Charting Surface Structure

Peter T. Sander
INRIA — Rocquencourt
78153 Le Chesnay Cedex, France
and
INRIA — Sophia-Antipolis
06565 Valbonne Cedex, France

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

Abstract

Computing surface curvature would seem to be a simple application of differential geometry, but problems arise due to noise and the quantized nature of digital images. We present a method for determining principal curvatures and directions of surfaces estimated from three-dimensional images. We use smoothness constraints to *connect* different surface points and by then comparing information over local neighbourhoods we iteratively update the information at each point to ensure that this information is consistent over the estimated surface.

1 Introduction

Differential geometry supplies the tools necessary for the computation of surface tangent and curvature. However, applying these tools to the determination of the structure of surfaces in digital three-dimensional images is problematic since surfaces are only implicit in image intensities (the "edge detection" problem), and computations are adversely affected by image quantization and noise. In this paper, we are concerned with estimating the following from 3-D images: the surface trace (that is, the image points lying on a surface); the surface normal (or tangent plane) at each surface trace point; and the principal curvatures and directions of the surface.

A simple paradigm for computing surface structure consists of the following steps: (i) extract the surface trace points from the image; (ii) fit some surface(s) to the points; (iii) compute tangent and curvature for the fit surface(s); and (iv) take this information as belonging to the underlying surface. In practice, this straightforward approach generally leads to unsatisfactory results due to well-known problems with the so-called "edge detectors" for computing (i), and to instabilities in the surface fitting process (ii). The general conclusion has been that curvature computations are somewhat delicate.

Support for this research was partially provided by NSERC Grant A4470 and MRC Grant MA 6125.
*Fellow, Canadian Institute for Advanced Research

We show that this need not be the case, that curvature computation can indeed be made reliable. We largely follow the above scheme, but consider that it yields only an initial estimate of local surface characteristics which must be further refined. In particular, we consider that local surfaces fit in (ii) correspond to the differential geometric notion of a *chart,* and that the surface itself consists of overlapping charts. The fundamental idea then is to look for such a surface within a class of admissible surfaces which minimizes and appropriate functional ensuring that local curvature information is consistent over neighbourhoods. The implementation of these ideas of local curvature consistency is by iterative minimization with similarities to relaxation labeling procedures [1]. At each step of the iterative refinement process, information at each trace point is updated by making it more consistent with information over local neighbourhoods. This leads to more consistent information over the whole surface.

In this paper, we present experiments in applying these methods to clinical 3-D magnetic resonance (MR) imagery. Work related to ours is found primarily in the domain of laser rangefinder image understanding, e.g., [2,3,4], with the difference that their surface trace points, albeit noisy, are available *a priori* (see Ferrie *et al* [5] for an application of methods based on ours to rangefinder images).

2 Local surface parametrization

Our working definition of a surface $\mathcal{S}$ (technically a *2-dimensional differentiable manifold* [6]) is a set of points such that for every $r \in \mathcal{S}$, there exist open sets $U \subset \mathbb{R}^2$, and $V \subset \mathbb{R}^3$ with $r \in V$, and a diffeomorphism $\phi : U \rightarrow V \cap \mathcal{S}$ (recall that a diffeomorphism is a smooth bijective map with a smooth inverse). The pair (ϕ, U) is a *parametrization of* $\mathcal{S}$ *at* r, $\phi(U) = \mathcal{O}$ is a *coordinate neighbourhood,* and ϕ^{-1} is a *chart* which assigns coordinates at r. Note that when r is in two different coordinate neighbourhoods $r \in \mathcal{O}_1 \cap \mathcal{O}_2$ (we also say that r is in overlapping charts or parametrizations), then the *change of coordinate* mapping taking (part of) $\mathbb{R}^2$ into $\mathbb{R}^2$

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

is smooth with a smooth inverse. Our refinement methods of §3 are based in an essential way on the change of coordinate mappings at overlapping parametrizations.

We depart slightly from the above notation by explicitly including a set of right-handed orthonormal axes in the notion of local chart, 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 r, and $(\vec{P}, \vec{Q})$ spans the tangent plane so that ξ is a *Darboux trihedron* or *frame* [6] of $\mathcal{S}$ bound to r. 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 at r, then ξ is the *principal direction frame*, and we term such a chart a *principal chart.*

In this paper, we can only present a brief sketch of the process of instantiating the local charts from the input image (the reader is referred to [7] for details). Following the scheme outlined in the introduction, step (i) estimates surface trace points as the thresholded output of a 3-D gradient operator [8]. The operator also returns an estimate of the surface normal which allows us to set up a frame representing the local tangent

plane coordinate system. There is a wide choice for the mapping ϕ of step (ii), and we take it to be a *parabolic (*or *non-central) quadric* surface

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

where p, q, n are in the local coordinates $(\vec{P}, \vec{Q}, \vec{N})$. This is the simplest smooth surface from which curvature can be computed. We instantiate a chart at each trace point r by fitting a parabolic quadric via least-squares error minimization to the positions and normals of neighbouring estimated trace points. Note that for the quadric surface fit at r, $\phi(0,0) = r = (0,0,0)$. In step (iii), we compute principal curvatures and directions of the chart from Eq. (1). We use the principal directions as a new basis for the tangent plane at r so that the frame ξ becomes $(\vec{M}, \vec{m}, \vec{N})$ where $\vec{M}, \vec{m}$ are the principal directions corresponding to the maximum and minimum principal curvatures κ, λ respectively. In this principal chart, the mapping ϕ assumes a particularly significant form with parameters $a = \kappa, b = 0, c = \lambda$.

As we show in §4, the results obtained by steps (i)-(iii) do not warrant step (iv) directly — they should only be treated as rough estimates of the structure of the underlying surface. In the following section, we develop a method for refining such estimates.

3 Refining estimated local information

Our goal is to determine a surface of the form introduced above which minimizes a residual functional based on local tangent and curvature compatibilities. In practice, our approach to refining estimated surface structure is similar to relaxation labeling [1] in that the information at each candidate surface point is updated based on information from neighbouring surface points. It differs in that we don't use explicit quantized labels but keep everything in the continuous domain.

We will develop a rule $(\phi_r, U_r, \xi_r)^i \mapsto (\phi_r, U_r, \xi_r)^{i+1}$ for iteratively updating the principal chart at each estimated surface point $r \in \mathcal{S}^i$ so that the outcome of iteration i is to effectively update the whole estimated surface $\mathcal{S}^i \mapsto \mathcal{S}^{i+1}$. Three issues must be specified:

1. which neighbouring points s should contribute to the support of the principal chart at r, i.e., which points determine its *contextual neighbourhood* $\mathcal{N}_r$;
2. how the information at point $s \in \mathcal{N}_r$ supports the principal chart at r;
3. given 1 and 2, how to update the principal chart at r.

For consistent notation, let r refer to the point being refined and s be a point of its supporting contextual neighbourhood.

3.1 Contextual neighbourhood

Recall that the principal chart (ϕ_s, U_s, ξ_s) at $s \in \mathcal{S}$ determines the quadric surface patch $\mathcal{O}_s = \phi_s(u,v) = \left(u, v, \frac{1}{2}(\kappa_s u^2 + \lambda_s v^2)\right)$, $(u,v) \in U_s \subset \text{span}(\vec{M}_s, \vec{m}_s)$. We take the contextual neighbourhood of $r \in \mathcal{S}$ to be the set

$$\mathcal{N}_r = \{s \in \mathcal{S} : \|r - s\| \leq d \text{ and } r \in \mathcal{O}_s\},$$

for some neighbourhood radius d ($\|\bullet\|$ is the Euclidean norm).

3.2 Local support (compatibility)

Given that s is in the contextual neighbourhood of r, we wish to determine how compatible the (principal) chart at r is with the (principal) chart at s. This cannot be computed directly since the charts are bound to different points. (Note the crucial fact that, e.g., the inner product $\langle \vec{M}_r, \vec{M}_s \rangle$ is meaningless since $\vec{M}_r, \vec{M}_s$ are bound to different surface points, r and s respectively — computing the inner product would assume that the surface was flat. See [7] for a discussion of the underlying notion of parallel transport.) Thus, we set up a mapping "transporting" the chart at s to the point r, which we denote $(\phi_s, U_s, \xi_s) \mapsto (\phi_{(s,r)}, U_{(s,r)}, \xi_{(s,r)})$, as explained below.

Support is developed with reference to the following commutative diagram. γ, δ are the principal curvature and direction computations respectively in a chart, e.g., from Eq. (1, and R is the rotation taking the principal directions from tangent plane coordinates into the image space coordinate system.

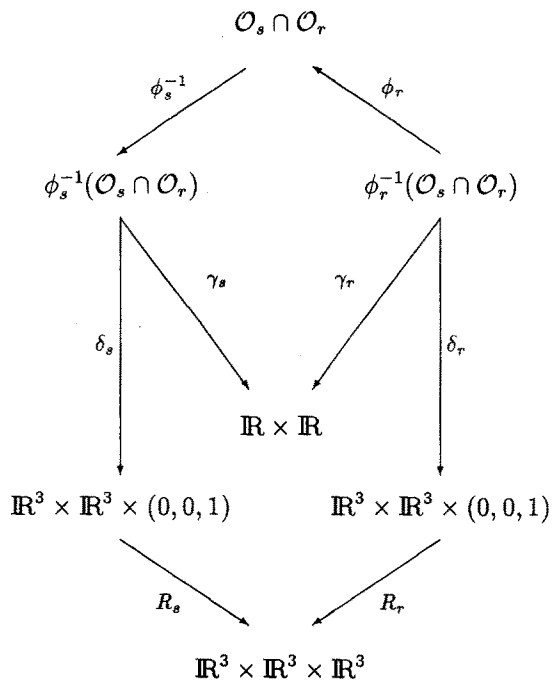

By the assumption that $s \in \mathcal{N}_r$ we have $r \in \mathcal{O}_s \cap \mathcal{O}_r$, and if we consider that $r \in \mathcal{O}_r$ then coordinates assigned by $\phi_r^{-1}(r)$. In this chart, the principal curvatures and directions at r are computed along the right-hand side of the above diagram as $\gamma_r \circ \phi_r^{-1}(r)$ and $R_r \circ \delta_r \circ \phi_r^{-1}(r)$ respectively. We have "direct" access to the principal curvatures and directions here since these are updated from the previous iteration, and are available in r's principal chart itself as the coefficients κ_r, λ_r of ϕ_r, and as ξ_r respectively.

On the other hand, $r \in \mathcal{O}_s$ and hence its coordinates in the chart at s are given by $\phi_s^{-1}(r)$, with the corresponding principal curvatures and directions computed along the left-hand side of the diagram as $\gamma_s \circ \phi_s^{-1}(r)$ and $R_s \circ \delta_s \circ \phi_s^{-1}(r)$ respectively. These are less direct to determine than along the right-hand side since $r \neq \phi_s(0,0)$; details may be found in [7]. We consider this computation as the determination of the mapping

$(\phi_s, U_s, \xi_s) \mapsto (\phi_{(s,r)}, U_{(s,r)}, \xi_{(s,r)})$ since it says what r's structure ought to be as seen from s.

Now we have the principal curvatures and directions at r in its own principal chart along the right-hand side of the diagram, and also as a point of s's chart as computed along the left-hand side of the diagram. These principal curvature measures are directly comparable as scalars — indicated in the diagram by mappings γ into the same $\mathbb{R} \times \mathbb{R}$ space. The principal directions computed by δ as frames $\xi_r = (\vec{M}_r, \vec{m}_r, \vec{N}_r)$ at r as a point of $\mathcal{O}_r$, and $\xi_{(s,r)} = (\vec{M}_{(s,r)}, \vec{m}_{(s,r)}, \vec{N}_{(s,r)})$ at r as a point of $\mathcal{O}_s$ consist of vector components now bound to r and, once transformed by R into the common image space coordinate system, permit the computation of inner products $\langle \vec{M}_r, \vec{M}_{(s,r)} \rangle, \langle \vec{N}_r, \vec{N}_{(s,r)} \rangle$.

3.3 Updating local information

The mapping $(\phi_s, U_s, \xi_s) \mapsto (\phi_{(s,r)}, U_{(s,r)}, \xi_{(s,r)})$ transports the principal curvatures and directions along the left-hand side of the above diagram from supporting point s to the point r being refined. Curvatures computed along the right-hand side say what the current estimated data is at r. The contextual neighbourhood of r consisting of k points, say, thus contributes k estimates of the local information at r. We now use this support $\left\{(\phi_{(s,r)}, U_{(s,r)}, \xi_{(s,r)})^i, s \in \mathcal{N}_r\right\}$ to map $(\phi_r, U_r, \xi_r)^i \mapsto (\phi_r, U_r, \xi_r)^{i+1}$, producing a better local estimate of the principal direction chart at r which contributes to updating the surface $\mathcal{S}^i \mapsto \mathcal{S}^{i+1}$. Briefly, the new chart is the one that best fits, in the natural least-squares error sense, what the supporting points indicate that the principal chart at r "ought" to be. Some care must be taken with the fit due to the possibility of umbilic points (where $\kappa = \lambda$), since no unique principal direction frame exists there. However, the normal direction remains valid for the computation of the best fit $\vec{N}$. Thus we decompose the fit into two steps:

1. all k surface normals from $\mathcal{N}_r$ are used to determine $\vec{N}$, the surface normal component of the new frame; then
2. the $\hat{k}$ non-umbilic points 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 2 shows the increasing support of a representative r from the MR image in Fig. 1 with its neighbouring support points over the course of several iterations.

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 (this assumption is supported in [7]).

Let us write the k charts from the $s_\alpha \in \mathcal{N}_r, \alpha = 1, \ldots, k$ as $(\phi_{(\alpha,r)}, U_{(\alpha,r)}, \xi_{(\alpha,r)})$, with $(N_{(\alpha,r)x}, N_{(\alpha,r)y}, N_{(\alpha,r)z}) = \vec{N}_{(\alpha,r)} \in \xi_{(\alpha,r)}$ being the components of the surface normal, etc. Step 1 above determines their best fit unit normal $\vec{N}$ as

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

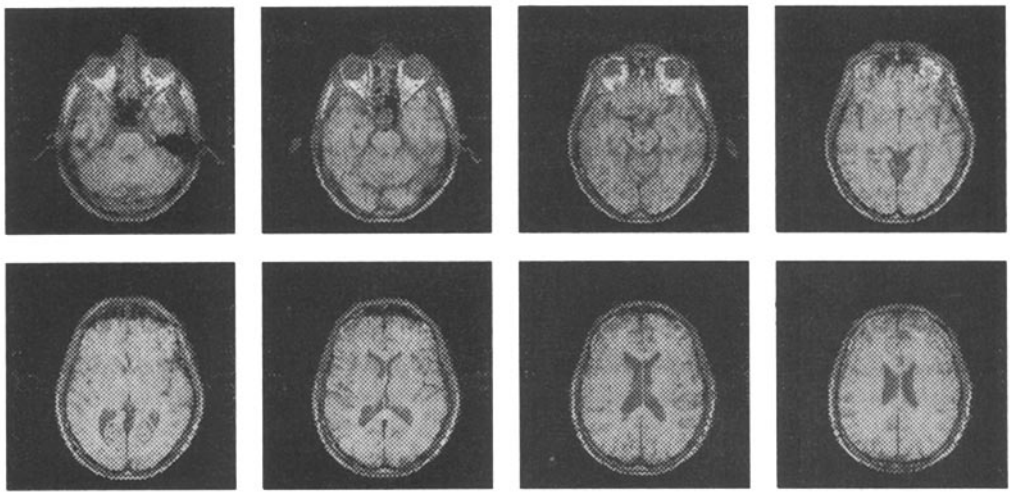

Figure 1: Eight images from a sequence of 41 slices from an MR image of the head. The images are 256×256 pixels in the xy plane, with a resolution of 0.2 cm. in the z direction. (Courtesy Dr. T. Peters, Montreal Neurological Institute.)

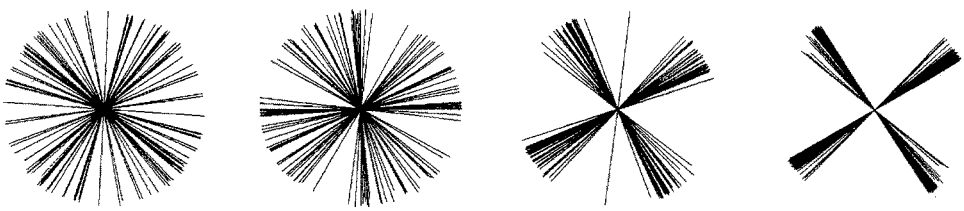

Figure 2: Increasing support over the contextual neighbourhood for the two principal directions at a typical point of the MR image. Initial estimate, one, three, and five iterations.

where

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

by minimizing

$$E_N^2 = \sum_{\alpha=1}^{k} \left\| \vec{N} - \vec{N}_{(\alpha,r)} \right\|^2$$

subject to the constraint $\|\vec{N}\| = 1$ (by the method of Lagrange multipliers, see [7] for details). 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

$$E_M^2 = \sum_{\alpha=1}^{\hat{k}} \left\| \vec{M} - \vec{M}_{(\alpha,r)} \right\|^2$$

subject to the constraints $\langle \vec{M}, \vec{M} \rangle = 1$ and $\langle \vec{M}, \vec{N} \rangle = 0$. This gives

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

where

$$\begin{aligned} \lambda_G &= 2(N_x \sum_{\alpha=1}^{\hat{k}} M_{(\alpha,r)x} + N_y \sum_{\alpha=1}^{\hat{k}} M_{(\alpha,r)y} + N_z \sum_{\alpha=1}^{\hat{k}} M_{(\alpha,r)z}), \\ d &= \left(\sum_{\alpha=1}^{\hat{k}} M_{(\alpha,r)x} \right)^2 + \left(\sum_{\alpha=1}^{\hat{k}} M_{(\alpha,r)y} \right)^2 + \left(\sum_{\alpha=1}^{\hat{k}} M_{(\alpha,r)z} \right)^2. \end{aligned}$$

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

$$E_\kappa^2 = \sum_{\alpha=1}^{k} \left[\frac{\left(\kappa - \kappa_{(\alpha,r)}\right)^2}{\max(|\kappa - \kappa_{(\alpha,r)}|)} + \frac{\left(\lambda - \lambda_{(\alpha,r)}\right)^2}{\max(|\lambda - \lambda_{(\alpha,r)}|)} \right]$$

of the best fit principal curvatures κ, λ, so that

$$E^2((\phi_r, U_r, \xi_r)^i) = E_N^2 + E_M^2 + E_\kappa^2$$

gives an "overall squared residual" of the fit of the principal chart at r.

It is evident that when E^2 is small, the contextual neighbourhood of r is strongly supportive of the estimated chart there, so we take

$$\Phi(\{(\phi_r, U_r, \xi_r)^i, r \in \mathcal{S}^i\}) = \sum_{r \in \mathcal{S}^i} E^2$$

as the functional to be minimized by updating the surface $\mathcal{S}^i \mapsto \mathcal{S}^{i+1}$.

4 Experiments

The Gaussian curvature at a point of a surface is the product $K = \kappa\lambda$ of the principal curvatures, and its sign can be used to segment the surface into regions of elliptic ($K > 0$) and hyperbolic ($K < 0$) points. Figures 3, 4 show the iterative improvement of the Gaussian curvature segmentation in the face and lateral venticle regions of the MR image (lateral venticles are internal structures in the brain seen as the dark grey areas near the centre of the images in the lower panel of Fig. 1).

Figures 5, 6 show the progressive improvement in the principal direction estimates over the course of several iterations until (approximate) convergence of the algorithm.

5 Discussion

In this paper, we have shown that surface trace, tangent and curvature can indeed be computed from digital images. However, that most commonly used tool of differential

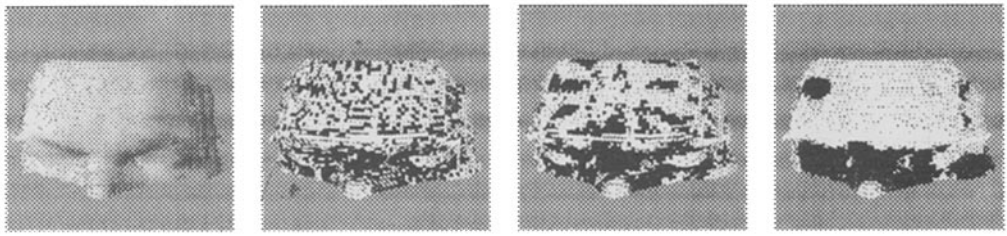

Figure 3: Coarse segmentation by Gaussian curvature classes of the face region of the MR image. Bright points are elliptic and dark are hyperbolic. Shaded tangent plane orientations, initial segmentation, one, and seven refinement iterations.

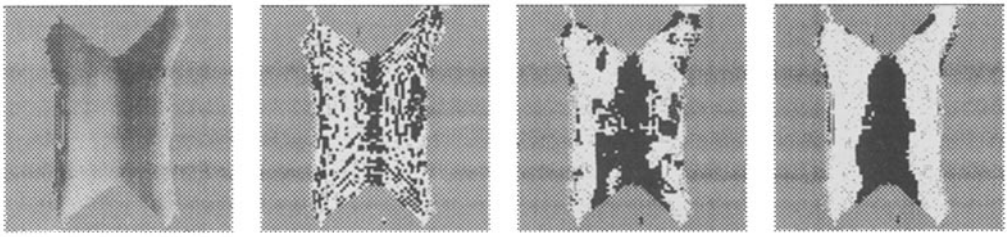

Figure 4: Curvature labeling of the top surface of the lateral ventricles. Shaded tangent plane orientations, initial labeling, two, and five refinement iterations.

geometry, differentiation, is insufficient for the task, and other notions must be exploited as well. Specifically, the problem is how to develop a *connection* between different points, how to relate estimated information at different estimated surface points.

We have presented an approach to this problem by considering the surface as a collection of local surface patches with smoothness constraints between overlapping patches imposing curvature consistency over neighbourhoods. Our methods are implemented as an iterative functional minimization process seeking a surface whose information at each point is maximally consistent with structure over a local neighbourhood.

We presented experiments showing that this approach works well in computing surface structure from clinical three-dimensional magnetic resonance imagery.

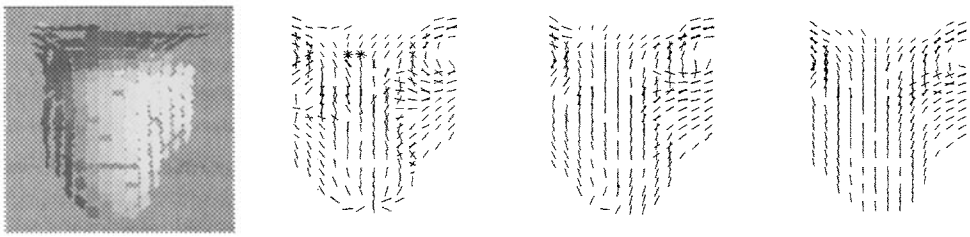

Figure 5: One of the principal direction field in the nose region of the MR image. Shaded tangent plane orientations, initial estimates, iterations one, and five. Apparent multiple principal directions are due to projection.

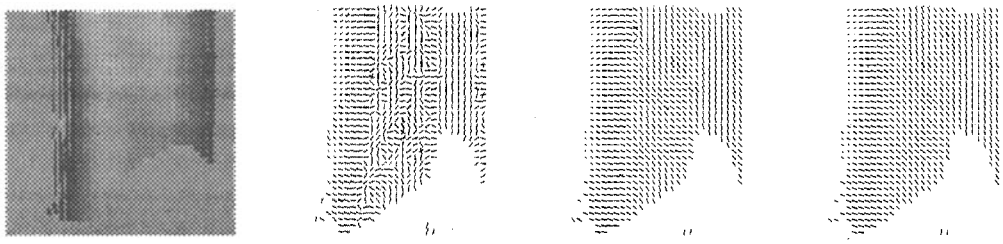

Figure 6: Principal direction field on a region of the surface of the lateral ventricles. Shaded tangent plane orientations, initial field, two, and three iterations.

References

[1] 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.

[2] Paul Besl and Ramesh Jain. Intrinsic and extrinsic surface characteristics. In *IEEE Proceedings on Computer Vision and Pattern Recognition*, pages 226–233, June 1985.

[3] 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.

[4] 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.

[5] Frank P. Ferrie, Jean Lagarde, and Peter Whaite. Recovery of volumetric object descriptions from laser rangefinder images. These proceedings.

[6] Manfredo P. do Carmo. *Differential Geometry of Curves and Surfaces.* Prentice-Hall, Englewood Cliffs, 1976.

[7] Peter T. Sander and Steven W. Zucker. Inferring surface trace and differential structure from 3-D images. *IEEE Transactions on Pattern Analysis and Machine Intelligence.* To appear.

[8] 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.