

Shape Priors in Medical Image Analysis: Extensions of the Level Set Method

Technical Report MS-CIS-02-08

Author: Albert A. Montillo
Department of Computer and Information Science
University of Pennsylvania

November 19, 2000

Abstract

The 3D medical image segmentation problem typically involves assigning labels to 3D pixels, called voxels, which comprise a given medical volume. In its simplest form the segmentation problem involves assigning the labels “part of the structure of interest” or “not part of the structure” to each voxel using locally measured properties and prior knowledge of human anatomy. Robust segmentation remains an open research problem today due to the significant challenges in the task including: partial volume averaging, overlapping intensity distributions and image noise. In the face of these challenges prior knowledge needs to be added to make the segmentation methods more robust. Active contours were introduced in the late 1980's mainly to address situations in which the object to be segmented had a single closed boundary. To address situations in which the object(s) to be segmented have unknown topology the level set framework was recently introduced to segment medical images. Unlike active contours, the level set method relies on an implicit shape representation rather than an explicit shape representation and hence new methods to impose prior knowledge about expected shape have to be devised for the new framework. This paper explores recent segmentation methods from four research groups which address the task of imposing prior knowledge of shape for object boundary segmentation. Three of the methods impose priors onto the level set technique and one employs a medial axis shape representation and Statistical shape information to guide a model-based segmentation. All of the methods include a notion of a statistical shape distribution. Each method is described, analyzed for its strengths and weaknesses. The paper concludes with a comparison of all four methods and recommendations for their applicability.

Table Of Contents

1	Introduction.....	4
2	Background.....	5
2.1	Manual vs Automatic Segmentation.....	5
2.2	Evolution of Segmentation Methods towards Shape Priors.....	5
2.3	Active Contour Methods.....	6
2.3.1	Lagrangian Active Contours.....	6
2.3.2	Eulerian Active Contours.....	7
2.3.2.1	The Speed Function is Application Dependent.....	9
3	The Coupled Surface Propagation Method.....	10
3.1	The Problem Domain.....	10
3.2	Overview of the Method.....	10
3.3	Two surfaces are Better Than One.....	11
3.4	Deriving Image Information to Guide the Level Set Evolution.....	11
3.5	Inner and Outer Surface Evolution.....	13
3.6	The Speed Term Couples the Surfaces.....	13
3.7	Results and Critical Analysis.....	15
3.7.1	Critical Analysis of the Method.....	18
4	A Method to Maintain Distance Map Integrity during Level Set Evolution.....	20
4.1	Theoretical and Practical Reasons to Maintain the Distance Map.....	22
4.2	Deriving a New PDE which Maintains the Distance Map.....	22
4.3	An Implementation Strategy for the New PDE.....	25
4.3.1	The Discrete Update Equation.....	25
4.3.2	Handling Discrete Sampling.....	25
4.3.3	Computing \mathbf{b} at Voxels Adjacent to the Zero Level Set.....	26
4.3.4	Computing \mathbf{b} on the Whole Narrow Band.....	26
4.4	Results and Critical Analysis of the Method.....	28
4.4.1	Revisiting the Coupled Surface Propagation.....	28
4.4.2	Extracting the Medial Manifold of an Evolving Surface.....	29
5	A Method to Guide Geodesic Active Contour Evolution using a Statistical Shape Model.....	32
5.1	The Shape Model.....	32
5.1.1	Shape Representation.....	32
5.1.2	Addressing the Correspondence Problem.....	36
5.2	Geodesic Active contours.....	36
5.3	Evolving the Surface with Shape Influence.....	37
5.4	Estimating the Final Shape and Pose.....	38
5.4.1	Inside Term: $P(u \mathbf{a}, p)$	40
5.4.2	Gradient Term: $P(\nabla I \mathbf{a}, p, u)$	40
5.4.3	Shape and Pose Terms: $P(\mathbf{a})$ and $P(p)$	41
5.5	Results.....	41
5.6	Critical Analysis of the Results and Method.....	43
6	A Model-based Segmentation Method Using Medial Axis Shape Representation.....	45
6.1	The Shape Model.....	45
6.1.1	Gross Shape Characteristics are represented using Medial and End Primitives.....	45
6.1.2	Coarse to Fine Representation.....	46
6.1.3	Boundary Primitives.....	46
6.1.4	Manual Construction of the Shape Model.....	47

6.2	Segmenting Novel Structures Using the Shape Model.....	47
6.3	Critical Analysis of Results and Method.....	48
6.3.1	Results.....	48
6.3.2	A Shape's Deformable Shape Loci Representation is Not Unique	49
6.3.3	Segmentation Limited to Shapes with Deterministic Topology.....	49
6.3.4	Psychophysical Basis.....	49
6.3.5	Space and Time are $O(N)$ but N is Increases with Accuracy Used in Shape Model.....	50
7	General Conclusions	51
7.1	Comparison of the methods.....	51
7.1.1	Pizer's Method vs Leventon's Method.....	51
7.1.2	Overall Comparison.....	51
8	References.....	52

1 Introduction

In this paper three methods [Zeng98, Gomes00, Leventon00] which are described and analyzed use the level set method to segment anatomical structures from a medical image. This segmentation problem is to label the voxels that comprise the image with the label “part of the structure of interest” or the label “not part of the structure” using measured properties such as tissue magnetization and anatomical knowledge of normal subjects. For comparison with the level set methods a different segmentation method [Pizer99] is also analyzed in which the medial axis shape representation is used to guide a model-based segmentation. PDE methods such as the level set method have been used by several researchers for a variety of purposes not just medical image segmentation including: registration of anatomical data sets, shape from stereo and the computation of the medial manifold of a shape. New applications are constantly being found and there are texts and conferences dedicated to the development of the theory and application of these PDE methods.

For the purposes of discussion, this paper assumes that the methods are applied to MR images although in many cases the methods can be used with little or no modification in other modalities. MR is a widely used medical imaging modality because it is safe yet produces high resolution and high contrast images of key soft tissue structures including the brain and soft organs.

Segmentation of structures in medical modalities is a particularly challenging due to:

1. Partial volume averaging
 - For manipulation in a computer, the images are spatially sampled. The resolution of MR is typically on the order of 1mm per side of a cube-shaped voxel. In MR, the value assigned to a voxel comes from the average magnetization measured in that voxel of tissue. When there is more than one tissue in the voxel the value measured is an average of the magnetization of the tissues in the voxels. Partial voluming can lead to missing or incomplete tissue boundaries.
2. Overlapping intensity distributions
 - The material properties of many structures in the body are similar; as a consequence, the measurements of many structures have overlapping intensity distributions. Also many structures of interested are composed of multiple tissue types, reducing the distance in intensities between different structures..
3. Image Noise
 - The signal can be degraded by motion artifacts from patient movement during the scan and in MR, images include nonlinear gain artifacts due to inhomogeneities in the transmitter or receiver coils.

One of the most important structures in the human body is the cortical layer and two of the methods [Zeng98, Gomes00] analyzed target this structure in particular. These methods impose a global shape constraint by using two evolving surfaces coupled to extract volumetric layers of constant thickness. However, the majority of anatomical structures are not volumetric layers of nearly constant thickness; the other methods analyzed [Leventon00, Pizer99] can segment a wider class of shapes. Leventon develops a method which drives the level set evolution using image data and a statistical shape model based on training shape vectors, while Pizer’s shape model uses a medial axis shape representation.

2 Background

2.1 Manual vs Automatic Segmentation

Segmentation methods can be roughly categorized into:

- manual segmentation methods
- semi-automatic segmentation methods
- automatic segmentation methods

Manual segmentation involves the tracing of contours by an expert anatomist. For example the expert may outline the boundary of a deep-seated brain tumor in MRI images and optic nerve so that a neurosurgeon can plan a safe surgical procedure preoperatively. Manual segmentation is typically achieved by the tracing the contours on each image in the “stack” from the volume acquisition. This procedure can take 3 – 18 hours depending on the complexity and number of objects to be traced and the resolution of the data. [Kelemen99, Zeng98]. To speed the process, *semi-automated* segmentation methods have been designed. The expert traces a rough contour and then presses the “snap” button or similar on the program’s interface. This causes the application of a local image-gradient-following algorithm to move the contour onto the closest image boundary. These algorithms employ only local information and hence they have a short capture range. Thus the expert’s contour must be fairly close to the actual contour or the contour will lock onto the wrong image boundary. However even this segmentation can take 2 hours. As a consequence, *fully automated segmentation* methods have been developed.

Fully automated methods can perform the segmentation slice by slice or by examining the full 3D data all at once to derive a bounding surface in 3D instead of a bounding contour in 2D. These methods offer several distinct advantages over manual and semi-automatic segmentation including:

- Segmentations are *consistent* and *objective* interpretations of the data
- Results can be obtained in a *fraction of the time* of manual or semi-automatic methods

It has been shown that expert segmentation can deviate by as much as 10% [Zeng98]. For certain applications such as in the assessment of the drug effectiveness it is vital to have consistent objective segmentations of the before and after images. Fully automated segmentation meets this need.

Moreover the reduction in segmentation time makes possible *many new exciting applications for medicine*. For example it enables mass screening of a population to reduce mortality rates for certain diseases and it facilitates computer assisted surgery and telesurgery and makes plausible, robotic surgery.

2.2 Evolution of Segmentation Methods towards Shape Priors

The complex segmentation methods under development today have evolved gradually from the inadequacies of simpler strategies. The most basic segmentation strategy used on MR images assigns voxel labels using only the intensity (magnetization) information at that voxel. An example of such a point level operation [Jain95] is *global thresholding*. While these methods are useful for structures with non overlapping intensity distributions such as a bone in an image of soft tissue, many structures have overlapping intensity distributions with the surrounding structures, and moreover in almost all cases partial voluming and image noise prevent point based methods from reliably segmenting many anatomical structures completely. Local level operations include information in the surround of the voxel. These include the boundary based methods and their complement, region based methods.

Boundary based methods seek to identify the boundary of the structure of interest using image gradient information. *Region based methods* attempt to identify the voxels which constitute the structure of interest by grouping voxels of similar properties. Hybrids of these methods have also

been tried. However the local level methods fail when the boundary or region information is incomplete due to overlapping intensity distributions, partial voluming and image noise.

To meet these challenges, global information about the structure to be segmented is used in the most recently developed segmentation methods. Often knowledge about the shape of the structure to be segmented is known a priori. To encode this shape information for the segmentation process, many shape representations have been used including, but by no means limited to: finite meshes, Fourier shape descriptors, and medial manifolds. By utilizing global shape information in conjunction with image data, overlapping intensity distributions, partial voluming and image noise are less apt to foil the segmentation. However, a single static shape template can rarely be used to segment anatomical structures reliably due to normal and pathological variability in subject anatomy. Therefore the global shape information must only be a guide for the probable shape and not used as a rigid template for the final segmentation results.

Methods using deformable models begin with an initial shape and pose information of the targeted structure and evolve this information so that it is both *faithful to the image data* and *satisfies internal and external shape constraints*. These evolution can be implemented as an energy minimization problem. Fidelity to image data is often measured as how well the proposed boundary approximates the image edge information. Internal shape constraints can be implemented as an energy term which is proportional to the curvature of the shape –which biases the final segmentation to smooth curves. External shape constraints can be implemented as fidelity of the segmented shape to a shape probability model – which biases the final segmented shape towards those which are more likely to occur given in shape distribution. Deformable models are also called active contours and two of the possible implementation strategies are described in next section.

2.3 Active Contour Methods

The development of active contour methods to simulate curve evolution for segmentation, tracking and registration has come from the efforts of many scientists and over a decade of research. Methods to evolve these contours were introduced to computer vision using a Lagrangian formulation by Kass, Witkin and Terzopoulos [Kass88] and then reformulated in the context of PDE-driven surfaces by Caselles, Kimmel and Sapiro in [Caselles95]. Research addressing the theoretical aspects of these PDE's including uniqueness and existence of solutions has been extensive. Level set methods were introduced by Osher and Sethian [Osher88], [Sethian99] as a theoretical framework and practical tool to solve the type of PDE's which govern curve evolution.

2.3.1 Lagrangian Active Contours

In one form of active contours, also called snakes (2D) or balloons (3D) the surface is modeled by a *set* of markers or particles of constant cardinality. The particle motion is modeled with a Lagrangian formulation. Below is a brief synopsis of the Lagrangian motion formulation in the 2D case.

- Given the parameterized evolving closed curve $x(s, t) = [x(s, t), y(s, t)]^T$ where s is arc length and t is time,
- Given the curve evolves with movement along the normal to the curve
- From these statements we have:
 - (0) the tangent vector to the curve at time t :

$$\hat{t}(s, t) = \left[\frac{x_s}{(x_s^2 + y_s^2)^{1/2}}, \frac{y_s}{(x_s^2 + y_s^2)^{1/2}} \right]^T \quad (\text{Intro1})$$

(1) $\hat{n}(s, t)$ is the unit normal vector at time t at point s on the curve and is the derivative of the tangent vector:

$$\hat{n}(s, t) = \frac{\hat{t}'(s, t)}{|\hat{t}'(s, t)|} \quad (\text{Intro2})$$

(2) $k(s, t)$ = the local curvature of the curve at time t at point s on the curve for a given time t :

$$k(s) = \frac{y_{ss}x_s - x_{ss}y_s}{(x_s^2 + y_s^2)^{3/2}} \Big|_s \quad (\text{Intro3})$$

Thus the motion of the particle is: $\bar{x}_t = F\hat{n}$ or in expanded form:

$$\bar{x}_t = \left[F(k(s)) \frac{y_s}{(x_s^2 + y_s^2)^{1/2}}, F(k(s)) \frac{x_s}{(x_s^2 + y_s^2)^{1/2}} \right]^T \quad (\text{Intro4})$$

There are several limitations if this method is used and they are difficult to overcome. One limitation is that the *curve cannot readily change topology* as it evolves. Consider the case in which the initial contour is a single small ellipse but the image contains 2 round objects side by side. If we expand a small ellipse starting from the gap in between these objects we will need to split the contour into two contours when it wraps around the two objects and meets itself on the other side. This means we need to end up more contours than were initialized. *Merging and splitting of contours is not handled* by the Lagrangian formulation since, inherent in the method, is that the number of contours does not change. In medical images there is often unknown topology. Another problem with the Lagrangian formulation is that there is a *loss of curve approximation fidelity as the contour expands* because the number of markers remains fixed. An alternative formulation, which overcomes these limitations, is an *Eulerian formulation*.

2.3.2 Eulerian Active Contours

The main idea is to embed the evolving boundary $C(t)$ as the zero level set of higher dimensional function $u(\bar{x}, t)$ which is defined as:

$$u(\bar{x}, t = 0) = \text{signed distance between } \bar{x} \text{ and } C(t = 0) \quad (\text{Intro5})$$

This means that u is a function which maps points in \mathfrak{R}^n (when tracking a n -D curve or surface) to a distance value: $u(\bar{x}, t = 0): \mathfrak{R}^n \rightarrow \mathfrak{R}$. In particular the value of u at \bar{x} is the distance to the boundary. Points outside the closed curve are defined to be a positive distance from the curve while points inside are a negative distance from the curve. Points on the front itself are those points such that the distance is zero:

$$C(t = 0) = \{\bar{x} | u(\bar{x}, t = 0) = 0\} \quad (\text{Intro6})$$

For example, if $C(t=0)$ is the unit circle, then $u(\bar{x}, t=0)$ is a circular cone with the point at $z = -1$ and intersecting the unit circle on the xy plane.

To track the boundary we can **evolve** $u(\bar{x}, t)$ over time, rather than $C(t)$. A cross-section of $u(\bar{x}, t)$ at time t , yields $C(t)$. Moreover since $u(\bar{x}, t)$ is simply the distance to $C(t)$, $u(\bar{x}, t)$ **will remain a surface** while its cross section in the xy plane can change topology. To derive an equation governing the evolution of $u(\bar{x}, t)$ we can look at the path, $\bar{x}(t)$, of a point on the moving surface: since this particle moves along the normal \bar{n} to the boundary we have the magnitude of the projection of the partial of \bar{x} with respect to t onto \bar{n} is the speed:

$$\bar{x}_t \cdot \bar{n} = F(\bar{x}(t)) \quad (\text{Intro7})$$

In other words the component of \bar{x}_t in the unit \bar{n} direction can be extracted using the dot product. Since the zero level set of the evolving function u must always match the propagating hypersurface we have :

$$u(\bar{x}(t), t) = 0 \quad (\text{Intro8})$$

Differentiating both sides with respect to t using the chain rule we get the evolution constraint for u :

$$\begin{aligned} \frac{d}{dt} \{u(\bar{x}(t), t)\} &= \frac{d}{dt} 0 \\ \frac{\partial u}{\partial \bar{x}} \frac{d\bar{x}}{dt} + \frac{\partial u}{\partial t} &= 0 \\ \left[\frac{\partial u}{\partial x_1} \quad \frac{\partial u}{\partial x_2} \right] \begin{bmatrix} \frac{dx_1}{dt} \\ \frac{dx_2}{dt} \end{bmatrix} + \frac{\partial u}{\partial t} &= 0 \quad (\text{Intro9}) \\ \frac{\partial u}{\partial x_1} \frac{dx_1}{dt} + \frac{\partial u}{\partial x_2} \frac{dx_2}{dt} + \frac{\partial u}{\partial t} &= 0 \\ \nabla u \cdot \bar{x}_t + u_t &= 0 \end{aligned}$$

The dependence on $\bar{x}(t)$ can be removed by noting that the unit normal vector is given by:

$$\bar{n} = \frac{\nabla u}{|\nabla u|} \quad (\text{Intro10})$$

and then substituting equation (Intro10) into equation (Intro7) yields:

$$\begin{aligned} \frac{d\bar{x}}{dt} \cdot \bar{n} &= F \\ \frac{d\bar{x}}{dt} \cdot \frac{\nabla u}{|\nabla u|} &= F \quad (\text{Intro11}) \\ \nabla u \cdot \bar{x}_t &= F |\nabla u| \end{aligned}$$

followed by substituting equation (Intro11) into equation (Intro9) removes the dependence and yields the level set equation:

$$u_t + F |\nabla u| = 0 \quad (\text{Intro12a})$$

and using a numeric method such as the Euler method yields the discrete form of the update equation:

$$u(t+1) = u(t) - F|\nabla u(t)| \quad (\text{Intro12b})$$

2.3.2.1 The Speed Function is Application Dependent

As shown in equation (Intro12a), the level set equation is a first order, non-linear, homogeneous partial differential equation which depends on the speed term, F. F is a function of:

- the local properties of the boundary such as curvature and normal direction
- the global properties of the boundary which depend on the shape and position of the front
- properties which are independent of the shape of the front such as the underlying “fluid” velocity that passively inflates the boundary and image information such as pixel intensity and gradient.

Thus the speed function links the evolution to the image and other shape properties.

3 The Coupled Surface Propagation Method

3.1 The Problem Domain

The most important organs in the human body are arguably the brain and the heart. In the brain the cortical layer is one of the most important structures for it contains the pyramidal cells which are generally thought to be those responsible for higher order thinking processes. In the heart the left ventricle is vital because this chamber is responsible for pumping blood throughout the body. Both the cortical layer and the myocardium of the left ventricle are volumetric structures of roughly constant thickness (much more roughly for the myocardium). Segmenting a thin volumetric layer by propagating *one* surface is inherently difficult. Besides the canonical segmentation challenges stated above there is the added complication of the presence of two nearly parallel edges in close proximity. Using just one surface, it is likely for the segmented surface to contain part of one boundary in other places and part of the other interface. Zeng leverages two properties of these structures to achieve more robust segmentation [Zeng98]. These properties are that the volumetric layer has roughly constant thickness throughout and that the material property contained within the layer is homogenous. To embed the thickness constraint into the method Zeng uses the level set method to evolve two surfaces simultaneously while maintaining a constant separation. One surface locates the outer bounding surface and the other the inner surface. To incorporate the homogeneity property into the algorithm, Zeng uses statistically likely gray level voxel intensities of the layer instead of just image gradient information to drive the surface propagation.

3.2 Overview of the Method

Both region and boundary-based segmentation methods have been applied to segment the cortical layer and myocardium with limited success. In [Cline90] a region-based method using multispectral data and connectivity is used to classify each voxel into different tissue types. A material mixture model is assumed for the region intensity distributions. However the method requires significant manual processing to group regions into a single anatomic structure. Moreover in practice multispectral data is unavailable due to the added time and/or cost for acquisition. Several attempts have been made to use boundary finding methods to segment these structures. Snakes and balloon models tend to be insufficient for cortical layer extraction because they require an initialization which is quite close to the final result. Deformable surface models based on the finite element method have been used to “shrink wrap” a surface to a skull-stripped 3D MR brain image in order to find the gray matter – cerebral spinal fluid (CSF) interface. Davatzikos et al. [Davatzikos96] developed a model which is targeted specifically at finding an individual sulcus. His ribbon model employed knowledge about the expected thickness of the sulcus and an active contour method for locating the spline of the ribbon and is extended to 3D. The deformable surface models and the ribbon model typically require human intervention to override local smoothness to force the surface down into the deep narrow folds of the sulci.


Zeng circumvents the initialization and coherency problems found in the traditional boundary and region based methods by modeling more properties of the volumetric layer to be segmented. Zeng recognizes that the layer is bounded by two surfaces which are separated by a finite, fairly constant, thickness and that there exists a homogeneity of certain voxel statistics in the region between the bounding surfaces. The method evolves two surfaces simultaneously via the level set approach. The evolution of each surface is driven by its own image-derived information and guided by the location of the companion surface. The image-derived information is based on a statistical measure of the homogeneity of the intensities of the voxels between the two evolving surfaces.

3.3 Two surfaces are Better Than One

Field inhomogeneities and partial volume averaging often create regions in which there is insufficient information to clearly define either the inner bounding surface or the outer surface. In these regions one clearly defined boundary helps to resolve the location of the other. For example in MR images of the brain, in some places the white matter-gray matter boundary is not well shown. Partial volume averaging blends the white and gray voxels while the CSF-gray matter boundary may still appear clearly. Using just a single expanding surface to locate the former boundary would cause the inner cortical surface expand into CSF. In other places such as where the eye sockets are, the CSF doesn't appear in the image. The bone of the eye socket appears with high intensity vales, so using a single surface expanded to capture the gray matter –CSF interface would cause the surface to expand into the eye socket. Using the coupled surface approach these errors can be avoided: a minimum and maximum distance between the inner and outer surfaces is maintained during propagation and thus one well defined boundary anchors the possible locations for the other. This “information sharing” or *coupling* of the surfaces during evolution is effected through the design of speed terms dependent on the surface separation.

3.4 Deriving Image Information to Guide the Level Set Evolution

In the coupled level set approach a local image operator is defined which measures the likelihood that a voxel lies on the boundary between two tissue types A and B. The surface is assumed to have finite curvature therefore, at a fine scale, the local orientation of the surface can be approximated by a tangent plane. In figure B1 a plane through the point S is shown. This plane divides a small cube of voxels around S into two regions R1 and R2. The plane can be



[Zeng98] Figure B1: Local Image Operator

parameterized by its orientation \vec{q} , for example by latitude and longitude angles on the unit sphere. The probability that \vec{q} defines a plane through S on the boundary between surface A and B is the probability that $(R_1 \in A \text{ and } R_2 \in B)$ or $(R_1 \in B \text{ and } R_2 \in A)$. In this level set method the initial surfaces must be closed and thus the notion of inside and outside the front is well defined. When we know that A is always on one side –say the outside of the evolving surface, then we can reduce the event to one case: $(R_1 \in A \text{ and } R_2 \in B)$. The probability of this event occurring can be expressed as:

$$p_{AB}(\vec{q}) = p(R_1 \in \text{TissueA})p(R_2 \in \text{TissueB}) \quad (\text{B1})$$

Let \vec{q}^* be the estimate of the normal direction of the possible boundary through S. Then $p(\vec{q}^*)$ is the estimate of the probability that S lies on the boundary between A and B. To find \vec{q}^* the method first generates the vector $P = [p(\vec{q}_1), p(\vec{q}_2), \dots, p(\vec{q}_k)]^T$ and \vec{q}^* is the angle vector, \vec{q}_i , $1 \leq i \leq k$, which corresponds to the largest element in P. Zeng does not say how to select the \vec{q}_i however one possible

method is to discretize the space of possible orientations by the orientations of the facets of an icosahedron:





Figure B1b: Nearly regular tessellations of the sphere with icosahedrons. The icosahedron subdivision illustrates the tunable discretizations for generating the \vec{q}_i

The method assumes that the intensity distribution for each tissue can be modeled as a Gaussian. Let $N(\mu_a, \sigma_a)$ be the Gaussian distribution for tissue A and $N(\mu_b, \sigma_b)$ be the Gaussian distribution for tissue B. Assuming that voxels in tissue A are independently drawn from $N(\mu_a, \sigma_a)$ and the voxels in tissue B are independently drawn from $N(\mu_b, \sigma_b)$ then the probability that \vec{q} defines the orientation of a plane through S separating two tissues A and B is

$$p_{AB}(\vec{q}) = \left(\prod_{r \in R_1} \frac{1}{\sqrt{2\pi} \sigma_A} \exp\left(-\frac{(x_r - \mu_A)^2}{2\sigma_A^2}\right) \right) \left(\prod_{r \in R_2} \frac{1}{\sqrt{2\pi} \sigma_B} \exp\left(-\frac{(x_r - \mu_B)^2}{2\sigma_B^2}\right) \right) \quad (\text{B2})$$

Results of applying this operator to MR images of the brain can be seen in figure B2. (a) and (e) are



[Zeng98] Figure B2: Results of applying the homogeneity operator to MR images of the brain

the original axial and coronal slices. To the right of each is the magnitude of the gradient image; below on the left is the probability of the gray matter-white matter interface and below on the right is the probability of the CSF-gray matter interface.

3.5 Inner and Outer Surface Evolution

The level set method is used to evolve initial inner and outer surfaces to their final locations. In this Eulerian formulation for the motion, each surface propagates along its normal direction with speed F :

$$u_{in_t} + F_{in} |\nabla u_{in}| = 0 \quad (B4)$$

$$u_{out_t} + F_{out} |\nabla u_{out}| = 0 \quad (B5)$$


where u_{in} , u_{out} and F_{in} and F_{out} are the distance maps and speed terms for the inner and outer distance maps respectively. It is assumed that the bounding surfaces are initially completely inside the inner bounding surface and with an offset in between.

3.6 The Speed Term Couples the Surfaces

The speed term in equation (B4) and equation (B5) serves both to link the evolution to the image data and to couple the evolution. The speed terms are designed so the interfaces propagate along the normal direction and stop at the desired place while maintaining an acceptable distance

between them. Where the distance between the interfaces is within the normal range the image-derived information regulates the propagation and where the distance between the interfaces is outside the normal range the propagation is constrained.

Use of the level set method for front propagation affords an easy way to compute the distances between the interfaces. For any point p on the inner moving interface the distance to the outer surface is $|u_{out}(p)|$. For any point q on the outer moving surface the distance to the inner surface is $|u_{in}(q)|$. The later relationship is illustrated in figure B2a.




[Montillo-00] Figure B2a: The distance map for $u_{in}(q)=d$ gives the distance d from the inner surface to the outer surface at q

The speed is modeled as:

$$\begin{aligned} F_{in} &= g(p(\vec{q}^*))h(u_{out}) \\ F_{out} &= g(p(\vec{q}^*))h(u_{in}) \end{aligned} \tag{B6), (B7)}$$

where the functions g and h are functions which are shaped like the curves shown in figure B3.



[Zeng98] Figure B3: the functions of the speed term


The g function reduces the front propagation speed where the probability is high that the front is at the interface based on the image information. The g function stops the movement when the probability reaches $P_{\text{threshold}}$ which is near 1.0. Where the probability is low the g function tends toward a positive constant. The h function reduces the front propagation speed where the distance to the companion curve (volumetric layer thickness) reaches the limits of the normal range: $[T_{\text{min}}, T_{\text{max}}]$. The h function stops the movement when the distance is either too great (the layer would become too thick) or too small (the layer would become too thin). Since the h and g terms are multiplied, the front propagation stops when *either* term is zero. While the speed term only has this meaning on the front it is extended to the other points in the distance map using the level set derivation:

$$u_{in}(t+1) = u_{in}(t) - g(p(\bar{\mathbf{q}}^*))h(u_{out})|\nabla u_{in}(t)| \quad (\text{B8})$$

$$u_{out}(t+1) = u_{out}(t) - g(p(\bar{\mathbf{q}}^*))h(u_{in})|\nabla u_{out}(t)| \quad (\text{B9})$$

3.7 Results and Critical Analysis

Zeng reports the results of a synthetic test which illustrates the advantage of employing coupled surfaces to capture a layer of finite thickness. In figure B4 a synthetic image of a ribbon is shown. This ribbon is the mid-level intensity which surrounds the center white region. In the image the ribbon is of constant thickness in all places except the folds in the left and top of the figure in (a).




[Zeng98] Figure B4: Synthetic test illustrates the advantage coupled surface propagation

Image (b) shows the initial location of the inner and outer boundary surfaces. In (d) the method has converged. It has captured the gray region with a ribbon of constant thickness.


The coupled surface segmentation method has also been tested on volumetric MR images of the brain. Fourteen subjects were imaged using a GE 1.5T scanner and a 3D SPGR pulse sequence. These volumes contained voxels which measured 1.2mm on a side. The images were intensity normalized to reduce the effects of bias field inhomogeneities using a standard non-linear map. To segment a volumetric layer in a 3D data set, the method requires 3D initial surfaces. Zeng defined the initial surfaces to be several small spheres inside the brain volumes. Each initial sphere was composed of an outer and inner sphere. In this test the inner and outer spheres are set to be 3.5mm apart and allowed to separate during evolution by as much as -20% to $+30\%$ that distance. The outer spheres expand to the gray matter – CSF interface while the inner spheres expand to the gray matter –

white matter interface. Figure B5 shows the initial outer and inner spheres in the top row, two intermediate steps in the evolution in the next two rows and the final interfaces in the bottom row.



[Zeng98] Figure B5: Propagation of the outer (left column) and inner (right column) surfaces.

We can see qualitatively from the middle two rows, that the level set method is handling merging without a problem as the level set theory suggests it would. Moreover images like those shown in figure B6 give further qualitative evidence of the correctness of the final segmentation of the cortical layer.



[Zeng98] Figure B6: Qualitative evidence of correctness of the method


Quantitative validation is preferred for this critical segmentation however and Zeng gives a numerical measure of the accuracy of the method by comparing the individual voxels labeled by the method with those labeled manually by expert neuroanatomists who traced the inner and outer boundaries slice by slice. To compare the coupled surface segmentation to the manually segmentation, the number of false positive voxels and the number of false negative voxels were counted. A false positive voxel is a voxel which was included in the coupled surface layer but is not in the manual layer. A false negative voxel is a voxel which is missing from the coupled surface layer but is in the manual layer. The *true positive rate* is the number of voxels in common normalized by the number in the manually segmented volume (a.k.a. 1- (false negative rate)). The *false negative rate* is the number of false positives normalized by the number of voxels in the manually segmented volume. As shown in table B1, in the first 2 columns the *whole brain* volumes are compared. Zeng defines the whole brain volume as all voxels inside the outer bounding surface (the inner surface is ignored). Thus the whole brain volume includes structures such as the brain stem and the corpus callosum. Including these structures. The whole brain true positive rate ranges from 93.4% to 97% for the 14 brains tested.

brain volume TP rate(%)	brain volume FP rate(%)	cortex TP rate(%) (axial)	cortex TP rate(%) (coronal)
94.3	3.2	86.9	86.6
93.4	3.8	87.1	85.9
95.3	4.0	87.0	85.6
95.4	3.3	86.2	85.7
94.6	5.1	83.7	82.6
97.2	5.1	88.2	86.8
95.3	3.9	87.3	86.2
95.5	3.7	86.6	85.3
95.6	3.8	88.6	85.6
95.0	3.4	87.9	88.8
94.0	2.0	89.7	89.5
94.5	4.8	87.0	84.4
95.3	4.0	84.7	85.4
97.2	3.9	87.0	87.8

[Zeng98] Table B1: Quantitative comparison of whole brain and cortical layer error rates. Whole brain error rates in the first two columns; cortical layer errors in the right most two columns.

Zeng adds that, qualitatively, the missing portions of the brain are mainly regions in the brain stem and cerebellum where the constant thickness assumptions is violated. The false positive rate is 2%-5.1% which is partly due to the variability in expert tracing.

Comparing only the cortical gray matter segmentation reveals similar though somewhat less accurate results. In the third and fourth columns of the table, the results are counted from an axial slice and coronal slice containing only the cortical layer. The true positive rate varies from 83.6% to 89.7% which is similar to the repeatability of the manually tracing of the cortical layer by experts: 90%-95%. Figure B8 shows the cortical gray slices compared which were counted. In (a) the original



[Zeng98] Figure B8: Axial and Coronal slices used to compute error rates for the cortical layer

slices are shown; in (b) the manually traced volumes and in (c) the coupled surface volumes.

In many neurological studies, total brain volume and the volume of the constituent parts are compared among different subjects. To get a sense of how reliable the algorithm's volume measurements are, it is important to compare how large the coupled surface volume error is relative to the deviation in the volume when different experts segment the same brain or when the same expert re-segments the same brain. In one study Zeng reports that whole brain volume on 15 subjects by two independent expert tracers produced very low error- average volumes differed by 0.6%. When whole brain volumes measured by the algorithm were compared to the expert tracings average volumes differed by 1.35%. For the gray matter volume alone, a single expert was able to reproduce the volume across seven subjects with an average of 3.1% difference in volume. The algorithm's gray matter volumes differed from the tracings by 4.4% on the average. Thus the discrepancy between the algorithm's volumes and the expert traced volumes is roughly equal to the discrepancy that an expert would find if he/she retraced the same brain. Therefore the algorithm can be considered quite reliable for studies which are whole brain volume or gray matter volume measures.

3.7.1 Critical Analysis of the Method

Advantages Of the Method

Use of the level set and coupled surface and coupled surface propagation enabled tight folds in the cortical layer to be captured accurately without human intervention. This is a significant improvement over earlier attempts by [Davatzikos] and others employed the use of a deformable model which collapsed around the cortex and needed to be manually guided down into the deep sulci.

The initial bounding surfaces can be placed anywhere inside the subject's brain volume. In practice this is an easy task. Regardless of which interior location is chosen, the algorithm will still generate the same segmentation. This stems from the use of the distance map $u(x,t)$ to represent the evolving surface. $u(x,t)$ always remains a single-valued function, however the propagating hypersurface $C(t)$ may change topology, break, merge and form sharp corners as u evolves. Thus a key advantage of an implicit representation of shape over an explicit one at least in this application is robustness with respect to initialization.

Another advantage is that the distance map can be used to compute geometric properties of the level set surface:

- The normal to the surface is given by: $\vec{n}(x) = \nabla u(x)$
- The curvature of the surface is the divergence of the gradient of the unit normal vector to the front:

$$k = \nabla \cdot \frac{\nabla u}{|\nabla u|}$$

- The thickness map between the inner and outer surfaces can be computed from u_{in} and u_{out} . The cortical thickness map is useful to detect neurological diseases which cause the layer to atrophy.

The method is clearly faster than manual tracing: the manual tracing required a minimum of 9-12 hours per brain while the coupled surface method required only about 1/10 th that time on a R10000 SGI Indigo2.

Limitations of the Method

The method is not without its limitations. In regions where *both* boundaries are obscured by partial voluming or field inhomogeneities the level set method is still prone to “leak through” into surrounding structures.

The method is limited to volumetric layers of fairly constant thickness. Where this assumption is violated by the structure the method will produce an inaccurate segmentation. Even if presented with a training set of hand segmented cortical layers, it is not readily apparent how one would override this constraint in areas likely to have the violation. The method by Leventon presented in a later section in this paper gives some ideas along the lines of a statistically guided level set evolution.

The quantification of the errors based on true positive and false negative rates and the volume comparisons do not capture a notion of where the errors occurred or how sizeable any given error region is. For example one cannot discern from these numbers whether all of the voxels in error were spread evenly throughout the volume or clumped into one place. Also, it is not easy to compute from these numbers what the maximum distance of separation was between the manually identified boundaries and the corresponding coupled surface identified boundary. These metrics should be used so that these questions can be resolved. In general in order to compare competing algorithms from different researchers a suite of bench marks needs to be identified and used consistently among researchers. For example Leventon uses the partial hausdorff distance but not volumetric voxel true positive and false negative rates.

Other Comments Concerning the Level Set Method In General

- Tunable precision can be achieved by adjusting the order (number of taylor series terms retained) in the numeric method used to solve the PDE. The precision can also be adjusted by changing the grid spacing step used to compute u and the time step dt for successive iterations
- Fast numeric solutions are possible by computing only the portion of u near the xy plane.
- The method can be readily extended to even higher dimensions

4 A Method to Maintain Distance Map Integrity during Level Set Evolution

Like Zeng, Gomes is also concerned with the simulation of the evolution of a surface represented implicitly and governed by a partial differential equation. As in the coupled surface propagation, Gomes also chooses the distance map as the implicit surface representation, however in [Gomes99] the primary focus is on the reformulation of the level set methods introduced by Osher and Sethian in [Osher88] and in demonstrating the theoretical and practical implications of the reformulation.

Consider the family of hypersurfaces $S(\mathbf{p}, t)$ in \mathfrak{R}^3 , in which \mathbf{p} parameterizes the surface S and t is time, that evolve along the direction normal to the boundary. This evolution can be expressed as the PDE:

$$\frac{\partial S}{\partial t} = \mathbf{b}(\bar{x})\hat{n} \quad (C1)$$

where \hat{n} is the *inward* unit normal vector of S and $\mathbf{b}(\bar{x})$ is the speed of the evolution at point $\bar{x} \in S$ and with the initial condition: $S(t=0) = S_0$ where S_0 is some initial closed surface. Level set methods were introduced by Osher and Sethian [Osher88], as a theoretical framework and practical tool to solve the type of PDE's which govern curve evolution. An idea introduced in [Osher88] is the use of an embedding function u . When used in the evolution of the surface S in \mathfrak{R}^3 this map can be expressed as $u: \mathfrak{R}^3 \times \mathfrak{R} \rightarrow \mathfrak{R}$, that is the function maps a point in \mathfrak{R}^3 at a given time to a measure in \mathfrak{R} . Osher and Sethian proposed the criterion for the embedding function:

$$u(S, t) = 0 \quad \forall t$$

The derivation of the PDE which governs the evolution of u was described in the Background section of this paper. The PDE derived there is (with F renamed as \mathbf{b}):

$$\frac{\partial u(\bar{x}, t)}{\partial t} = \mathbf{b}(\bar{x})|\nabla u(\bar{x}, t)| \quad (C2)$$

This first order PDE is called a Hamilton-Jacobi equation because the unknown function u does not appear explicitly in the equation. The initial condition for the PDE in equation (C3) is

$$u(S, t=0) = u_0(\bar{x}) \quad (C3)$$

Where u_0 is an initial function $\mathfrak{R}^3 \rightarrow \mathfrak{R}$ such that $u_0(S_0) = 0$ to satisfy equation (C2). It has been proven that for a large class of functions u and u_0 , the zero level set at time t of equation (C3) is also the solution to equation (C1) at t . But this need not always be the case since we are free to choose u to be any embedding function of the form $u: \mathfrak{R}^3 \times \mathfrak{R} \rightarrow \mathfrak{R}$. The most often chosen embedding function is the signed distance map to S_0 . This map can be characterized by the following two equations:

$$\{\bar{x} \in \mathfrak{R}^3, u_0(\bar{x}) = 0\} = S_0 \text{ and } \nabla u_0(\bar{x}) = 1 \quad \forall \bar{x} \in \mathfrak{R}^3$$

The later condition can be understood when you consider that $|\nabla u_0(\bar{x})|$ is equal to the magnitude of the distance function from S_0 in the direction normal to S_0 which has a magnitude of one.


Use of the distance map is an intuitive choice and it facilitates the creation of practical algorithms such as Zeng's coupled surface propagation method, however it has been proven in [Barles93] that when the signed distance map is chosen as the embedding function, the solution u to

equation (C3) is not the signed distance function of the solution S of equation (C1). This causes several problems which are addressed by Gomes.

The first problem is that as u is evolved in accordance with equation (C3) it begins to over estimate the distance to the surface S at points not on S. This causes the need to recompute the distance map periodically in order to maintain such information during evolution.

The second problem is that the equation (C3) requires the speed function to be defined for all points $\bar{x} \in \mathfrak{R}^3$ while equation (C1) only requires the speed function to be defined on the surface S. Therefore there is the need to extend the definition of \mathbf{b} from S to the domain of u , namely \mathfrak{R}^3 . Depending on how one defines \mathbf{b} on S, the extension to \mathfrak{R}^3 may or may not be natural. If $\mathbf{b} = H_S$ the mean curvature of S then one may choose $\mathbf{b}(\bar{x}) = H_{u(\bar{x})}$ where $H_{u(\bar{x})}$ is the mean curvature of the isosurface of u passing through \bar{x} which can be computed through a method described in the Appendices. For other cases, such as Zeng's coupling dependent speed term, the extension is not so natural. In this case one may assign $\mathbf{b}(\bar{x}) = \mathbf{b}(\bar{y})$ where \bar{y} is the closest point to \bar{x} on the surface S. However if one uses such an extension mechanism for \mathbf{b} then the Hamilton-Jacobi equation (C3) can no longer be used to evolve u because \mathbf{b} will depend on u . Gomes works out what this dependency is and uses it to improve the results of Zeng and to compute the medial manifold of the outer cortical surface.

While the fact that driving all of the level sets with equation (C3) cannot conserve the distance map and leads to unbounded values of $|\nabla u|$ has been proven formally in [Barles93], a geometric argument can also be used to illustrate the reasoning. Suppose that \mathbf{b} is extended using the natural method explained above and one considers the case in which we have a point M on the real number line at coordinate $x \in \mathfrak{R}$. Furthermore suppose M is moving along the number line so as to minimize its energy $E(x) = \frac{x^2}{2}$. Let the coordinate of M be 1 at the initial time t_0 . ∇u is one at time t_0 as shown by the dotted line passing through $M(t_0)$ in figure (C2).



[Gomes00] Figure C2: $|\nabla u|$ increases as the point M is evolved.

Now the initial distance map is simply $u_0(x) = x - 1$ and it is evolved according to equation (C3):

$\frac{\partial u}{\partial t} = \mathbf{b}|\nabla u| = x \frac{\partial u}{\partial x}$ if we take $\mathbf{b} = x$. The solution to this PDE is $u(x,t) = e^t x - x_0$. Figure (C2) shows that the zero level set of u (the point M) travels toward $x=0$ as expected to minimize the energy however the distance map is not preserved, instead, $\nabla u = \frac{\partial u}{\partial x} = e^t$ which increases as t increases.

Allowing this condition to remain not only adds complexity to the implementation in the form of a required periodic reinitialization of the distance map but also reduces the theoretical and practical power of the method.

4.1 Theoretical and Practical Reasons to Maintain the Distance Map

From the *theoretical* standpoint, the use of the distance map to represent the parameterized surface is a reasonable choice because these two descriptions are equivalent when S is seen as a subset of \mathfrak{R}^3 . Any surface S defines a unique signed distance map and any implicit function u such that $|\nabla u| = 1$ is the signed distance function to a surface plus a constant –which is taken to be zero on S . So long as the two descriptions are equivalent properties of one can be transposed to the other. For example when the descriptions are equivalent, the distance map u has converged *if and only if* $S(t)$ has converged. Maintaining equivalent representations is also important for several *practical* reasons. Intrinsic properties of $S(t)$ can be determined from local information in u . For example the second fundamental form of the surface $S(t)$ can be computed from the derivatives of the squared distance map as detailed in [Adalsteinsson95]. Also distances between coupled surfaces can be computed directly from their distance maps as described in [Zeng98]. Finally, the singularities in the distance map can be used to compute medial manifold of an evolving surface as described in the analysis section of this method.

4.2 Deriving a New PDE which Maintains the Distance Map

In order to change the PDE given in equation (C3) so that for all t , $u(\bar{x},t)$ is the signed distance map of the solution S of equation (C1), let's first assume that we begin with a signed distance map. That is $u(\bar{x},t=0) = u_0$ is the signed distance function to the initial surface S_0 . What we need is a new right hand side to equation (C3). Let's also call this new right hand side B . Now from equation (C3) we have $\frac{\partial u}{\partial t} = B$. Since in general $\frac{\partial u}{\partial t}$ can change with both position and time, Gomes looks for a function B which is of the form $\mathfrak{R}^3 \times \mathfrak{R} \rightarrow \mathfrak{R}$ and furthermore B must satisfy the two constraints (i) $\bar{x} \rightarrow u(\bar{x},t)$ is a distance function for all t (ii) the zero level set of u evolves along the normal direction to the front with speed $\mathbf{b}(\bar{x})$. These constraints can be expressed as a system of PDE's:

$$B(\bar{x},t) \Big|_{\{\bar{x} | u(\bar{x},t)=0\}} = \mathbf{b}(\bar{x}) \quad (C5)$$

$$\frac{\partial u(\bar{x},t)}{\partial t} = B(\bar{x},t) \quad (C6)$$

$$|\nabla u(\bar{x})| = 1 \quad (C8)$$

The subscript notation in equation (C5) denotes the restriction to the zero level set of u . Differentiating both sides of equation (C6) by the del operator yields:

$$\nabla \frac{\partial u}{\partial t} = \nabla B \quad (\text{C8a})$$

Differentiating both sides of equation (C7) by $\frac{\partial}{\partial t}$ yields:

$$\frac{\partial}{\partial t} |\nabla u| = \frac{\partial}{\partial t} 1 \quad (\text{C8b})$$

$$\frac{\nabla u}{|\nabla u|} \cdot \frac{\partial \nabla u}{\partial t} = 0 \quad (\text{C8c})$$

Factoring out $\frac{1}{|\nabla u|}$ leaves

$$\nabla u \cdot \frac{\partial \nabla u}{\partial t} = 0 \quad (\text{C8d})$$

Now the second term $\frac{\partial \nabla u}{\partial t}$ in equation (C8d) is related to the left hand side $\nabla \frac{\partial u}{\partial t}$ of equation (C8a).

If we expand both of these expressions we see that they are equal:

$$\frac{\partial \nabla u}{\partial t} = \frac{\partial}{\partial t} \left[\frac{\partial u}{\partial x_1} \quad \frac{\partial u}{\partial x_2} \quad \frac{\partial u}{\partial x_3} \right]^T = \left[\frac{\partial^2 u}{\partial t \partial x_1} \quad \frac{\partial^2 u}{\partial t \partial x_2} \quad \frac{\partial^2 u}{\partial t \partial x_3} \right]^T \quad (\text{C8e})$$

$$\nabla \frac{\partial u}{\partial t} = \left[\frac{\partial}{\partial x_1} \quad \frac{\partial}{\partial x_2} \quad \frac{\partial}{\partial x_3} \right]^T \frac{\partial u}{\partial t} = \left[\frac{\partial^2 u}{\partial x_1 \partial t} \quad \frac{\partial^2 u}{\partial x_2 \partial t} \quad \frac{\partial^2 u}{\partial x_3 \partial t} \right]^T \quad (\text{C8f})$$


This is known as the Schwartz equality. Therefore we can substitute ∇B in equation (C8d) leaving the new system of PDE's which determine B:

$$\nabla u(\bar{x}, t) \cdot \nabla B(\bar{x}, t) = 0 \quad (\text{C9})$$

$$B(\bar{x}, t) \Big|_{\{\bar{x} | u(\bar{x}, t) = 0\}} = \mathbf{b}(\bar{x}) \quad (\text{C5})$$

$$\frac{\partial u(\bar{x}, t)}{\partial t} = B(\bar{x}, t) \quad (\text{C6})$$

Equation (C9) states that ∇B and ∇u are perpendicular and that $B(\bar{x}, t)$ does not vary along integral curves of ∇u . These integral curves of ∇u are the characteristics of u , the curves along which ∇u doesn't change. In [Arnold83] it is shown that the characteristics of distance functions are straight lines. Since $|\nabla u|=1$ we have the characteristics lines are the lines which are orthogonal to the zero level set. Equation (C9) also implies that the characteristics of u and B are orthogonal. Thus we have the situation shown in figure C3.



[Gomes00] Figure C3: The characteristic curves of the field ∇u are straight lines orthogonal to the zero level set.

B is constant along these straight lines *and is known* at the point where the rays intersect the zero level set. The value of B is \mathbf{b} at these points as given in equation (C5). Therefore the value of B is propagated through the domain of u along the characteristic curves of the field ∇u .

We are now ready to flesh out B . Let $\bar{x} \in \mathfrak{R}^3$, then the equation of the characteristic line of u passing through \bar{x} is all points on the line $\bar{x} - \mathbf{I} \nabla u(\bar{x}, t)$ where $\mathbf{I} \in \mathfrak{R}$. Where $\nabla u(\bar{x}, t)$ is defined, the closest point to \bar{x} on the zero level set occurs along this line. Specifically, the distance from \bar{x} to the zero level set is $u(\bar{x})$ and $|\nabla u(\bar{x}, t)| = 1$ therefore the point $\bar{y} = \bar{x} - u(\bar{x}, t) \nabla u(\bar{x}, t)$ is the closest point to \bar{x} on the zero level set of u , i.e. \bar{y} is the closest point to \bar{x} such that $u(\bar{y}) = 0$. Thus we have $B(\bar{y}) = B(\bar{x}) = \mathbf{b}(\bar{x} - u \nabla u)$. Having developed an expression for B *the reformulated PDE* can now be stated as:

$$\frac{\partial u}{\partial t} = \mathbf{b}(\bar{x} - u \nabla u) \quad (\text{C10})$$


whose initial condition is $u(\bar{x}, t=0) = u_0(\bar{x})$. Equation (C10) should be used to evolve the distance map instead of equation (C3). One way of interpreting equation (C10) is that the zero level set of u is driven by $\frac{\partial u}{\partial t} = B$ which implies that the evolution of this surface defines, through propagation, the evolution of **all** other level set surfaces in the domain of u .

Solutions to equation (C10) are of the form

$$u(S + \mathbf{I} \hat{n}) = \mathbf{I} \quad \forall t, \mathbf{I} \quad (\text{C10a})$$

which can be proven by substitution and differentiation with respect to t and λ . Examining (C10a) one can see that the surface parallel to S and λ units away from S should be the λ level set of u . This makes intuitive sense since we have set out to develop a new PDE to maintain the distance map. The constraint given in equation (C10a) can be compared with the less rigorous constraint $u(S, t) = 0 \quad \forall t$ given by Osher and Sethian.

As shown in figure C4 there are points $\bar{x} \in \mathfrak{R}^3$ at which $\nabla u(\bar{x}, t)$ is not defined. This can occur when the closest point \bar{y} to \bar{x} on the zero level set is not unique.



[Gomes00] Figure C4: Portions of the skeleton of the zero level set is shown by the dotted curves.

Such points form the skeleton or medial manifold of S . A medial manifold representation (or medial axis in the 2D case) is used in many computer vision methods as a compact shape description which can represent a wide variety of shapes. It is used as the shape representation in the medical image segmentation methods used by Pizer et al. in [Pizer99]. With an appropriate implementation of the new level set evolution equation, the skeleton of the evolving surface can be computed concomitant with evolution.

4.3 An Implementation Strategy for the New PDE

4.3.1 The Discrete Update Equation

Unless the new PDE in equation (C10) can be implemented it is of little practical value. Gomes, however, describes an implementation strategy. It is assumed that u is initialized to the signed distance function to S_0 . Then at each step in the evolution, $B(\bar{x}) = \mathbf{b}(\bar{x} - u\nabla u)$ is computed on a narrow band around the zero level set. Finally u is updated with the Euler method:

$$u(\bar{x}, t+1) = u(\bar{x}, t) + B(\bar{x}, t) \quad (\text{C10b})$$

Two steps are required to compute B and they correspond to equation (C5) and equation (C9). The main remaining challenge is generating real values results when the distance map is known only at discrete samples.

4.3.2 Handling Discrete Sampling

The strategy for handling the sampling effects can be better described if we first set up several definitions and conventions.

- Let lowercase letter such as \bar{x} denote real valued points in \mathfrak{R}^3
- Let upper case letters such as X denote integer values voxels in \mathbb{N}^3
- Let $V(\bar{x})$ be the set of 9 voxels whose coordinates are formed from the floor and ceil of the components of \bar{x} .
- Let the subscript “l” as in $F_l(\bar{x})$ denote the trilinear approximation of the given function F at \bar{x} . For linear interpolation of the function $p(\cdot)$ we have: $p(x') = \sum_{i=1..2} (1-T_i)p_i$ where $T_i = |x'-x_i|$ and where the weight given to p_i is equivalent to the distance to the opposite point. Here every pair of

consecutive data points has a separation of one unit between them thus the opposite point is $1 - T_i$. For trilinear interpolation this sum is extended over all points:

$$p(x', y', z') = \sum_{i=1..2} \sum_{j=1..2} \sum_{k=1..2} (1 - T_i)(1 - R_j)(1 - S_k) p_{ijk} \quad (\text{C10c})$$

where $T_i = |x' - x_i|$, $R_j = |y' - y_j|$ and $S_k = |z' - z_k|$.

- Let $A(X)$ be the 26 connected neighborhood of the voxel X
- Let Z be the set of voxels adjacent to the zero level set of u . Z must be computed carefully since the zero level set of u is composed of real points. In order to determine which voxels are adjacent to the zero level set:
- Let C_u be a function defined on the voxels of u such that:

$$C_u(X) = \begin{cases} 0 & \text{if } u(X) > 0 \\ 1 & \text{if } u(X) \leq 0 \end{cases} \quad (\text{C10d})$$

then a voxel X is adjacent to the zero level set if there exists a voxel Y in the neighborhood of $A(X)$ such that $C_u(X) \neq C_u(Y)$.

4.3.3 Computing \mathbf{b} at Voxels Adjacent to the Zero Level Set

Since we need to propagate the value of \mathbf{b} on the level set to the narrow band around it, the first step is to compute the value of \mathbf{b} on Z . There are 2 cases which typically occur in practice. In the first \mathbf{b} is defined on \mathcal{R}^3 so that no matter where the level set is we can compute the speed for all points along it. If this is the case then B^Z (read “the value of B on Z ”) can be expressed as:

$$B^Z(X) = \mathbf{b}(X) \quad \forall X \in Z \quad (\text{C10e})$$

In the second case \mathbf{b} is defined only on the nodes of a mesh describing the zero level set of u . In this case B^Z can be defined as:

$$B^Z(X) = \mathbf{b}(v_i) \quad \forall X \in Z \quad (\text{C10f})$$


where v_i is the closest node of the mesh to the voxel $X \in Z$. In either case $B^Z(X)$ is just an intermediate value used to compute $B(X)$ on the whole narrow band.

4.3.4 Computing \mathbf{b} on the Whole Narrow Band

In order to propagate values from the discrete level set, Z , to the whole narrow band we need a method to compute, for each point X in the narrow band, the closest point \bar{y} in Z which lies on the same characteristic of u as X . Once we have computed \bar{y} then the propagation follows as:

$$B(X, t) = B^Z(\bar{y}, t)$$

where $u_t(\bar{y}) = 0$. The first inclination is to compute \bar{y} from $X - u(X)\nabla u(X)$ however this suffers from numerical instability since small errors in $\nabla u(X)$ will introduce large errors in \bar{y} . To minimize these errors Gomes proposes that the characteristic passing through X be followed in unit steps as shown in figure C5



[Gomes00] Figure C5: Following unit steps from B to the zero level set is generally more stable

The algorithm to march along the characteristic is as follows:

1. Let $\bar{y}_0 = X$
2. Repeat $\bar{y}_{n+1} = \bar{y}_n - \begin{cases} \max \{u_l(\bar{y}_n), \text{sign}(u_l(\bar{y}_n))\} \nabla u_l(\bar{y}_n) & \text{if } u_l(\bar{y}_n) < 0 \\ \min \{u_l(\bar{y}_n), \text{sign}(u_l(\bar{y}_n))\} \nabla u_l(\bar{y}_n) & \text{if } u_l(\bar{y}_n) > 0 \end{cases}$
3. Until $u_l(\bar{y}_n) = 0$

The marching procedure is used to compute values of B for all voxels in the narrow band, even those in Z. Notice that in the first step, when $n=0$, the method computes $\nabla u(X)$. If X is situated on the skeleton of S then $\nabla u(X)$ is undefined. To avoid such difficulties, the presence of singularity has to be detected. Gomes' method to detect these singularities requires switch functions which turn on or off when a singularity is detected.

To estimate ∇u first define 6 discrete derivative operators:

$$\begin{aligned} D_x^+ u &= u(i+1, j, k) - u(i, j, k) \\ D_x^- u &= u(i, j, k) - u(i-1, j, k) \end{aligned} \tag{C10h}$$

and likewise for D_y and D_z . Then eight estimators D^i $1 \leq i \leq 8$ of ∇u :

$$\begin{aligned} D^1 u &= (D_x^+ u, D_y^+ u, D_z^+ u) \\ D^2 u &= (D_x^+ u, D_y^+ u, D_z^- u) \\ &\vdots \\ D^8 u &= (D_x^- u, D_y^- u, D_z^- u) \end{aligned} \tag{C10i}$$

Finally the approximation of ∇u is given as:

$$\nabla u(X) = \arg \max_i \left(\left| D^i u(X) \right| \right) \tag{C10j}$$

The direction of the characteristic line towards the closest point to X on the zero level set of u is the direction of maximum slope at X. The presence of the skeleton at X can be detected by comparing the eight estimators as described in the skeleton computation analysis below.

4.4 Results and Critical Analysis of the Method


4.4.1 Revisiting the Coupled Surface Propagation

Gomes tested this implementation strategy of the new PDE on the evolution of coupled surfaces method proposed by Zeng. Recall that in this method the inner surface S_{in} locates the boundary between white and gray matter and the outer surface S_{out} captures the interface between gray matter and CSF. Gomes defines new expressions for the speeds of these two surfaces:

$$\mathbf{b}_{in} = f(I - I_{in}) + C(u_{out} + \mathbf{e}) \quad (C11)$$

$$\mathbf{b}_{out} = f(I - I_{out}) + C(u_{in} - \mathbf{e}) \quad (C12)$$

where I is the MR voxel intensity, I_{in} and I_{out} are the fixed positive thresholds for white and gray matter, \mathbf{e} is the expected cortical thickness and f and C are functions which have the shapes shown in figure C6.




[Gomes00] Figure C6: Shapes of the f and C functions. δ_I and δ_d are fixed tolerances.

We can investigate how Gomes' equations govern the motion of evolution by examining each of the terms in equation (C11) and (C12). We can see that the first term in each equation enforces a degree of data fidelity while the second implements a coupling constraint. The surfaces are evolved inward due to the definition of an *inward* pointing normal, \hat{n} , given in equation (C1). This contrasts with Zeng's outward propagation. For the purposes of discussion, let's assume that we want to understand the behavior of the evolution of S_{in} . The objective of S_{in} is to locate the gray-white boundary. Assuming further that white matter has higher intensity values than gray matter we would want S_{in} to continue to collapse inward from an initial state in the CSF outside the convex hull of the cortical layer, whenever the intensity on S_{in} is less than I_{in} . This would indicate that S_{in} lies over CSF or gray matter which is typically darker than white matter on T2 weighted MR and that S_{in} has not yet reached the gray-white boundary. This is exactly what f does when $I - I_{in} < 0$, f is positive and when $I - I_{in} > 0$ then f is negative indicating that S_{in} has gone too far and should move back. We can also analyze the second term $C(u_{out} + \mathbf{e})$. When u_{out} is evaluated at a point p on S_{in} , it yields the distance to S_{out} from that point on S_{in} . We have $u_{out}(p) < 0$ since points in the interior of a contour are defined to be a negative distance from the contour. When S_{in} is too far inside S_{out} then $u_{out} + \mathbf{e} < 0$ and we want S_{in} to slow down and move back (i.e. outward). When we look at C we see exactly this behavior since $C(w) < 0$ when $w < 0$. A similar analysis can be done for S_{out} on equation (C12).

If the coupled surface propagation is implemented with equation (C3) then the magnitudes of u_{in} and u_{out} will increase with time. This will result in an overestimation of the amount of separation between the surfaces. Since in equation (C11) the coupling term is negative while it is positive in equation (C12), the surfaces will move closer and closer to one another until the distance function is corrected. Moreover while the coupling term is non zero the full effect of the data term is not seen in the segmentation.

Figure C7 shows illustrates Gomes' results on original MR data from 3 different



[Gomes00] Figure C7: Segmentation results using different algorithms

areas of the brain, shown in column 1 with concavities to compare circled and labeled A,B,C. Column 2 shows simple thresholding at T1 and T2, column 3 the cross sections of final S_{in} and S_{out} surfaces when coupling terms are not used. Column 4 shows the results when the coupling terms are used and the evolutions are implemented with equation (C3). While the distance constraint is maintained, the data coupling term fails to have full effect. Finally in the last column the evolution is driven by the new PDE. The concavities A,B,C are better captured because of the preservation of the distance map enables the coupling term to stop having effect when the distance constraint is satisfied and the data term to drive the surfaces in accord with image intensities.

Gomes' method imposes weaker constraints for the thickness of the cortical layer compared to Zeng's. In particular in Gomes' formulation there is no maximum distance only a minimum distance of separation. Accurate localization of these surfaces is important since atrophy in cortical thickness is often touted as a indicator of neuropathology. The results show what appears to be* a more accurate localization of the inner and outer cortical surfaces. *In [Gomes00] Gomes does not perform a quantitative comparison against expert segmentation and relies solely on the qualitative analysis shown in figure C7.

Initialization of a surface which involves inward is more difficult than one which evolves outward. More initial points must be specified to specify the initial surface which must lie completely outside the cortical layer. Since the cortical layer is surrounded by several other layers if finite thickness, such as the three meninges surrounding the brain, inward evolution is prone to error – any one of these layers could be misidentified as a cortical layer. Thus a more robust approach for initialization seems to be outward evolution.

4.4.2 Extracting the Medial Manifold of an Evolving Surface

As described above the set of points at which ∇u is undefined form the skeleton of the zero level set of u . Since the new PDE preserves the distance map throughout its evolution, when u is properly initialized to the distance map of S_0 , the skeleton can be computed at any time step of the surface's evolution. Skeletons are practical shape representations for many computer vision applications such as object segmentation, object recognition and registration because of the wide range of shapes that can be represented and because of the compactness afforded by representing

opposing surfaces as one surface. Fortunately, skeletons can be computed from the information formed in the evolution using the new PDE.

Points on the skeleton of the level set of u are points where the variations in the direction of ∇u are large. The eight estimators $D^i u$ of ∇u will differ significantly at such points therefore Gomes proposes the use of a global threshold on the following function for the identification of skeleton points:

$$\sum_{i=1}^8 \left(\frac{D^i u}{|D^i u|} \cdot \frac{D\bar{u}}{|D\bar{u}|} \right)^2 \quad (\text{C13})$$

where $D\bar{u} = \frac{1}{8} \sum_{i=1}^8 D^i u$


Gomes test this skeletal computation method by evolving the final S_{out} surface of the cortex found in the previously described results. During this second round evolution the cortical layer is inflated: regions of high curvature are smoothed by inflating them outward. Thus type of surface simplification is called mean curvature flow and is governed by the PDE:

$$\frac{\partial S}{\partial t} = H\hat{n} \quad (\text{C14})$$

where the speed term $\mathbf{b} = H$ where H is the mean curvature of the surface S . Thus we have for evolution:

$$\frac{\partial u}{\partial t} = H(\bar{x} - u(\bar{x}, t))\nabla u(\bar{x}, t) \quad (\text{C15})$$

An example of the S_{out} evolution is shown in the left column of figure C8. In examining these qualitative results, we can see that the sulci are inflated outward and the surface as a whole get smoother and smoother through the evolution. This puts the cortex in a form in which the gray matter in the folds are more visible and facilitates intersubject comparison. The corresponding skeletons of the evolving surface shown in the right column. All of the skeletons in figure C8 appear to be reasonable approximations to a medial manifold. A qualitative test on a structure with a known medial manifold would help confirm these findings. As a minor note the last 3 pairs do not appear to correspond well—for example notice that in the 5th row, the bottom of the shape on the left has a one one downward protrusion while the left shape has two. This is probably just an error in forming the collage for the figure, and in any case the theory and mechanism for detecting the skeletal points seems sound.



[Gomes00] Figure C8: Evolution of the cortical surface S_{out} and simultaneous skeleton computation.
Note that the each surface is rescaled to fit the entire image.

5 A Method to Guide Geodesic Active Contour Evolution using a Statistical Shape Model

The research of Zeng [Zeng98] demonstrate favorable results in the application of the level set propagation framework using coupled surfaces to segment structures with constant thickness including the cerebral cortex and the myocardium of the left ventricle. Lorigo [Lorigo99] extended the single surface level set propagation framework to segment structures of co-dimension 2 including the tubular vascular branching structures in MRA. In “Statistical Shape Influence in Geodesic Active Contours” Leventon broadens the class of shapes which the level set method can segment reliably from simply shapes of constant thickness or tubular structures. The method can be applied to segment any shape (2D or 3D) provided the variance in the shape distribution is not too large and provided the training set is representative of the shape distribution. This method is exciting because many anatomical structures have a characteristic mean shape and pose in the human body. For example the corpus callosum, the arched bridge of nervous tissue that connects the two cerebral hemispheres, allowing communication between the right and left sides of the brain, typically appears in a classic “C” shape with the ends of the “C” inferior and the curved part of the C superior. This structure also typically appears at the center of the brain and oriented with the long axis extending from anterior to posterior. To segment anatomical structures such as the corpus callosum, an initial curve is embedded as the zero level set of a higher dimensional surface. Next a maximum a posterior estimate of the shape and pose of the object is formed using the current zero level set and the shape model. Then the surface is evolved globally towards the MAP estimate and locally using image gradient and surface curvature. Lastly the process of estimation and surface evolution is repeated until the surface converges.

There are four salient aspects of Leventon’s method:

- Parameterization of shape and pose and formation of shape and pose probability distribution models from the parameterized training instances
- A method for estimating which shape and pose from the distribution to guide the evolution
- Selection of a level set evolution equation to extend
- A strategy for steering traditional level set evolution in the direction of the an estimated global shape

In the next section these aspects will be described individually and then the integration of these methods into an overall segmentation strategy will be described. The results of this method on real and synthetic data will be analyzed and critical analysis of the method will be given including suggestions for extending the method.

5.1 The Shape Model

The Leventon method assumes the anatomical structure to be segmented has a shape which is drawn from a shape distribution function. This distribution is computed from a given set of hand labeled training instances. To construct the shape prior model a shape representation is selected and then all of the training instances are coded with that shape representation. A probability function over the parameters of the representation is then computed from the coded training instances.

5.1.1 Shape Representation

The shape representation is computed from the training instances using a signed distance function. Let the shape be either a curve or surface, C , in \mathcal{R}^d . The signed distance function is computed after quantizing \mathcal{R}^d into N samples in each of the d dimensions. At these N^d samples the Euclidean distance to the surface is computed. The sign is positive if the sample is outside C and

negative if inside. Now each distance map, u , computed in this way is actually a higher dimensional surface, one dimension higher than C . The distance map is the shape representation and the shape that it represents can be extracted from the map by taking the locus of the points for which $u=0$. This shape is the zero level set of the distance map.

With the shape representation chosen, the shape distribution is characterized through the analysis of a training set T of shapes. This training set consists of a set of n surfaces:

$T = \{u_1, u_2, \dots, u_n\}$ which are the distance maps which are uniquely determined by each hand drawn shape. A cloud of n points is formed from stacking the columns from u_i into one long column vector. Since this space is so large, a reduction in dimensionality is sought for in which only the principal modes of shape variation are retained from the point cloud. Justification of the dimensionality reduction is two fold. First there is redundancy in the distance map because each map is determined uniquely from a lower dimensional zero level set. Second, the collection of curves in the training set should be strongly correlated because they are shapes of the same category of object. The cloud is assumed to be Gaussian in nature and it is assumed to have high variances in only a small fraction of the dimensions of the \mathcal{R}^{N^d} space.

A multivariate Gaussian distribution is parameterized by a mean vector \mathbf{m} and a covariance matrix V . The mean shape is found by taking the mean of the training distance maps:

$$\mathbf{m} = \frac{1}{n} \sum_{i=1}^n u_i \quad (\text{A0})$$

The variance in shape is computed using Principal Component Analysis. The mean shape is subtracted from each training distance map u_i to create a mean offset map \hat{u}_i . These n $N^d \times 1$ column vectors are used to form the columns of a $N^d \times n$ matrix M where the i^{th} column of M is \hat{u}_i . The covariance matrix can be expressed succinctly in matrix form as:

$$V = \frac{1}{n} M M^T \quad (\text{A1a})$$

where V is an $N^d \times N^d$ symmetric positive semi-definite matrix. Through Singular Value decomposition V can be factored into the product:


$$V = U \Sigma U^T \quad (\text{A1b})$$

where U is a $N^d \times N^d$ matrix whose columns are orthogonal and which represent the modes of shape variation and where Σ is an $N^d \times N^d$ diagonal matrix of singular values. The i^{th} singular value σ_{ii} corresponds to the i^{th} mode of shape variation and to the i^{th} column of U . The larger the singular value the greater the spread of the shape cloud in the direction of the i^{th} column. Using only the k modes of shape variation corresponding to the k^{th} largest singular values leads to a dimensionality reduction when $k \ll N^d$. Therefore a shape can be approximated by k principal components in a k -dimensional vector of coefficients α . α can be computed from:

$$\mathbf{a} = U_k^T (u - \mathbf{m}) \quad (\text{A2})$$

where U_k is the matrix consisting of the 1^{st} k columns of U and is used above to project the surface into the eigen-space. A full $N^d \times N^d$ representation \tilde{u} can be constructed from α by:

$$\tilde{u} = U_k \mathbf{a} + u \quad (\text{A3})$$



[Montillo-00] Figure A0: Constructing a distance map from a k-vector α . Just two of the N^d dimensions are shown for diagrammatic purposes. These are represented by the horizontal and vertical axes.


In general \tilde{u} will not be a true distance function because convex linear combinations of distance maps do not produce distance maps, but the surfaces \tilde{u} will be smooth and have zero level sets consistent with combinations of the original curves.

With the shape represented in the k-coordinate vector α in eigen-space and under the assumption of a Gaussian distribution of shape, the probability of a given shape α can be computed as:

$$P(\mathbf{a}) = \frac{1}{\sqrt{(2\pi)^k |\Sigma_k|}} \exp\left(-\frac{1}{2} \mathbf{a}^T \Sigma_k^{-1} \mathbf{a}\right) \quad (\text{A4})$$

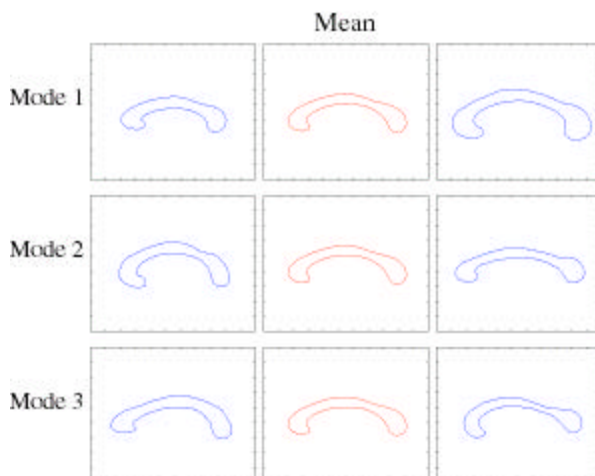
where Σ_k contains the first k rows (and columns) of Σ .

Experiments on real medical data illustrate the process. In a relatively simple 2D case the training set consisted of 49 training curves hand drawn to the boundary of the corpus callosum. Some of these curves are shown in figure A2. The curves are initially roughly aligned using centroids and second moments. Then the signed distance map is computed. It is shown in the background in the images: roughly speaking, whiter shades are positive distances from the curve while dark shades are negative distances to the curve.




[Leventon00] Figure A2: several training curves for the corpus callosum

The mean level set of these distance maps is shown in the center column of figure A3. The three rows in the figure show the three primary modes of variance about the mean.



[Leventon00] Figure A3: Primary modes of variance of shape about the mean

The method can be readily extended to 3D, albeit with longer computation time. In the top row of figure A4 seven thoracic vertebrae T3-T9 of one subject used as training data are shown with a triangular surface model and after a rigid alignment. Signed distance maps are computed for the vertebrae with positive distances outside the vertebrae and negative distances inside the vertebrae. In the second row the two primary modes of variation are shown.



[Leventon00] Figure A4: Training vertebrae and two primary modes of variation

In both the 2D corpus callosum test and the 3D vertebrae test the primary modes of variation explain a fair amount of the variation in the shape distribution. It is a somewhat subjective matter however to tell if they relate to an underlying physiological cause or if they represent change along a medically significant dimension. Since *any* matrix can be decomposed using SVD, even errors in alignment could be explained by the modes. In the case of the corpus callosum, the first mode seems to capture the overall size and the second the overall bending of the corpus. It is less clear what the modes represent in the 3D case, though they appear reasonable given the training data in the top row of the figure.

5.1.2 Addressing the Correspondence Problem

In order to accurately characterize a shape distribution from a set of static training data exemplars, corresponding parts of the objects must be compared. In the rotationally variant signed distance representation misalignment could corrupt the primary modes. For example in the case of the vertebrae if a process of one is overlapping a concavity of another then the discrepancy will contribute to a false mode of shape distribution.

Certainly one solution to the correspondence problem is to identify corresponding points on the shapes in the training set. This would ensure that the corresponding parts of the objects are compared. Manually labeling the correspondences is tedious especially in voluminous 3D imagery and not applicable to the creation of a fully automated algorithm. Leventon's method *does not* generate corresponding feature points automatically nor does it use any hand labeled correspondences. Instead the method relies on the inherent tolerance of the signed distance map to slight misalignments of object features. Slightly misaligned pixels in the distance maps are generally still highly correlated. Leventon uses only rough rigid alignment of the training shapes. Thus rigid pose parameters are considered extrinsic to the shape representation but scale and affine warp are left as intrinsic and thereby contribute to the modes of variation.

5.2 Geodesic Active contours

Having characterized a shape distribution from a set of training instances, Leventon then borrows an existing segmentation strategy and adapts it so that the segmentation result tends towards probable shapes in the shape model. The segmentation strategy which is selected is geodesic active contours. This method is derived from the standard snake methodology in which Kass et al. define the optimum segmentation as the curve C which minimizes the internal and external energy of the curve:

$$E(C) = \mathbf{b} \int |C'(q)|^2 dq - \int |\nabla I(C(q))| dq \quad (\text{A5})$$

where the segmentation result C^* is the curve which minimizes $E(C)$:

$$C^* = \min_{C(q)} E(C) \quad (\text{A6a})$$

Caselles et al. [Caselles] derive the equivalence of the geodesic active contours method to the energy based snakes framework. The minimization problem is reduced to the form:

$$\min_{C(q)} \int g(|\nabla I(C(q))|) |C'(q)| dq \quad (\text{A6b})$$

where g is a function of the form

$$\frac{1}{1 + |\nabla I|^2} \cdot \quad (\text{A6c})$$

As stated neither of these formulations describe a solution strategy: a naïve brute force approach, trying all possible curves would fail because the search space is simply too large. However assumptions can be made to make the problem manageable. In this paper these assumptions are:

1. At least one interior point of the object to be segmented is known
2. The object to be segmented has a closed boundary

For many anatomical structures, these assumptions are reasonable. The assumptions imply that we could initialize a closed contour inside the object to be segmented and that in the absence of other information this contour should be expanded outwards to reach the object boundary.

Caselles uses the Euler-Lagrange method to convert this minimization problem into an equation for the evolution of the curve C:

$$\frac{\partial C(t)}{\partial t} = g \mathbf{k} N - (\nabla g \cdot N) N \quad (\text{A7})$$

where \mathbf{k} is curvature along the curve and N is the unit normal. The update function for the signed distance map u becomes:

$$\frac{\partial u}{\partial t} = g(c + \mathbf{k}) |\nabla u| + \nabla u \cdot \nabla g \quad (\text{A8})$$

where c is the image-dependent balloon force which is added to force the contour outward. In this level set framework the surface u evolves at every point perpendicular to the level set curve and at a rate dependent on local curvature and image gradient.

5.3 Evolving the Surface with Shape Influence

The surface at time $t = 0$ is assumed to be a closed contour which lies inside the object to be segmented. Given this surface $u(t)$ the method seeks an evolution step which deforms the surface using local gradient and global shape information so that it moves closer to the correct final segmentation. The level set update equation A8 provides a way to evolve the surface so that it satisfies the original curve minimization problem given by equation A6. Using the Euler method equation A8 is transformed into an expression to compute the evolution numerically at discrete time increments:

$$u(t+1) = u(t) + \mathbf{I}_1 (g(c + \mathbf{k}) |\nabla u(t)| + \nabla u(t) \cdot \nabla g) \quad (\text{A16})$$


where λ_1 is a parameter defining the constant scale factor to apply to the direction indicated by the curvature and image gradient information.

To pull the surface in the direction of a probable final shape, a maximum a posteriori final shape u^* is estimated given the location of the curve at time t and a term is added to the level set update equation. $u^*(t)$ is defined in the next section. If we only wanted to evolve the surface in the direction of the MAP estimate the following simple update equation would be used:

$$u(t+1) = u(t) + \mathbf{I}_2 (u^*(t) - u(t)) \quad (\text{A17})$$

where $\mathbf{I}_2 \in [0,1]$ is a constant coefficient which determines how much trust to place in $u^*(t)$.

As shown in figure A8 the gradient-curvature based influence can be added using standard vector addition to the shape model influence:



[Leventon00] Figure A8: Red and blue makes purple: *Vector addition* combines the gradient-curvature force and the shape prior force vector

Combining these influences yields the expression for updating the surface:

$$u(t+1) = u(t) + I_1(g(c + \mathbf{k})|\nabla u(t)| + \nabla u(t) \cdot \nabla g) + I_2(u^*(t) - u(t)) \quad (\text{A18})$$


The parameters λ_1 and λ_2 weight the influence of the curvature-gradient model and shape model on the evolution. λ_1 is large relative to λ_2 if more faith is placed in the image quality and λ_2 is large relative to λ_1 if there is much trust placed in the shape model. The λ_1 and λ_2 effect a trade off between the shape and image influences.

5.4 Estimating the Final Shape and Pose

In order to add a global shape term to the evolution equation the final distance map $u^*(t)$ must be estimated at each time step. This map is estimated using a maximum a posteriori approach:

$$u_{MAP}^* = \arg \max_{u^*} P(u^* | u, \nabla I) \quad (\text{A10})$$


where u is the evolving surface at time t whose zero level set is the curve that is segmenting the image and ∇I is the gradient of the image to be segmented. For example in figure A5 the mean curve is shown in A. It is the initial estimate of the final segmented shape α . Part B of the figure shows the image containing the shape to be segmented. In part C the same solid elliptical curve is shown in different location within the object and the dashed line depict a later step in the evolution which is dependent upon the solid curve's position and shape.



[Leventon00] Figure A5: A: expected curve. B: image to be segmented C: 3 example initial conditions and subsequent evolutions

It is important to realize that coordinate frame for u and the coordinate frame of the shape model are generally not aligned. This occurs because the pose of structures in the body vary from subject to subject so training instances will not be aligned to test structures. Thus in order for the shape model to adequately constrain the evolution. The pose of the evolving curve with respect to the shape model must be known. In figure A6 three stages in the evolution of the blue curve are shown in

the top row with evolution proceeding from left to right. In the second row the curve is matched to the expected curve to obtain a PDF over pose. Notice that there is ambiguity in the first stage which is resolved in the second and further refined in the third.



[Leventon00] Figure A6: The evolving curve is matched to the expected curve to obtain a PDF

The map u^* is completely determined by the shape parameters α and pose parameters p of the final shape. Therefore we also have

$$\langle \mathbf{a}_{MAP}, p_{MAP} \rangle = \arg \max_{\mathbf{a}, p} P(\mathbf{a}, p | u, \nabla I) \quad (\text{A9})$$

To compute the maximum a posteriori final shape, the terms in equation 9 are expanded using Bayes' rule:

$$P(\mathbf{a}, p | u, \nabla I) = \frac{P(u, \nabla I | \mathbf{a}, p) P(\mathbf{a}, p)}{P(u, \nabla I)} \quad (\text{A11a})$$

which says that the a posteriori probability of the shape and pose parameters given the observed image gradient and current distance map is equal to the product of the conditional probability of the distance map and image gradient given the shape and pose parameters times the prior probability of the shape and pose parameters normalized by the probability of the distance map and image gradient. Assuming that pose and shape are independent equation A11a can be written as:

$$P(\mathbf{a}, p | u, \nabla I) = \frac{P(u | \mathbf{a}, p) P(\nabla I | \mathbf{a}, p, u) P(\mathbf{a}) P(p)}{P(u, \nabla I)} \quad (\text{A11b})$$

Leventon does not attempt to capture the relationships between pose and shape. For MAP estimation of α and p the normalized term in the denominator can be discarded because it does not vary with respect to shape or pose.

The main work involved in developing this method is to define the remaining terms with mathematical expressions in terms of the observations or terms derived from the observations.

5.4.1 Inside Term: $P(u | \mathbf{a}, p)$

The term in equation A11b computes the probability of an evolving curve (represented as the distance map u) given shape α and pose p of the final curve. The term does not include image information. In this implementation of level sets, the initial curve must be initialized to a contour which lies completely inside the object to be segmented and the only external force applied during evolution is an inflationary balloon force. Thus the curve tends to remain inside the object throughout the evolution: expanding until it reaches the boundary of the object. On these reasons Leventon gives curves evolving which lies completely inside the final curve estimate a higher probability than an evolving curve which lies partially or fully outside the final curve:


$$P(u | \mathbf{a}, p) = \exp(-V_{outside}) \quad (A12)$$

where $V_{outside}$ is the volume of the curve u that lies outside the curve u^* . This term implies that any curve u lying inside u^* is equally likely and that the curve can evolve along any path towards the boundary of the object.

5.4.2 Gradient Term: $P(\nabla I | \mathbf{a}, p, u)$

The second term in equation A11b computes the probability of the image gradient given the current and final curve. Consider figure A7, which shows a portion of a gray object on a white background and the normal \hat{n} and tangent \hat{t} vectors to the object boundary at the point q . Assuming that the object boundary is a smoothed step edge, then $|\nabla I|$ approximates a Gaussian along the normal at q . When the distance map u^* correctly outlines the boundary of the object, u^* is linear along the normal direction to the curve at any boundary pt q . Also along this direction

$$\begin{aligned} u^*(q) &= 0 \text{ and} \\ u^*(\text{point inside the object}) &< 0 \text{ and} \\ u^*(\text{point outside the object}) &> 0. \end{aligned}$$



[Leventon00] Figure A7: $|\nabla I|$ vs $u^*(q)$ along the normal direction to the boundary

Since u^* is linear along the normal direction, it is likely that there is a Gaussian relationship between $|\nabla I|$ and u^* . In figure A7 the scatter plot of $|\nabla I|$ vs $u^*(q)$ along the normal direction to the boundary provides empirical data which supports this belief. Letting $h(u^*)$ be the best fit Gaussian to the samples $(u^*, |\nabla I|)$, the gradient term is modeled as a Laplacian of the goodness of fit of the Gaussian:

$$P(\nabla I | u^*, u) = \exp\left(-\|h(u^*) - |\nabla I|\|^2\right) \quad (A13a)$$

5.4.3 Shape and Pose Terms: $P(\mathbf{a})$ and $P(p)$

The last two terms in equation A11b are based on prior models of shape and pose. The shape prior is the multivariate Gaussian distribution over the shape parameter vector α with shape variance Σ_k :

$$P(\mathbf{a}) = \frac{1}{\sqrt{(2\pi)^k |\Sigma_k|}} \exp\left(-\frac{1}{2} \mathbf{a}^T \Sigma_k^{-1} \mathbf{a}\right) \quad (\text{A14})$$

Leventon's method does not utilize information in the training instances about the probability of an object appearing with a given pose in the image. Instead the method assumes a uniform distribution over pose parameters:

$$P(p) = U(-\infty, \infty) \quad (\text{A15})$$

Thus the prior pose probability term does not affect the MAP estimate of α and p because it does not vary with those parameters.

New shape and pose parameters are computed from a given current curve u and image gradient ∇I and previous α and p using standard gradient ascent on the log of the probability function given in equation A11b.

Expressing equation A13a in terms of α and p :


$$P(\nabla I | \mathbf{a}, p, u) = \exp\left(-|h(\mathbf{a}, p) - \nabla I|^2\right) \quad (\text{A15b})$$

we can see that this probability is a function of the shape and pose estimate and hence this term will contribute shape and pose terms in the gradient ascent method used to find the MAP estimate. For efficiency the terms in the MAP estimate of α and p are computed only in a narrow band around the zero level set of the evolving surface. Moreover the MAP pose and estimate are re-estimated at each evolution step.

5.5 Results

Each ascent of the gradient ascent method may produce a local maximum, however empirical results show that continuous re-estimation of these parameters as the surface evolves, Leventon reports, usually results in convergence on the desired boundary shape and pose parameters.

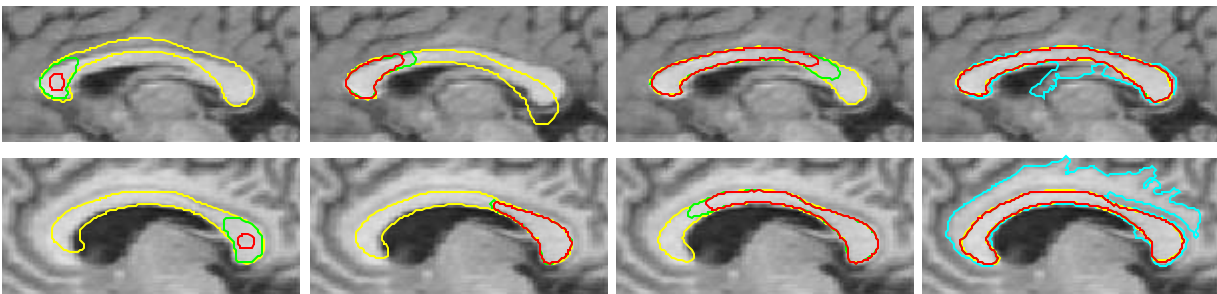
The algorithm was tested on both synthetic objects and real objects in 2D and 3D. During the synthetic tests rombi of various sizes and aspect ratios were generated for the training the shape model . The test images contained two random rombi and additive Gaussian speckle noise as well as a low frequency field bias. Two of these test images are shown in figure (A9) along with four stages in the evolution of the initial curve.



[Leventon00] Figure A9: Controlled tests with synthetic rombi.


The initial curve is shown in red on the left most image in each of the two sets. The green curve is the next step in the evolution of the curve. The yellow curve is the MAP estimate of the shape and pose. The final segmentations are shown in the right most image in each of the two sets. Notice that the images shown the size and position of the MAP estimate can change during the course of curve evolution.

The algorithm was also tested on real objects including the corpus callosum, femur and vertebrae. 51 slices of a corpus callosum were acquired, two were left out for testing and remaining 49 were used for the training set. The segmentations on the two left out are shown in figure (A11) with steps from the first callosum shown in the top row and those steps from the second callosum in the second row.



[Leventon00] Figure A11: Corpus callosum segmentations. Both the pose and shape of the MAP estimate change during curve evolution.

Two different initial curves were used. They are shown in the left most column as the small red contour –on the left in the first row and on the right in the second row. Green curves are the next step in the curve evolution and the yellow curve is the MAP estimate of shape and pose. In these segmentaitons it can see that the initial MAP estimate converges on the boundary of the object. The cyan contour in the right most column shows the segmentation results when only the gradient-curvature based term is used during the evolution. This curve corresponds to a contour along which the image gradient is high and was found through expanding the same intial contour, however it does not correspond to the desired segmentation of the of the corpus callosum and provides support for the need for shape model constraints – particularly in the absence of boundaries with gaps or weak edge regions. Similar tests and results are shown on the femur slice segmentation shown in figure (A10)




[Leventon00] Figure A10: Femur segmentation. Final contour locks onto the femur boundary.

Segmentation of the femure and corpora converged in under a minute on a 550 MHz Pentium.

3D segmentation tests were conducted on the vertebra T7. the training set consisted of vertebrae T3-T9 with the T7 left out. The initial surface was a small sphere completely contained inside the vertebra. As shown in figure (A12) the red contour is a slice through the zero level set of the evolving hypersurface. In the figure three steps in the evolution are shown.

In each 2x2 set of images an axial, sagittal and coronal slice is shown and in the lower right cell of the 2x2 grid the 3D triangulated surface is shown. The yellow overlay is the MAP estimate of shape and pose. These segmentation results took approximatesly 6 minutes.



[Leventon00] Figure A12: 3D segmentations on the T7 vertebra

Quantitative analysis of the segmentation results was measured with the undirected partial Hausdorff distance between the boundary of the computed segmentation and a manually segmented ground truth. This distance is defined as:

$$H_k(A, B) = \max(h_k(A, B), H_k(B, A)) \quad (A19a)$$

where the directed partial Hausdorff distance is $h_k(A, B) = k^{th} \min_{a \in A, b \in B} \|a - b\|$. The results are shown in table 1

K	Corpus 1	Corpus 2	Vertebra
95%	1.3 mm	1.5 mm	2.7 mm
99%	1.6 mm	2.0 mm	4.4 mm

[Leventon00] Table A1: Undirected Partial Hausdorff distance for 2 corpus callosum and the T7 vertebra

5.6 Critical Analysis of the Results and Method

The overall method of employing a shape model which is represented by a Gaussian distribution over shape parameters has intuitive appeal. Expert segmenters rely upon such information when they are segmenting a shape manually and the resultant manual segmentation is a curve which

has a fair amount of data fidelity, smoothness and employs prior knowledge of the typical shape of the object where clear boundary information is absent in the image.

Overall, the segmentation results look good and the table 1 shows that almost all of the boundary points lay within 1 or 2 2mm voxels of the ground truth. In the test results shown in shown figure A11 bottom right image, where there was no boundary the constrained deformation did the job that its was intended to do. Namely it constrained the segmentation where there was a lack on boundary information but allowed the boundaries which were part of the shape to be used in the final segmentation.

The real object tests which Leventon describes in this paper do not include inter-subject training sets, instead the training sets come from the same subject as the test cases. For example the vertebrae from the same subject are divided in to a training set and a set of test cases. In order to better understand the applicability of the method to segmenting structures from *novel* subjects it would be better to test the method on structures from novel subjects which are not in the training set.

Embedding the training curves as the zero level set of a higher dimensional surface or distance map u and subsequently representing u as a vector in \mathfrak{R}^{N^d} effects two positive results. First it enables a point cloud view of shape which can be analyzed using SVD and parameterized by a multivariate Gaussian. Second it induces the notion of a directional vector which represents shape change. This allows shape to be represented as a vector and shape change as standard vector addition albeit in a very high dimensional space. One drawback of this shape representation is that it has many parameters even in the k -vector (α) form and the parameters do not correspond to a meaningful notion of shape change such as twisting or bending [J.Park,D.Metaxas].

The paper leaves out a few details such as how many degrees of freedom where used to parameterize the shape model in order to achieve the illustrated segmentation results and the function $g()$ is only defined as a function of gradient but not defined precisely. Pose representation is not discussed. Leventon does not say which curvature(s) uses for a 3D surface in equation A7.

6 A Model-based Segmentation Method Using Medial Axis Shape Representation


Leventon's method utilizes a shape parameterization based on eigenshapes which makes the representation less intuitive than one based on a more familiar geometric shape representation. In [Pizer99] the later approach is taken for shape representation. Dominant shape aspects are represented by a set of medial points, while fine-scale boundary aspects are represented by a set of boundary points. The points in each set are linked in certain ways and then compound shape are formed by linking multiple sets to one another. Information about the angles between the links is also stored. In this way a more geometrically intuitive shape representation is formed. In [Pizer99] the 2D methods presented in [Fritsch97] are generalized to 3D. Pizer's method is composed of a similar set of salient aspects which will be discussed:

- Parameterization of shape and pose and formation of shape probability distribution models from parameterized training instances
- A method for estimating which shape and pose from the distribution to guide the evolution
- segmenting novel structures using the shape model

6.1 The Shape Model

6.1.1 Gross Shape Characteristics are represented using Medial and End Primitives

For a single 2D shape, the bounding surface is a curve and the medial component is a 1-manifold. For a single 3D shape the bounding surface is a surface in R^3 and the medial component is a 2-manifold. A *medial primitive* is a medial point with two boundary-pointing vectors of equal length. As shown in figure D2, let the boundary of a portion of a shape be that shown in (a).



[Pizer99] Figure D2: Medial Primitives imply a bounding surface for a shape

To represent this shape, a linked list of medial primitives is formed as shown in (c). Each white circle is a medial primitive. It is shown in detail in (b) and is composed of several parameters including: the medial point vector position, \bar{x} , the scalar distance to the boundary "r", a unit vector \bar{b} bisecting the boundary-pointing vectors and an angle θ between \bar{b} and the two boundary pointing vectors. The boundary points can be recovered when necessary using:

$$\bar{x} + rR(\mathbf{q})\bar{b} \text{ and } \bar{x} + rR(-\mathbf{q})\bar{b} \quad (\text{D1})$$

where R is the familiar rotation matrix. The boundary pointing vectors define a boundary fragments normal to the vectors.

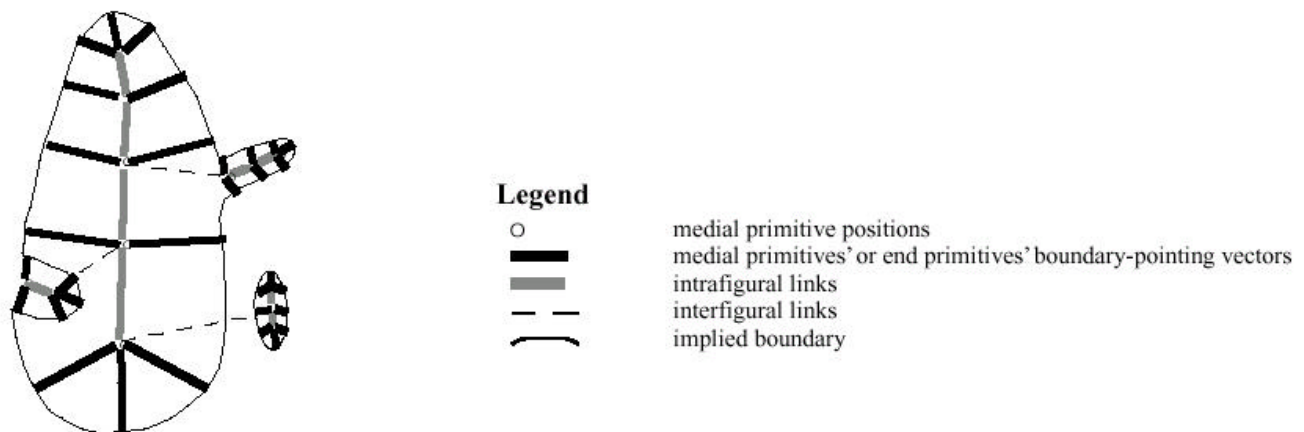
A *figure* (shape) is the geometric object implied by a sequence of medial primitives such that each of the two implied boundary point sequences given by equation (D1) is ordered along the implied boundary in the same order as the primitives. The sequence may be closed on either or both ends by an end primitive.

End primitives are used to cap each end of the list of medial primitives. They are similar to medial primitives however they have a third boundary-pointing vector. An end primitive is parameterized by:

$$\bar{x} + rR(\mathbf{q})\bar{b} \text{ and } \bar{x} + rR(-\mathbf{q})\bar{b} \text{ and } \bar{x} + \mathbf{a}r\bar{b} \quad (\text{D2})$$

6.1.2 Coarse to Fine Representation

In a manner similar to the computer graphics method called constructive solid geometry [Foley90] complex shapes can be represented by adding and subtracting shapes.



[Pizer99] Figure D3: A complex shape together with a nearby shape is described by linking simple shapes forming an acyclic graph

As shown in figure D3, medial primitives can also be linked across figures. One of the figures is considered the parent and the others are considered children, and the whole structure forms a DAG. The structure of the graph is taken to indicate its topology and is assumed to be deterministic. Moreover each relation between a figure and a subfigure (child) is represented by a single link between a medial point on the parent shape to an endpoint of the child.

6.1.3 Boundary Primitives

The actual boundary of the shape is considered to be a fine scale detail. The boundary is defined at a small scale by displacements relative to, and along the normal to, the mean boundary implied by the linked list of medial primitives. Thus *boundary primitives* are scalar displacements. A boundary *net* represents the surface in R^3 and is formed by “piecing together” boundary primitives between subfigures and parent figures. Details of the boundary tessellation process is given in [Chen98] and is beyond the scope of this analysis.

6.1.4 Manual Construction of the Shape Model

Expert manual segmentation is used to construct the shape model. The expert traces the contour of the shape to be segmented and software computes the medial primitives and connects them to the boundary primitives sampled from the trace. The training patient data sets are aligned and scaled, and then the mean and variance of the location of the medial points is computed in an object centered coordinate frame. In addition the mean and variance of the angle between the medial axis line segments and boundary points is computed and stored in the model.

6.2 Segmenting Novel Structures Using the Shape Model

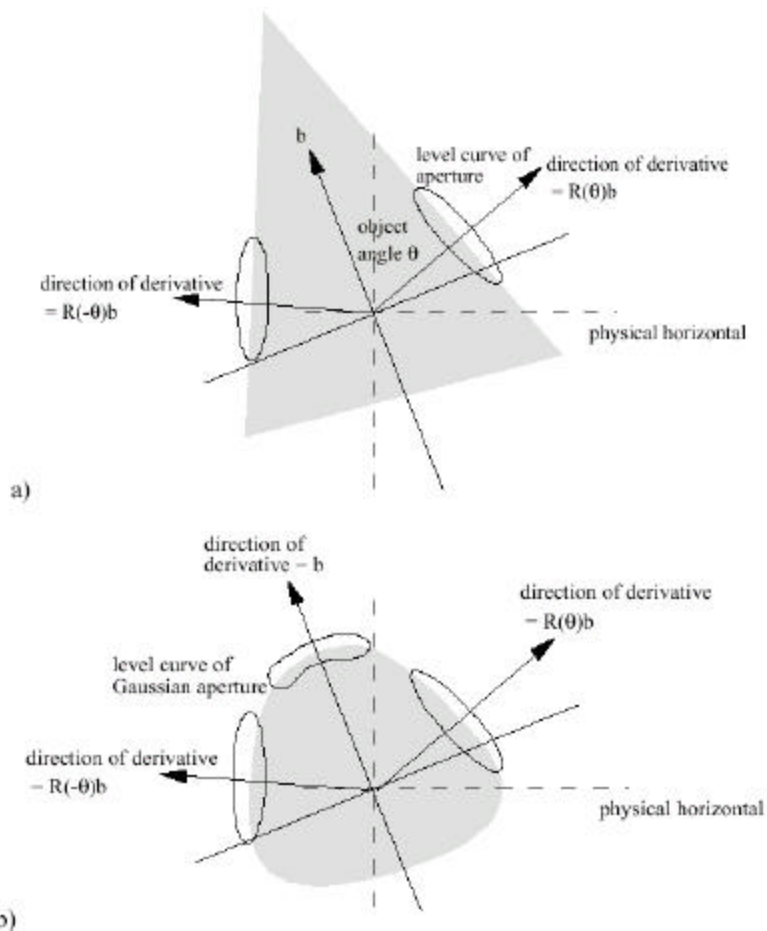
In [Fritsch97], Pizer and Fritsch present a Bayesian method for segmenting objects based on this shape representation method. In a method similar to Leventon's method, Bayes rule is used to express a posteriori probabilities in terms of a priori probabilities. The a posteriori probability is the probability of the shape parameters given the image, the prior probability is the probability of the shape parameters, and the conditional probability term is the density of the image given a set of shape parameters. Taking the log of this relation yields the objective function:

$$-\log P = w_b \sum_{i=1}^{N_b} B(b_i) + w_m \sum_{i=1}^{N_m} M(m_i) - w_l \sum_{i=1}^{N_l+1} \|l_i - l_i^{ref}\| \quad (D3)$$

to optimize. Finding the shape parameters which optimize equation (D3) yield the parameters of the shape which are most likely given the image. These are the shape parameters which define the final segmentation. Briefly stated, the first two terms $\sum_{i=1}^{N_b} B(b_i)$ and $\sum_{i=1}^{N_m} M(m_i)$ measure the density of the image given the a particular shape. The former measures the match between the image and boundary primitives while the latter measures the match between the image and medial primitives. The third term, $\sum_{i=1}^{N_l+1} \|l_i - l_i^{ref}\|$, measures the prior probability of the model and contains no image information. In these equations

- w_b, w_m, w_l are scalar weighting factors which give the relative strength of the terms
- b_i = the displacement of the i^{th} boundary primitive (a scalar)
- $m_i = [x_i, r_i, \bar{b}_i, \mathbf{q}_i]^T$ is the i^{th} medial primitive
- $l_j = \left[\frac{\Delta x_j}{r_j}, \frac{\Delta r_j}{r_j}, \Delta \bar{b}_j, \mathbf{q}_j, \Delta \mathbf{q}_j \right]^T$ is the j^{th} shape primitive link
- $M(m)$ is a measure of the agreement of the image with medial primitives

$$M(m) = (R(\mathbf{q})\bar{b}) \cdot \nabla A(y, r) I(x + rR(\mathbf{q})\bar{b} + y) + (R(-\mathbf{q})\bar{b}) \cdot \nabla A(y, r) I(x + rR(-\mathbf{q})\bar{b} + y)$$
 where A is an aperture function as shown in figure D7
- $B(\cdot)$ is a statistical measure of agreement of the image with the model in the normal direction and at a smaller scale than used in medial primitive match measurements




[Pizer99] Figure D7: Aperture function for (a) medial primitive, (b) end primitive

6.3 Critical Analysis of Results and Method

6.3.1 Results

The results in this paper are qualitative in nature. In [Fritsch97] the automated segmentation method is used successfully on a variety of objects including the left and right ventricles of the brain in MRI data as well as abdominal structures such as the kidneys in CT data. The speed of the algorithm on 2D shape segmentation varied from 20 seconds to a few minutes depending on the complexity of the figure to segment. In general the fitting algorithm runs in order $O(N)$ time where N is the number of loci: medial + boundary points. Three steps in the segmentation of part of the brain stem is shown in figure D11. Part (a) shows the initial shape and pose of the model.



[Pizer99] Figure D11: Brain stem segmentation

Part (b) shows an intermediate step in the segmentation showing how the medial points are aligned to the image first. The medial points tend to vary the least from subject to subject so their alignment gives a crude initial segmentation. Next, the boundary points are aligned by moving them iteratively until the gray level profile (along the boundary line segments) matches that of the model. These boundary point movements are constrained by the deformation mean and variances in the shape model. A quantitative comparison with expert segmented brain stem is not presented in the paper. This is a significant omission and make comparisons with other methods difficult.

6.3.2 A Shape's Deformable Shape Loci Representation is Not Unique

Given Pizer's shape representation method (a.k.a. DSL), distinct DSL descriptions can be formed from the same shape. For example, consider an object which is shaped like the human hand. It could be modeled as a parent palm shape with 5 children shapes representing the fingers and thumb and the extensions of the fingernails represented as boundary displacements. However the fingernails could be instead be modeled as subfigures of the fingers, or, the hand itself could be modeled as one parent mitten shape with 4 subfigures representing the 4 concavities between the fingers and thumb. Having multiple representations for the same shape makes comparing shapes difficult.

6.3.3 Segmentation Limited to Shapes with Deterministic Topology

Like Leventon's shape segmentation method, this method can segment many shapes – theoretically any shape that has a non-degenerate medial manifold, which includes most if not all anatomical shapes of interest, could be a candidate for the representation and subsequent segmentation using Pizer's method. However, in order for the model to work, interfigure homology is required across subjects. This means that a given shape may **not** have changes in the number of protrusions, holes, nor in the ordering of the protrusions on the medial manifold. Many anatomic shapes fall into this latter category. For example: blood vessel trees, bronchial trees and the volumetric cortical layer do not exhibit interfigural homology across subjects. Instead the subfigure structure is stochastic rather than deterministic. Since Pizer's method does not provide for the addition/subtraction of subfigures or their reordering, these shapes would not be segmented properly.

6.3.4 Psychophysical Basis

This shape representation has *much intuitive appeal* and it has been shown to be closely linked to the way in which humans perceive and recognize shape. An intuitive example is that when we say something is kidney bean shaped – it has an overall appearance, which is mostly captured by the

midline followed by a minor appearance, which is captured by the boundary points. Slight deviations in the boundary points generally do not change our description of the shape.

6.3.5 Space and Time are $O(N)$ but N is Increases with Accuracy Used in Shape Model

Pizer claims that the deformable shape loci method of shape representation is simultaneously is a rich shape description and efficient. Instead it seems that the more accurately one wishes to represent a complex shape, the more shape primitives are needed. Pizer does not state what metric he is measuring the efficiency and which alternative shape representation he is comparing. To illustrate this point, Pizer claims that the representation requires only an amount of space proportional to the number of primitives. I believe this statement is correct, but hardly something to boast about: one can also approximate a sphere with a icosahedron tessellation as shown in figure B1b and the amount of space required for the sphere approximation is proportional to the number of primitives used. However as the approximation error goes to zero (we could measure error say as the maximum distance from the sphere boundary to the tessellated sphere) the number of primitives goes to infinity. It seems that for a shape with a complex boundary the deformable shape loci would also require more and more loci to drive the error to zero, which doesn't constitute an space efficiency advantage over even a finite tessellation.

7 General Conclusions

7.1 Comparison of the methods

7.1.1 Pizer's Method vs Leventon's Method

Ease of initialization appears to be better in Leventon's method: any point inside the interior can be used to initialize the evolution, while in Pizer's method the model must be initialized close to the final result: the capture range is typically about half the width of the parent figure.

Pizer's method is somewhat MR resolution invariant while Leventon's would require either a retraining at the test resolution or resampling the test image to the resolution of the shape model.

In certain cases, training could be automated in Leventon's method. Certain structures, such as the vertebrae, are relatively easy to segment using a global threshold when good MR pulse sequence is used. Thus training shapes could be nearly automatically generated from simple global thresholding on intensity. These shapes could be used for testing and tuning Leventon's algorithm without having to trace boundary by hand as in Pizer's method.

Both the distance map and deformable shape loci shape representations can represent most any shape. In the distance map all constituent points can be treated alike while the deformable shape loci contain a heterogeneous mixture of primitives which complicates the segmentation strategy. However the richer deformable shape loci shape description has essential feature points encoded in it which can be used to solve the correspondence problem, provided corresponding feature points exist in the structures to match. The feature points which are present do not have necessarily anatomical or physiological significant but they could be coerced to have some relation during the model construction stage.

Both methods use prior shape information in a weighted fashion so that there exists the possibility of the segmentation of a pathological shape if the image information is strong enough.

7.1.2 Overall Comparison

An overall comparison of the segmentation methods is given in the table below.

Feature	Zeng	Gomes	Pizer	Leventon
Large Capture Range?	Yes	Yes	No	Yes
Can Segment a Wide Range of Shapes?	No	No	Yes	Yes
Training Required?	No	No	Yes	Yes
Can Shapes be compared with the method?	No	No	Yes – including where the shape change occurred	Yes - using distances in eigenshape space
Surface Curvature readily recoverable	Yes	Yes	No	Yes
Extracts Medial Manifold during segmentation?	No	Yes	Yes	Yes - if modified with New PDE

Green indicates an advantage, red a disadvantage.

8 References

- [Adalsteinsson95] D. Adalsteinsson, and J Sethian, "A Fast Level Set Method for Propagating Interfaces", *J. Computational Physics*, 118(2):269-277, 1995
- [Arnold83] V. Arnold, "Geometrical Methods in the Theory of Ordinary Differential Equations", Springer Verlag, New York, 1983
- [Barles93] G. Barles, H. Soner, P. Souganidis, "Front propagation and phase field Theory" *SIAM J. Control and Optimization*, 31(2):439-469, March 1993
- [Caselles95] V. Caselles, R. Kimmel, and G. Sapiro. Geodesic active contours. In *Proceedings of the 5th International Conference on Computer Vision*, 694-699, Boston, MA, June 1995, IEEE Computer Society Press
- [Caselles97] V. Caselles, R. Kimmel, and G. Sapiro. Geodesic active contours. *Intl'l J. Comp Vision*, 22(1):61-79, 1997
- [Chen98] D. Chen, "Volume Rendering Guided by Multiscale Medial Models", PhD Dissertation, Univ. of NC Dept of Comp Science, 1998
- [Cline90] H. Cline, W Lorensen, R. Kikinis, F. Jolesz, "Three-dimensional segmentation of MR images of the head using probability and connectivity", *Journal of Computer Assisted Tomography*, 14(6):1037-1045, Nov 1990
- [Davazikos96] C. Davazikos and R. Bryan, "Using a deformable surface model to obtain shape representation of cortex", *IEEE Transactions on Medical Imaging*, 15(6):785-795, 1996
- [Foley90] J. Foley, A. van Dam, S. Feiner, J. Hughes, "Computer Graphics", Addison-Wesley Pub. Co., 1990, pp 557-558
- [Ftitsch97] D. Fritsch, et al, "Localization and Segmentation of Medical Image Objects using Deformable Shape Loci", *Image Processing in Medical Imaging, Lecture Notes in Computer Science*, Volume 1230, Springer, 1997
- [Gomes00] J. Gomes, O. Faugeras, "Reconciling Distance Functions and Level Sets", *J. of Visual Communication and Image Representation* 11:209-223, 2000
- [Jain95] R. Jain, R. Kasturi, B. Schunck "Machine Vision", McGraw-Hill Inc., 1995
- [Kass88] M. Kass, A Witkin, D Terzopoulos, SNAKES: Active Contour Models. *The International J. Of Computer Vision*, 1:321-332, January 1988
- [Leventon00] M. Leventon, E. Grimson, et al. "Statistical Shape Influence in Geodesic Active Contours" *Comp. Vision and Patt. Recon. (CVPR)*, 2000.
- [Osher88] S. Osher, J. Sethian, "Fronts propagating with curvature dependent speed: algorithms based on the Hamilton-Jacobi formulation", *J of Computational Physics*, 79:12-49, 1988

[Pizer99] S. Pizer, D. Fritsch, P. Yushkevich, V Johnson, E Chaney, G Gerig. "Segmentation, Registration, and Measurement of Shape Variation via Image Object Shape". IEEE TMI, October 1999

[Sethian96] J. Sethian, "Level Set Methods", Cambridge University Press, 1996

[Sethian98] J. Sethian, "Level Set Methods for Shape Recovery",
<http://math.berkeley.edu/~sethian/levelset.html>, 1998

[Zeng98] X. Zeng, L.H.Staib, R.T.Schultz and J.S.Duncan, "Volumetric Layer Segmentation Using Coupled Surfaces Propagation", Comp. Vision and Patt. Recon. (CVPR) 1998

[Zeng98] X. Zeng, L. Staib, et al, "Segmentation and Measurement of the Cortex from 3D MR Images", Proceedings of the conference on Medical Image Computing and Computer Assisted Intervention, 1998