather, in [SZ], we develop the localization and characterization of umbilics based on singularity theory.




Figure 6: Support over the contextual neighbourhood for the principal direction frame at a typical point P of the magnetic resonance image. The lines are the two orthonormal principal directions bound at P by the mappings $\xi_\alpha \mapsto \xi_{P\alpha}$. (a)–(d) The increasingly consistent support from initialization through one, three, and five iterations. See text for a description of the updating formula.




Figure 7: Support over the contextual neighbourhood for the principal direction frame at an umbilic point P of a synthetic elliptic paraboloid. The lines are as in Fig. 6. The compatibility with supporting points does not change significantly from the initialization through five iterations.

where

$$d = \left(\sum_{\alpha=1}^k N_{P\alpha x} \right)^2 + \left(\sum_{\alpha=1}^k N_{P\alpha y} \right)^2 + \left(\sum_{\alpha=1}^k N_{P\alpha z} \right)^2,$$

by minimizing

$$E_N^2 = \sum_{\alpha=1}^k \left\| \vec{N} - \vec{N}_{P\alpha} \right\|^2$$

subject to the constraint $\|\vec{N}\| = \sqrt{\langle \vec{N}, \vec{N} \rangle} = 1$ (we don't give details since the computations are much the same as those for step (2), which follow; however, see [SZ87]).

Step (2) subsequently determines the principal direction corresponding to the maximal principal curvature, subject to unit length and orthogonality with the surface normal. We set this up as the problem of determining the $\vec{M} = (M_x, M_y, M_z)$ which minimizes

$$E_M^2 = \sum_{\alpha=1}^k \left\| \vec{M} - \vec{M}_{P\alpha} \right\|^2$$

subject to the constraints $\langle \vec{M}, \vec{M} \rangle = 1$ and $\langle \vec{M}, \vec{N} \rangle = 0$. Letting $F = \langle \vec{M}, \vec{M} \rangle - 1$, $G = \langle \vec{M}, \vec{N} \rangle$ and introducing the Lagrange multipliers λ_F , λ_G , the above constrained minimization is equivalent to finding the stationary point(s) of

$$\begin{aligned} J &= E_M^2 + \lambda_F F + \lambda_G G \\ &= \hat{k}(M_x^2 + M_y^2 + M_z^2) - 2 \sum_{\alpha=1}^{\hat{k}} (M_x M_{P\alpha x} + M_y M_{P\alpha y} + M_z M_{P\alpha z}) \\ &\quad + \lambda_F (M_x^2 + M_y^2 + M_z^2 - 1) + \lambda_G (M_x N_x + M_y N_y + M_z N_z) + \hat{k}. \end{aligned}$$

This is solved by $(M_x, M_y, M_z, \lambda_F, \lambda_G)$ at which

$$\frac{\partial J}{\partial M_x} = 2(\hat{k} + \lambda_F)M_x + \lambda_G N_x - 2 \sum_{\alpha=1}^{\hat{k}} M_{P\alpha x} = 0, \quad (7)$$

$$\frac{\partial J}{\partial M_y} = 2(\hat{k} + \lambda_F)M_y + \lambda_G N_y - 2 \sum_{\alpha=1}^{\hat{k}} M_{P\alpha y} = 0, \quad (8)$$

$$\frac{\partial J}{\partial M_z} = 2(\hat{k} + \lambda_F)M_z + \lambda_G N_z - 2 \sum_{\alpha=1}^{\hat{k}} M_{P\alpha z} = 0, \quad (9)$$

$$\frac{\partial J}{\partial \lambda_F} = M_x^2 + M_y^2 + M_z^2 - 1 = 0, \quad (10)$$

$$\frac{\partial J}{\partial \lambda_G} = M_x N_x + M_y N_y + M_z N_z = 0. \quad (11)$$

From Eqs. (7-9)

$$M_x = \frac{2 \sum_{\alpha=1}^{\hat{k}} M_{P\alpha x} - \lambda_G N_x}{2(\hat{k} + \lambda_F)}, \quad (12)$$

$$M_y = \frac{2 \sum_{\alpha=1}^{\hat{k}} M_{P\alpha y} - \lambda_G N_y}{2(\hat{k} + \lambda_F)}, \quad (13)$$

$$M_z = \frac{2 \sum_{\alpha=1}^{\hat{k}} M_{P\alpha z} - \lambda_G N_z}{2(\hat{k} + \lambda_F)}, \quad (14)$$

which on substitution into Eqs. (10, 11) permits solving for

$$\lambda_F = -\hat{k} - \frac{1}{2} \sqrt{4d - \lambda_G^2}$$

$$\lambda_G = 2(N_x \sum_{\alpha=1}^{\hat{k}} M_{P\alpha x} + N_y \sum_{\alpha=1}^{\hat{k}} M_{P\alpha y} + N_z \sum_{\alpha=1}^{\hat{k}} M_{P\alpha z}),$$

where

$$d = \left(\sum_{\alpha=1}^{\hat{k}} M_{P\alpha x} \right)^2 + \left(\sum_{\alpha=1}^{\hat{k}} M_{P\alpha y} \right)^2 + \left(\sum_{\alpha=1}^{\hat{k}} M_{P\alpha z} \right)^2.$$

Substituting the Lagrange multipliers back into Eqs. (12-14) gives the required best fit unit principal direction vector

$$\vec{M} = \left(\frac{\lambda_G N_x - 2 \sum_{\alpha=1}^{\hat{k}} M_{P\alpha x}}{\sqrt{4d - \lambda_G^2}}, \frac{\lambda_G N_y - 2 \sum_{\alpha=1}^{\hat{k}} M_{P\alpha y}}{\sqrt{4d - \lambda_G^2}}, \frac{\lambda_G N_z - 2 \sum_{\alpha=1}^{\hat{k}} M_{P\alpha z}}{\sqrt{4d - \lambda_G^2}} \right).$$

A measure of the quality of the new chart is given by the residual errors E_N^2, E_M^2 of the frame fit, and

$$E_\kappa^2 = \sum_{\alpha=1}^k \left[\frac{(\kappa_M - \kappa_{MP\alpha})^2}{|\max(\kappa_M, \kappa_{MP\alpha})|} + \frac{(\kappa_m - \kappa_{mP\alpha})^2}{|\max(\kappa_m, \kappa_{mP\alpha})|} \right]$$

of the curvature fit, with

$$\begin{aligned} r(\Gamma_P^{i+1}) &= E_N^2 + E_M^2 + E_\kappa^2 \\ &= \sum_{\alpha=1}^k \left[\|\vec{N} - \vec{N}_{P\alpha}\|^2 + \frac{(\kappa_M - \kappa_{MP\alpha})^2}{|\max(\kappa_M - \kappa_{MP\alpha})|} + \frac{(\kappa_m - \kappa_{mP\alpha})^2}{|\max(\kappa_m - \kappa_{mP\alpha})|} \right] \\ &\quad + \sum_{\alpha=1}^k \|\vec{M} - \vec{M}_{P\alpha}\|^2 \end{aligned} \quad (15)$$

giving an 'overall squared residual' of the fit of the new chart Γ_P^{i+1} at P .

It is evident that when $r(\Gamma_P^{i+1})$ is small, the contextual neighbourhood of P is strongly supportive of the estimated chart there (see Fig. 6). This observation leads naturally to the construction of the minimization functional as

$$\Phi(\mathcal{A}^i) = \sum_{P \in \mathcal{S}} r(\Gamma_P^{i+1}).$$

We have assumed the existence of a smooth underlying surface necessary for the differential geometry. Now, an empirical consideration of Eq. (15) supports the further additive Gaussian noise assumption about the form of the data which justifies the linear least-squares fit. For simplicity, we restrict discussion to the surface normal component \vec{N} . The problem of estimating the new normal \vec{N} at P from the sample data $\vec{N}_{P\alpha}$, $\alpha = 1, \dots, k$ is equivalent to the filtering problem of estimating the signal \vec{N} from samples $\vec{N}_{P\alpha} = \vec{N} + \nu_\alpha$ of the signal plus noise ν . Figure 8 shows histograms of $r(\Gamma_P^{i+1})$ over the supporting neighbourhoods at trace points P in the ventricle region of the magnetic resonance image. Note that each of the histograms resembles a χ^2 distribution; therefore, assuming that $\|\vec{N}_{P\alpha} - \vec{N}\|^2$ is χ^2 distributed, it may be supposed [Fre71] that $\vec{N}_{P\alpha} - \vec{N} (= \nu_\alpha)$ is a standard normal distribution, and hence that a linear least-squares estimator for \vec{N} is warranted.




Figure 8: The distribution of \hat{R} over the approximately 10,000 candidate trace points of the ventricle region of the magnetic resonance image after (a)–(d) iterations one, two, three, and five. Note that it resembles a χ^2 distribution, and hence is consistent with the assumption of Gaussian i.i.d. noise.

From the local convergence of the residuals $r(\Gamma_P^{i+1}), i = 1, 2, \dots$, it can be argued that the

sequence $\Phi(\mathcal{A}^i)$, $i = 1, 2, \dots$ converges as well⁸. Consider each $\Gamma_{P_a}^i$, $a = 1, 2, \dots, n$, as a normally distributed random variable with mean and variance μ_P, σ_P^2 , say. Then Γ_P^{i+1} determined from these random variables by the least-squares averaging is itself a random variable with mean μ_P and variance σ_P^2/k [Fre71]. A result of the iterated averaging is to bring information from an increasing number of points to bear on the computation of Γ_P^{i+1} , effectively increasing k . As the variance σ_P^2/k drives toward zero through repeated iteration, the mean of Γ_P^{i+1} approaches μ_P , ensuring that each residual term r_P approaches 0 (see Fig. 8 for empirical support). Thus the sum Φ of these residuals also approaches zero and, under the above assumptions, it may be supposed that the iterative constraint satisfaction algorithm converges⁹.

7 Results

The methods of inferring surface trace points and of estimating their local differential geometric structure presented in this paper were tested both on synthetic images and on clinical MRI data. Here we present examples of the results of the experiments in computing the principal direction fields and determining a coarse surface segmentation based on the Gaussian and mean curvatures. Remember that, since our ultimate goal will be a qualitative description of shape, we only require a degree of accuracy sufficient to permit computation of singularities (hence the use of an approximate convergence condition). We examine the special case of neighbourhoods of umbilic points in [SZ].

Figures 9, 10 show how well the contextual refinement techniques perform in recovering the structure in the synthetic bloblike image which has been degraded by the addition of Gaussian noise to the image intensities. Five refinement iterations on the degraded image are sufficient to permit a coarse curvature type segmentation, which is not significantly different from that on the noiseless image, and which clearly distinguishes the regions of elliptic points from hyperbolic points (Fig. 9). Figure 10 shows the increase in consistency in one of the principal direction fields over a region of the synthetic object.

Figures 2, 3(d) showed examples of the effectiveness of minimizing Φ on the recovery of the direction fields and on computing a coarse curvature segmentation on the face region of the magnetic resonance image. Figures 11, 12 continue this example by showing the progressive improvement in the principal direction and curvature estimates over the course of several iterations until (approximate) convergence of the algorithm. To test stability, forcing the procedure to continue ‘beyond convergence’ as defined in §6.4 produced no significant change in either the fields or the curvature classification.

The final example shows the extraction of the surface of the lateral ventricles (the lighter grey structures in Figs. 1 (f)–(i) whose tangent planes were in Fig. 3(c)). Fig. 13 shows the labeling of the points on the ventricle surfaces according to their curvature class, and Fig. 14 shows one of the estimated principal direction fields on a region of the surface.

⁸ Alternative approaches might be developed through the Euler equations of a global minimization of variational cross sections through the frame bundle [Bur85], or through considering the iterated averaging as a diffusion process.

⁹ As a practical termination condition in the examples, we take $\Phi(\mathcal{A}^i) - \Phi(\mathcal{A}^{i+1}) < \epsilon$ with $\epsilon \approx 2\%$ of $\Phi(\mathcal{A}^i)$.







Figure 9: Segmentation of surface points of a bloblike synthetic image degraded by noise. (a) The shading shows the orientations of the tangent planes; the light source is from the upper right. (b) Tangent planes extracted from the noiseless image. (c) A 2-D slice through the original noiseless blob, and (d) the same slice after the addition of zero mean Gaussian noise with standard deviation 20. (e) Coarse segmentation of the surface points of the original blob according to curvature type. Elliptic points are shown lighter than the background, hyperbolic points are darker (the far side of the object and the background show through the spaces between the planes). (f) The initial curvature type classification of the blob with additive noise, and (g) the refined curvature labeling after five iterations of the algorithm.




(a)



(b)




(c)




(d)


Figure 10: Principal directions on a region of the synthetic image. (a) The shading, with light source at upper left, shows the orientations of the tangent planes. (b) One of the principal direction fields (the one corresponding to the greater principal curvature) on a region of the noiseless object. (c) Initial estimate of the direction field when zero mean Gaussian noise of standard deviation 20 has been added to the image. (d) The direction field recovered after only two iterations.




(a)



(b)



(c)



(d)

Figure 11: The nose region of the magnetic resonance image showing one principal direction field. (a) The shading shows the orientation of the tangent planes—the light source is from the right. (b) The initial estimate of the principal direction field corresponding to the greater principal curvature. (c) The refined direction field after one, and (d) five iterations. The apparent multiple principal directions at some points is due to the many-to-one projection of a curved region of the surface (with trace points restricted to a grid) onto the viewing plane.




Figure 12: Coarse segmentation by curvature classes of the face region of the magnetic resonance image. (a) The orientation of the tangent planes is shown by the shading. (b) The initial segmentation with bright corresponding to elliptic points, and dark to hyperbolic points. (c) The segmentation after one, and (d) at convergence after seven refinement iterations.







Figure 13: Curvature labeling of the surface of the lateral ventricles, with shaded tangent planes in (a). (b) Initial labeling, and (c), (d) after two and five refinement iterations. Refer to Fig. 3 to see the hyperbolic character of the central portion of the ventricles.




(a)



(b)



(c)



(d)

Figure 14: Principal direction field on a region of the surface of the lateral ventricles (lower left corner of Fig. 13). (a) The light source for the shading of the tangent planes is from the right. (b) The initialization field, and (c), (d) refinements after two and three iterations respectively.

8 Fibre bundles and principal directions

An alternate language in which to cast the methods presented in this paper, and which in fact partially inspired their original development, is in terms of *fibre bundles* [Ste74, Spa66]. We give an informal presentation in this section, and refer the interested reader to [SZ88] for more detail.

A fibre bundle consists of three components: two spaces, the *base space* \mathcal{S} and the *bundle space* or *total space* \mathcal{B} , and a *projection* $\pi : \mathcal{B} \rightarrow \mathcal{S}$ assigning elements of the base space to elements of the total space. For $P \in \mathcal{S}$, the set of points $\pi^{-1}(P) \subset \mathcal{B}$ forming the inverse image of the projection is called the *fibre at P*, and this provides the sense in which the total space exists ‘over’ the base space. Certain additional constraints must be introduced to this structure, in particular, all the fibres are required to be topologically equivalent. In our case, each is in fact a copy of another space called the *standard fibre* \mathcal{F} .

For clarity, we simplify the situation somewhat and consider that the base space is given as a fixed set of surface points, and that the surface normal at each point is known (recall that both the point positions and normals are really only estimates). We are to determine the principal directions of the surface at the points of the base space. Our standard fibre \mathcal{F} then consists of the space of all right-handed orthonormal trihedra ξ , and the fibre over P is the space $\Xi_P = \{(\vec{P}, \vec{Q}, \vec{N})\}$, where \vec{N} is the surface normal at P and (\vec{P}, \vec{Q}) is any orthonormal basis for the tangent plane at P . This is shown schematically in Fig. 15 as the tangent planes at P with different bases stacked over P along the normal.

The above assigns a ‘vertical’ structure at each surface point. We now consider a ‘horizontal’ slice through the bundle space, called a (*cross*) *section* of \mathcal{B} , by taking a smooth mapping $\sigma : \mathcal{S} \rightarrow \mathcal{B}$, such that $\pi \circ \sigma(P) = P$. Essentially, σ assigns one element of the standard fibre to each element of the base space, and the smoothness condition ensures the orderly transition between elements at neighbouring points (a smooth vector field on a surface is an example of a cross section through a bundle space where the standard fibre consists of surface tangent vectors). In Fig. 15, the darker tangent planes represent a cross section through the bundle space.

Now with the above machinery in place, the solution to the problem of estimating the principal directions over a surface can be posed in terms of determining a particular cross section σ , where each $\sigma(P)$ corresponds to the principal direction frame at P . The solution method is to search for the cross section which minimizes a functional Φ based on the neighbourhood support r_P derived from connections between frames. This can be thought of figuratively in Fig. 15 as sliding along the fibre at P until the frame is found which is most compatible with the frames already estimated at neighbouring points. That frame is then chosen as $\sigma(P)$ for the next iteration, and so on until some (sufficiently) minimal value of Φ is attained.

9 Conclusions

Early image understanding seeks to derive analytic representations from image data. In this paper, we have presented steps towards this goal for true three-dimensional imagery by focusing on the inference of trace points and the estimation of associated differential structure. Trace points are the points which lie on surfaces, and the problem of trace point extraction is an extension of ‘edge detection’ applied to 3-D images. We developed a mathematically rigorous technique for the reliable extraction of such points based on an atlas of local surface charts, with overlapping charts used to increase the global coherence of differential structure. Note that something of a ‘chicken and egg’ problem arose here. The machinery of differential geometry could only be applied for the analysis of local structure once the trace points were identified, yet differential geometric tools seemed




Figure 15: The Cartesian product of S and all its possible tangent planes with orthonormal bases is an example of a fibre bundle over the surface. A cross section σ of the bundle is formed by 'smoothly choosing' one basis for each surface point, e.g., the darker planes.

appropriate for the identification of the trace points themselves. We proposed a solution to this quandary by casting the problem as functional minimization over atlases, where each principal chart met appropriate minimum criteria, and developed an implementation as an iterative constraint satisfaction procedure based on local surface smoothness properties.

Trace points must be inferred accurately since any surface fitting scheme is inherently limited by the quality of the data points; the problem is analogous to fitting a function to data points without knowing which data points are noise. However, it can be argued that for shape description and eventual object recognition, the exact recovery of local structure everywhere is less important than the identification of singular surface points which prove stable to noise and small surface perturbations. Among these are the parabolic and umbilic points, and we showed how to reliably compute coarse surface segmentations into elliptic and hyperbolic regions separated by parabolic points. In [SZ], we show how to compute umbilic points from the refined direction fields.

We tested the constraint satisfaction algorithm on several kinds of data. First, we used 'perfect' data where the principal curvature and direction fields were easily derived. Next, we added controlled amounts of noise to the images and investigated the performance in this more realistic situation. Finally, we used real clinical magnetic resonance images. In all three situations, as shown by the results, the procedures proved remarkably strong at estimating surface trace points and recovering differential structure and singularities. Ferrie et al [FLW89] have shown how these techniques can also be applied to laser rangefinder imagery.

Further to our overall goal of providing a stable basis for qualitative shape analysis, Gaussian and mean curvatures, considered as summaries of the directional curvature properties, provided an example of segmentation of the putative surface trace points into stable, qualitatively homogeneous, surface patches. Given the sampling and quantization noise inherent in real 3-D imagery, however, such partitions must necessarily be coarse. As such, they are useful for a rough partitioning of objects into parts, and for approximate matching into prototypical object classes. We believe that the results, taken together with the singularity computations in [SZ], can provide the foundations for improved object inferences.

References

- [AFH81] Ehud Artzy, Gideon Frieder, and Gabor T. Herman. The theory, design, implementation and evaluation of a three-dimensional surface detection algorithm. *Computer Graphics and Image Processing*, 15:1–24, 1981.
- [BB82] Dana H. Ballard and Christopher M. Brown. *Computer Vision*. Prentice-Hall, 1982.
- [BGM82] Thomas Banchoff, Terence Gaffney, and Clint McCrory. *Cusps of Gauss Mappings*. Pitman, Boston, 1982.
- [BJ85] Paul Besl and Ramesh Jain. Intrinsic and extrinsic surface characteristics. In *IEEE Proceedings on Computer Vision and Pattern Recognition*, pages 226–233, June 1985.
- [BJ86] Paul J. Besl and Ramesh C. Jain. Segmentation through symbolic surface descriptions. In *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pages 77–85, June 1986.
- [BLR83] Ruzena Bajcsy, Robert Lieberman, and Martin Reivich. A computerized system for the elastic matching of deformed radiographic images to idealized atlas images. *Journal of Computer Assisted Tomography*, 7(4):618–625, 1983.
- [Boi84] Jean-Daniel Boissonnat. Geometric structures for three-dimensional shape representation. *ACM Transactions on Graphics*, 3:266–286, October 1984.
- [BPYA85] Michael Brady, Jean Ponce, Alan Yuille, and Haruo Asada. Describing surfaces. In Hideo Hanafusa and Hirochika Inoue, editors, *Proceedings of the Second International Symposium on Robotics Research*, pages 5–16, MIT Press, Cambridge, Mass., 1985.
- [BTTU84] Linda J. Brewster, Sushma S. Trivedi, Heang K. Tuy, and Jayaram K. Udupa. Interactive surgical planning. *IEEE Computer Graphics and Applications*, 4(3):31–40, March 1984.
- [Bur85] William L. Burke. *Applied Differential Geometry*. Cambridge University Press, Cambridge, 1985.
- [Cap87] John Danilo Cappelletti. *Three-Dimensional Boundary Following*. Technical Report CAR-TR-314, Computer Vision Laboratory, Center for Automation Research, University of Maryland, College Park, August 1987.
- [dC76] Manfredo P. do Carmo. *Differential Geometry of Curves and Surfaces*. Prentice-Hall, Englewood Cliffs, 1976.

- [DH72] Richard O. Duda and Peter E. Hart. *Pattern Classification and Scene Analysis*. Wiley, New York, 1972.
- [Eis60] Luther Pfahler Eisenhart. *A Treatise on the Differential Geometry of Curves and Surfaces*. Dover, New York, 1960. originally published by Ginn and Company, 1909.
- [FH86] O.D. Faugeras and M. Hebert. The representation, recognition, and locating of 3-D objects. *International Journal of Robotics Research*, 5(3):27-52, Fall 1986.
- [FKU77] Henry Fuchs, Zvi M. Kedem, and Samuel P. Uselton. Optimal surface reconstruction from planar contours. *Communications of the ACM*, 20(10):693-702, October 1977.
- [FLW89] Frank Ferrie, Jean Lagarde, and Peter Whaite. Towards sensor-derived models of objects. In *Proceedings, Vision Interface '89*, London, Ont., June 1989.
- [Fre71] John Freund. *Mathematical Statistics*. Prentice-Hall, Englewood Cliffs, 1971.
- [GHJ75] Richard Gordon, Gabor T. Herman, and Steven A. Johnson. Image reconstruction from projections. *Scientific American*, 233(4):56-68, October 1975.
- [GP74] Victor Guillemin and Alan Pollack. *Differential Topology*. Prentice-Hall, Englewood Cliffs, 1974.
- [HL79] Gabor T. Herman and Hsun Kao Liu. Three-dimensional display of human organs from computed tomograms. *Computer Graphics and Image Processing*, 9:1-21, 1979.
- [HR84] D.D. Hoffman and W.A. Richards. Parts of recognition. *Cognition*, 18:65-96, 1984.
- [HW81] Robert M. Haralick and Layne Watson. A facet model for image data. *Computer Graphics and Image Processing*, 15:113-129, 1981.
- [HWL83] Robert M. Haralick, Layne T. Watson, and Thomas J. Laffey. The topographic primal sketch. *International Journal of Robotics Research*, 2:50-70, Spring 1983.
- [HZ83] Robert A. Hummel and Steven W. Zucker. On the foundations of relaxation labeling processes. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, PAMI-5:267-287, 1983.
- [KvD82] Jan J. Koenderink and Andrea J. van Doorn. The shape of smooth objects and the way contours end. *Perception*, 11:129-137, 1982.
- [LC87] William E. Lorensen and Harvey E. Cline. Marching cubes: a high resolution 3D surface reconstruction algorithm. *Computer Graphics*, 21(4):163-169, 1987. Proceedings of SIGGRAPH '87, Anaheim.
- [Lev85] Martin D. Levine. *Vision in Man and Machine*. McGraw-Hill, 1985.
- [Lue69] David G. Luenberger. *Optimization by Vector Space Methods*. Wiley, New York, 1969.
- [Mar82] David Marr. *Vision*. Freeman, 1982.
- [PB85] Jean Ponce and Michael Brady. Toward a surface primal sketch. In *Proceedings, IJCAI*, 1985.

- [PS78] Tim Poston and Ian Stewart. *Catastrophe Theory and its Applications*. Pitman, London, 1978.
- [Pyk82] Ian L. Pykett. NMR imaging in medicine. *Scientific American*, 246(5):78-88, May 1982.
- [PZ89] Pierre Parent and Steven W. Zucker. Trace inference, curvature consistency, and curve detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 11(8):823-839, August 1989.
- [RGA83] Michael L. Rhodes, William V. Glenn, and Yu-Ming Azzawi. Computer graphics and an interactive stereotactic system for CT-aided neurosurgery. *IEEE Computer Graphics and Applications*, 3(5):31-38, August 1983.
- [Ric] J. Rice. *IMSL Reference Edition*. McGraw-Hill.
- [RK81] Azriel Rosenfeld and A.C. Kak. *Digital Picture Processing*. Academic Press, 1981.
- [Spa66] Edwin Spanier. *Algebraic Topology*. McGraw-Hill, New York, 1966.
- [Ste74] Norman E. Steenrod. *Topology of Fibre Bundles*. Princeton University Press, Princeton, NJ, 1974.
- [SZ] Peter T. Sander and Steven W. Zucker. Singularities of principal direction fields from 3-D images. *IEEE Transactions on Pattern Analysis and Machine Intelligence*. To appear. Available as Technical Report CIM-88-7, McGill Research Center for Intelligent Machines, McGill University, Montréal.
- [SZ86] Peter T. Sander and Steven W. Zucker. Stable surface estimation. In *Proceedings of the Eighth International Conference on Pattern Recognition*, pages 1165-1167, Paris, October 1986.
- [SZ87] Peter T. Sander and Steven W. Zucker. Tracing surfaces for surfacing traces. In *Proceedings of the First International Conference on Computer Vision*, pages 241-249, London, June 1987.
- [SZ88] Peter T. Sander and Steven W. Zucker. *Inferring Differential Structure from 3-D Images: Smooth Cross Sections of Fibre Bundles*. Technical Report CIM-88-6, McGill Research Center for Intelligent Machines, McGill University, Montréal, 1988.
- [Tri86] Sushma S. Trivedi. Interactive manipulation of three dimensional binary scenes. *The Visual Computer*, 2(4):209-218, August 1986.
- [TRS80] Micheal M. Ter-Pogossian, Marcus E. Raichle, and Burton E. Sobel. Positron-emmission tomography. *Scientific American*, 243(4):170-181, October 1980.
- [VMW83] Michael W. Vannier, Jeffrey L. Marsh, and James O. Warren. Three dimensional computer graphics for craniofacial surgical planning and evaluation. *Computer Graphics*, 17(3):263-273, July 1983. Proceedings of SIGGRAPH '83.
- [YL87] N. Yokoya and M.D. Levine. *Range Image Segmentation Based on Differential Geometry: A Hybrid Approach*. Technical Report TR-CIM-87-16, McGill Research Center for Intelligent Machines, McGill University, Montréal, September 1987.

- [ZH81] S.W. Zucker and R.M. Hummel. A three-dimensional edge operator. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, PAMI-3(3):324–331, May 1981.
- [Zuc87] Steven W. Zucker. The emerging paradigm of computational vision. *Ann. Rev. Comput. Sci.*, 69–89, 1987.

A 3D gradient operator

The three-dimensional ‘edge detection’ preceding the local information computations consists of 2 steps: (i) applying a local digital gradient operator everywhere in the image to detect intensity differences; and (ii) interpreting the operator responses as evidence for edge elements. The 3-D gradient operator [ZH81] used for (i) is a direct generalization of the 2-D Sobel operator [DH72]. A ‘3-D edge’ separates a dark volume, say, from a light one. The operator template models such an ideal edge as an array of +1’s separated by a unit plane from an array of –1’s, with the orientation of the normal determining the orientation of the edge. The gradient operator is a discrete version of the functions

$$\begin{aligned}\phi_x &= \frac{x}{\sqrt{x^2 + y^2 + z^2}}, \\ \phi_y &= \frac{y}{\sqrt{x^2 + y^2 + z^2}}, \\ \phi_z &= \frac{z}{\sqrt{x^2 + y^2 + z^2}}\end{aligned}$$

which compute edges in the X , Y , Z directions respectively. Application of the operator consists of convolution of each of the functions with the image, giving estimated surface normal $\vec{N}(u, v, w) = (N_x(u, v, w), N_y(u, v, w), N_z(u, v, w))$,

$$\begin{aligned}N_x(u, v, w) &= \phi_x * I(u, v, w) = \iiint_{\delta} \phi_x(x, y, z) I(x - u, y - v, z - w) dx dy dz, \\ N_y(u, v, w) &= \phi_y * I(u, v, w) = \iiint_{\delta} \phi_y(x, y, z) I(x - u, y - v, z - w) dx dy dz, \\ N_z(u, v, w) &= \phi_z * I(u, v, w) = \iiint_{\delta} \phi_z(x, y, z) I(x - u, y - v, z - w) dx dy dz\end{aligned}$$

at image point (u, v, w) , where $\delta = \{(x, y, z) \mid x^2 + y^2 + z^2 \leq 1\}$.

Step (ii) of the edge detection involves the interpretation of the estimated normals \vec{N} . The simplest measure of the quality of the match to a high-contrast edge (a surface point) is the vector norm of \vec{N} . In the application, we use a slightly more complex evaluation and precede the norm calculation with a local maxima selection to give a somewhat better localization of the edge. Only those convolution responses which are above some lower threshold (necessary since the operator gives a response at all image locations) and which survive local maxima selection are likely to actually indicate surface trace points. Points with small responses are eliminated.

B Principal curvature and direction computations

For the surface $\phi(u, v) = (u, v, \frac{1}{2}(\kappa_M u^2 + \kappa_m v^2))$, the local tangent plane basis vectors (the Gauss frame) at (u, v) are

$$\vec{P}(u, v) = \frac{\partial \phi}{\partial u}(u, v) = \phi_u(u, v) = (1, 0, \kappa_M u),$$

$$\vec{Q}(u, v) = \frac{\partial \phi}{\partial v}(u, v) = \phi_v(u, v) = (0, 1, \kappa_m v)$$

in the parametrization $(\vec{M}, \vec{m}, \vec{N})$. The coefficients of the first and second fundamental forms are

$$\begin{aligned} E(u, v) &= \langle \phi_u, \phi_u \rangle = \kappa_M^2 u^2 + 1, \\ F(u, v) &= \langle \phi_u, \phi_v \rangle = \kappa_M \kappa_m uv, \\ G(u, v) &= \langle \phi_v, \phi_v \rangle = \kappa_m^2 v^2 + 1, \\ e(u, v) &= \langle \vec{N}, \phi_{uu} \rangle = \frac{\kappa_M}{\sqrt{\kappa_M^2 u^2 + \kappa_m^2 v^2 + 1}}, \\ f(u, v) &= \langle \vec{N}, \phi_{uv} \rangle = 0, \\ g(u, v) &= \langle \vec{N}, \phi_{vv} \rangle = \frac{\kappa_m}{\sqrt{\kappa_M^2 u^2 + \kappa_m^2 v^2 + 1}}, \end{aligned}$$

and

$$\vec{N}(u, v) = \frac{1}{\sqrt{\kappa_M^2 u^2 + \kappa_m^2 v^2 + 1}} (-\kappa_M u, -\kappa_m v, 1)$$

is the unit surface normal at P . Now, the differential of the surface normal at (u, v) is given by the map

$$D_1 \vec{N}|_{(u,v)} = \begin{pmatrix} a_{11} & a_{12} \\ a_{21} & a_{22} \end{pmatrix} \in \mathcal{L}(\mathbb{R}^2, \mathbb{R}^2),$$

where \mathcal{L} is the space of linear mappings $\mathbb{R}^2 \rightarrow \mathbb{R}^2$ and the $a_{ij}(u, v)$ are

$$\begin{aligned} a_{11} &= \frac{fF - eG}{EG - F^2} = -\frac{\kappa_M \kappa_m^2 v^2 + \kappa_M}{(\kappa_M^2 u^2 + \kappa_m^2 v^2 + 1)^{3/2}}, \\ a_{12} &= \frac{gF - fG}{EG - F^2} = \frac{\kappa_M \kappa_m^2 uv}{(\kappa_M^2 u^2 + \kappa_m^2 v^2 + 1)^{3/2}}, \\ a_{21} &= \frac{eF - fE}{EG - F^2} = \frac{\kappa_M^2 \kappa_m uv}{(\kappa_M^2 u^2 + \kappa_m^2 v^2 + 1)^{3/2}}, \\ a_{22} &= \frac{fF - gE}{EG - F^2} = -\frac{\kappa_M^2 \kappa_m u^2 + \kappa_m}{(\kappa_M^2 u^2 + \kappa_m^2 v^2 + 1)^{3/2}}. \end{aligned}$$

The principal curvatures, $\kappa_1 \geq \kappa_2$, of s at P are, for λ_1, λ_2 the eigenvalues of the differential $D_1 \vec{N}$,

$$\begin{aligned} -\kappa_1 &= \lambda_1 = \frac{1}{2} \left(a_{11} + a_{22} - \sqrt{(a_{11} - a_{22})^2 + 4a_{12}a_{21}} \right), \\ -\kappa_2 &= \lambda_2 = \frac{1}{2} \left(a_{11} + a_{22} + \sqrt{(a_{11} - a_{22})^2 + 4a_{12}a_{21}} \right). \end{aligned}$$

The expression $D_1 \vec{N} \vec{x}_i = -\kappa_i \vec{x}_i$, $i = 1, 2$, gives the eigenvectors

$$\vec{x}_1 = \begin{cases} \left(a_{12}, -\frac{1}{2}(a_{11} - a_{22} + \sqrt{(a_{11} - a_{22})^2 + 4a_{12}a_{21}}) \right) & \text{when } a_{11} \geq a_{22}, \\ \left(\frac{1}{2}(a_{11} - a_{22} - \sqrt{(a_{11} - a_{22})^2 + 4a_{12}a_{21}}), a_{21} \right) & \text{otherwise;} \end{cases}$$

$$\vec{x}_2 = \begin{cases} \left(\frac{1}{2}(a_{11} - a_{22} + \text{sqrt}), a_{21} \right) & \text{when } a_{11} \geq a_{22}, \\ \left(-a_{12}, \frac{1}{2}(a_{11} - a_{22} - \text{sqrt}) \right) & \text{otherwise,} \end{cases}$$

where $\text{sqrt} = \sqrt{(a_{11} - a_{22})^2 + 4a_{12}a_{21}}$.

Imprimé en France
par
l'Institut National de Recherche en Informatique et en Automatique

