A Framework for Level Set Segmentation of Volume Datasets

Ross Whitaker
School of Computing
University of Utah

David Breen, Ken Museth, and Neha Soni
Computer Science Department
California Institute of Technology

Abstract

This paper presents a framework for extracting surface models from a broad variety of volumetric datasets. These datasets are produced from standard 3D imaging devices, and are all noisy samplings of complex biological structures with boundaries that have low and often varying contrasts. The level set segmentation method, which is well documented in the literature, creates a new volume from the input data by solving an initial-value partial differential equation (PDE) with user-defined feature-extracting terms. Given the local/global nature of these terms, proper initialization of the level set algorithm is extremely important. Thus, level set deformations alone are not sufficient, they must be combined with powerful initialization techniques in order to produce successful segmentations. Our level set segmentation approach consists of defining a set of suitable pre-processing techniques for initialization and selecting/tuning different feature-extracting terms in the level set algorithm. This collection of techniques forms a toolkit that can be applied, under the guidance of a user, to segment a variety of volumetric data. Users can combine these methods in different ways and thereby access a wide range of functionalities, several of which are described in this paper and demonstrated on noisy volume data.

1 Introduction

As visualization tasks grow in size and complexity, the problem of presenting data effectively is accompanied by another, potentially more difficult problem—how to extract *presentable* data from the flood of raw information produced by large simulations and high resolution instruments. Thus, the field of *data visualization* is intimately tied to more traditional studies of *data analysis* such as signal and image processing, pattern recognition, artificial intelligence, and computer vision. However, in contrast to more conventional areas of data analysis, the field of visualization explicitly *includes the user* in the process of filtering, extracting, and rendering meaningful data.

This paper deals with a specific visualization problem—that is, how to build meaningful 3D models of complex structures from noisy datasets generated from 3D imaging devices. In certain circumstances such data can be visualized *directly* [1, 2, 3, 4]. While direct techniques can provide useful insights into volume data, they are insufficient for many problems. For instance, direct volume rendering techniques typically do not remove occluding structures, i.e., they do not allow one to "peel back" the various layers of the data to expose the inner structures that might be of interest. They also do not generate the models needed for quantitative study/analysis of the visualized structures. Furthermore, direct visualization techniques typically do not perform well when applied directly to noisy data, unless one filters the data first. Techniques for filtering noisy data are abundant in the literature, but there is a fundamental limitation—filtering that reduces noise tends to distort the shapes of the objects in the data. The challenge is to find methods which present the best tradeoff between fidelity and noise.

Level set segmentation relies on a surface-fitting strategy, which is effective for dealing with both small-scale noise and smoother intensity fluctuations in volume data. The level set segmentation method, which is well documented in the literature [5, 6, 7, 8], creates a new volume from the input data by solving an initial-value partial differential equation (PDE) with user-defined feature-extracting terms. Given the local/global nature of these terms, proper initialization of the level set algorithm is extremely important. Thus, level set deformations alone are not sufficient, they must be combined with powerful initialization techniques in order to produce successful segmentations. Our level set segmentation approach consists of defining a set of suitable pre-processing techniques for initialization and selecting/tuning different feature-extracting terms in the level set algorithm. We demonstrate that combining several pre-processing steps with level set deformations produces a powerful toolkit that can be applied, under the guidance of a user, to segment a wide variety of volumetric data.

There are more sophisticated strategies for isolating meaningful 3D structures in volume data. Indeed, the so called *segmentation problem* constitutes a significant fraction of the literature in image processing, computer vision, and medical image analysis. For instance, statistical approaches [9, 10, 11, 12] typically attempt to identify tissue types, voxel by voxel, using a collection of measurements at each voxel. Such strategies are best suited to problems where the data is inherently multi-valued or where there is sufficient prior knowledge [13] about the shape or intensity characteristics of the relevant anatomy. Alternatively, anatomical structures can be isolated by grouping voxels based on local image properties. Traditionally, image processing has relied on collections of edges, i.e. high-contrast boundaries, to distinguish regions of different types [14, 15, 16]. Furthermore deformable models, incorporating different degrees of domain-specific knowledge, can be *fitted* to the 3D input data [17, 18]. The work of this paper demonstrates a mathematical and computational framework which effectively combines or unifies classification, filtering, and surface-fitting approaches to modeling and visualizing 3D data.

2 Example Datasets

Our work is largely motivated by the desire to produce a semi-automatic segmentation approach which can partly or fully replace the tedious and extremely time-consuming process of manual data segmentation – a solution which to our initial surprise is widely used by colleagues in biology and medicine. Thus, to scientists working in these fields even an approximate scheme which can segment out approximately 90% of the model is immensely useful because it reduces the manual labor needed to produce a final result. We stress that there exists no fully automatic solution to the segmentation problem typically encounter in 3D imaging. For example, Figure 1(a) shows one of 270 slices of an electron tomography (ET) volume of a *spiny dendrite* provided by the National Center for Microscopy and Imaging Research, at UC San Diego. The complex structure of the dendrite and the noisy nature of the data make the rendering of such volume data difficult. Figure 1(b) shows the results of attempting to isolate the relevant structures in this dataset by extracting isosurfaces at greyscale value of 129. For this example we have blurred the data with a small Gaussian ker-

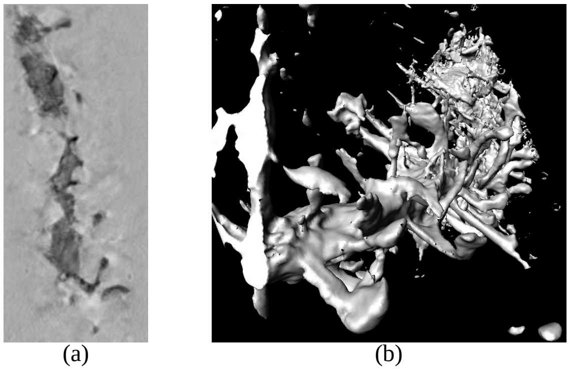

Figure 1: a) One slice of a 154 × 586 × 270 ET scan of a spiny dendrite shows low contrast and high noise content in a relatively complex data set. b) An isosurface rendering, with prefiltering, shows how noise and inhomogenieties in density interfere with visualizing the 3D structure of the dendrite.

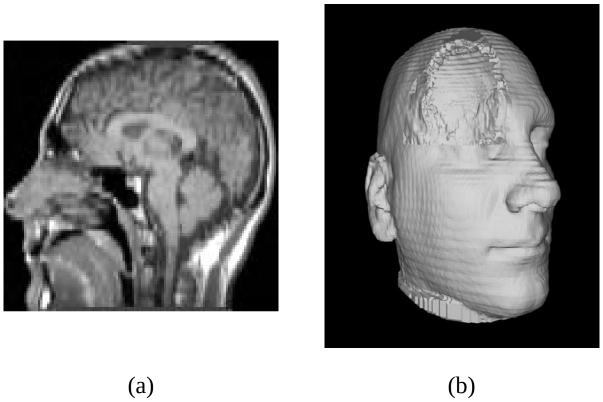

Figure 2: a) One slice of a 130 × 128 × 128 magnetic resonance (MR) volume of a human head shows high-contrast, relatively noise-free data with numerous internal structures. b) An isosurface rendering, with a small wedge removed for visualization, shows aliasing and internal structures that are not appropriate for the application.

nel ($\sigma = 1.0$) to try to improve the appearance of the isosurfaces. Despite the smoothing the isosurfaces are quite noisy, and contain many small, disconnected pieces that are not indicative of the structure of the dendrite. Furthermore, fluctuations in the tissue density both within and outside of the dendrite create a large number of distortions which prevent the isosurface from accurately representing the underlying shape of that structure.

Note that the image shonw in Figure 1(b) is produced in two stages: First, we compute the isosurface with the Marching Cubes algorithm [19] for a given isovalue. Next, the polygonal mesh is displayed using conventional graphics hardware. Alternatively we could visualize structures within the volume data using a one-stage direct method such as volume-rendering (e.g. ray casting with transfer functions or maximum intensity projection). Our choice of Marching Cubes for rendering isosurfaces of this and other datasets in this paper is not essential to the proposed method. The problems of noise and aliasing, present in the examples in this paper, would exist even if we used a direct volume rendering technique.

A second example, shown in Figure 2(a), is a magnetic resonance (MR) scan of a human head. Here the problem is not so much the quality of the data—isosurfaces can be used to visual-

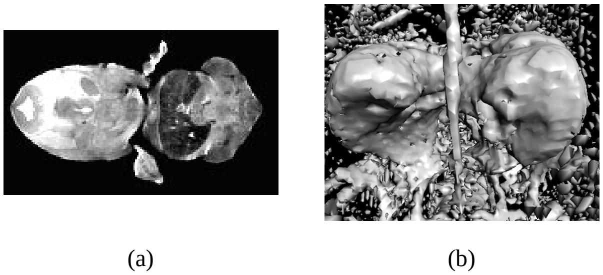

Figure 3: a) One slice of a 130 × 128 × 128 magnetic resonance (MR) volume of a human head shows high-contrast, relatively noise-free data with numerous internal structures. b) An isosurface rendering, with a small wedge removed for visualization, shows aliasing and internal structures that are not appropriate for the application.

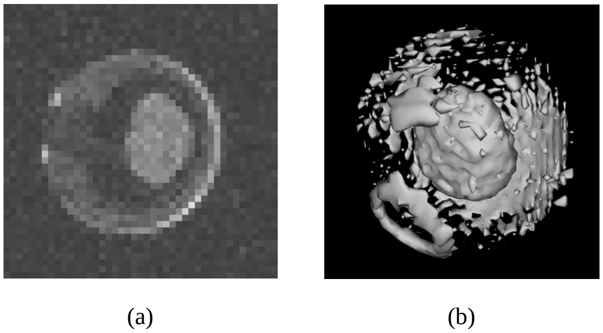

Figure 4: a) One slice of a 44 × 45 × 43 MR scan of a frog embryo. b) A Marching Cubes isosurface from the frog embryo volume. Isovalue = 60.

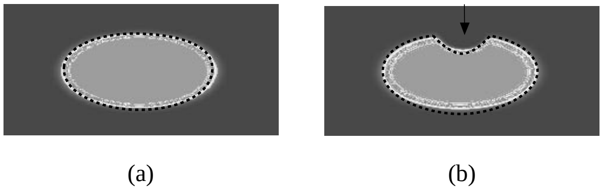

Figure 5: a) Level set models represent curves and surfaces implicitly using greyscale images. For example an ellipse is represented as the level set of an image shown here. b) To change the shape of the ellipse we modify the greyscale values of the image by solving a certain PDE.

ize the skin or skull. In this case the particular application [20] requires a relatively smooth, simple, closed surface, and will not tolerate significant aliasing. The application also requires that the fairly complicated structure of the inner head (usually unseen) be removed. Figure 2(b) shows a isosurface rendering, at a greyscale value of 30, which demonstrates the aliasing in the data. A small wedge has been removed to show the complex internal structures in this volume.

The third example, shown in Figure 3(a), is a 256 × 128 × 128 MR scan of a 12-day-old mouse embryo. Colleagues in the Caltech Biological Imaging Center (BIC) are using such images to develop a detailed atlas for the gestational development of these organisms. For this paper we will consider the specific task of isolating the liver, which is the dark, kidney-shaped area on the right. The liver, however, is not a single grey-scale value, and it is bordered by both more dense and less dense regions. Furthermore, the data contains noise. Therefore, the liver is not easily isolated by simple greyscale classification or isovalue schemes. Figure 3(b) shows an isosurface rendering which accommodates high and low thresholds associated with the liver, i.e., the zero crossings of $I' = \min(I - t_{\mathrm{lo}}, t_{\mathrm{hi}} - I)$, where I is the input volume, and $t_{\mathrm{hi/lo}}$ are the thresholds. The model constructed from the isosurface shows significant artifacts from noise and low-frequency fluctuations in the tissue. It also shows artifacts from the greyscale classification, which captures a large number of voxels in the transition between the skin and the surrounding regions. Smoothing further aggravates this problem.

The final example, shown in Figure 4(a), is a 44 × 45 × 43 *in vivo* MR scan of a frog embryo. This is one slice from one scan of a sequence of 22 volumes taken over a 24-hour period. Colleagues at the Caltech BIC are acquiring time-lapsed MR volume sequences in order to generate the first 3D staging sequence of a developing frog embryo. They require models of the dynamic structures that appear, move, change shape, merge and/or disappear over time within the embryo, as well as the outside shell. The individual structures do not necessarily have distinct signals in the MR scans, thus making it difficult to computationally isolate them. Figure 4(b) presents a polygonal isosurface (isovalue = 60) generated with the Marching Cubes algorithm [19]. At this isovalue two internal structures are produced, as well as a significant part of the outer shell.

3 Level Set Surface Models

When considering deformable models for segmenting 3D volume data, one is faced with a choice from a variety of surface representations, including triangle meshes [20, 21], superquadrics [22, 23, 24], and many others [18, 25, 26, 27, 28, 29, 30]. Another option is an implicit level set model, i.e., specifying the surface as a *level set* of a scalar volumetric function, $\phi : U \mapsto \mathrm{I\!R}$, where $U \subset \mathrm{I\!R}^3$ is the range of the surface model. Thus, a surface S is

$$S = \{s | \phi(s) = k\}, \tag{1}$$

and the choice of the isovalue, k, is arbitrary. In other words, S is the set of points s in $\mathrm{I\!R}^3$ that composes the k isosurface of ϕ. The embedding ϕ can be specified as a regular sampling on a rectilinear grid.

Our overall scheme for segmentation is largely based on the ideas of Osher and Sethian [31] that model propagating surfaces with (time-varying) curvature-dependent speeds. The surfaces are viewed as a specific level set of a higher-dimensional function ϕ – hence the name level set methods. These methods provide the mathematical and numerical mechanisms for computing surface deformations as isovalues of ϕ by solving a partial differential equation on the 3D grid. That is, the level set formulation provides a set of numerical methods that describes how to manipulate the greyscale values in a volume, so that the isosurfaces of ϕ move in a prescribed manner (shown in Figure 5). This paper does not present a comprehensive review of level set methods, but merely introduces the basic concepts and the notation used in successive sections. See [7] for more details.

There are two different approaches to defining deformable surface from a level set of a volumetric function as described in Equation 1. Either one can think of $\phi(s)$ as a *static* function and change the isovalue $k(t)$ or alternatively fix k and let the volumetric function *dynamically* change in time, i.e. $\phi(s,t)$. Thus, we can mathematically express the static and dynamic model respectively as

$$\phi(s) = k(t) \tag{2a}$$

$$\phi(s,t) = k. \tag{2b}$$

To transform these definitions into partial differential equations which can easily be solved by standard numerical techniques, we differentiate both sides of Equation 2 with respect to time t, and apply the chain rule:

$$\nabla\phi(s)\frac{\mathrm{d}s}{\mathrm{d}t} = \frac{\mathrm{d}k(t)}{\mathrm{d}t} \tag{3a}$$

$$\frac{\partial\phi(s,t)}{\partial t} + \nabla\phi(s,t)\cdot\frac{\mathrm{d}s}{\mathrm{d}t} = 0 \tag{3b}$$

The static Equation 3a is often referred to as the "Eikonal" equation and defines a boundary value problem for the time-independent volumetric function ϕ. This static level set approach has been solved [32, 33] using a "Fast Marching Method". However, it has some inherent limitations following the simple definition in Equation 2a. Because ϕ is a function (i.e. single-valued), isosurfaces cannot self intersect over time, i.e. shapes defined in the static model are strictly expanding or contracting over time. The dynamic level set approach of Equation 3b is much more flexible and shall serve as the basis of the segmentation scheme in this paper. Equation 3b is sometimes referred to as a "Hamilton-Jacobi-type" equation and defines an initial-value problem for the time-dependent ϕ. Throughout the remainder of this paper we shall for simplicity refer to this dynamical approach as the level set method – and completely ignore the static alternative.

Thus, to summarize the essence of the (dynamic) level set approach; let $\mathrm{d}s/\mathrm{d}t$ be the movement of a point on a surface as it deforms, such that it can be expressed in terms of the position of $s \in U$ and the geometry of the surface at that point, which is, in turn, a differential expression of the implicit function, ϕ. This gives a partial differential equation on ϕ: $s \equiv s(t)$

$$\frac{\partial\phi}{\partial t} = -\nabla\phi\cdot\frac{\mathrm{d}s}{\mathrm{d}t} \equiv -\nabla\phi\cdot F(s, \mathrm{D}\phi, \mathrm{D}^2\phi, \ldots), \tag{4}$$

where F is a user-defined "speed" term which depends on a set of order-n derivatives of ϕ, $D^n\phi$, evaluated at s, as well as other functions of s. Because this relationship applies to every level set of ϕ, i.e. all values of k, this equation can be applied to all of U, and therefore the movements of *all* the level set surfaces embedded in ϕ can be calculated from Equation 4. Such level set methods are well documented in the literature for applications such as computational physics [34], image processing [35, 36], computer vision [6, 37], medical image analysis [6, 38], shape morphing[39], and 3D reconstruction [40].

The level set representation has a number of practical and theoretical advantages over conventional surface models, especially in the context of deformation and segmentation. First, level set models are topologically flexible, they easily represent complicated surface shapes that can, form holes, split to form multiple objects, or merge with other objects to form a single structure. These models can incorporate many (millions) of degrees of freedom, and therefore they can accommodate complex shapes such as the dendrite in Figure 1.

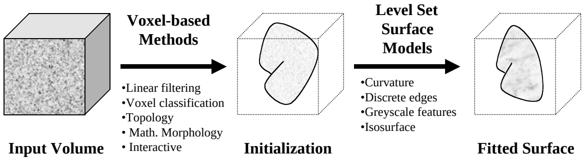

Figure 6: Level set segmentation stages – initialization and surface deformation.

Indeed, the shapes formed by the level sets of ϕ are restricted only by the resolution of the sampling. Thus, there is no need to reparameterize the model as it undergoes significant changes in shape.

Level set methods have been shown to be effective in extracting surface structures from biological and medical data. For instance Malladi et al. [6] propose a method in which the level sets form an expanding or contracting contour which tends to "cling" to interesting features in 2D angiograms. At the same time the contour is also influenced by its own curvature, and therefore remains smooth. Whitaker has shown [38, 41, 42] that level sets can be used to simulate conventional deformable surface models, and demonstrated this by extracting skin and tumors from *thick-sliced* (e.g. clinical) MR data, and by reconstructing a fetal face from 3D ultrasound. A variety of authors [36, 43, 44] have presented variations on the method with results for 2D and 3D data. Sethian [7] gives several examples of level set curves and surface for segmenting CT and MR data.

The purpose of this paper is to present a collection of initialization and level set mechanisms which form a "toolbox" for volume dataset segmentation. We also show how these methods can be combined to solve the problems presented in Figures 1–4. These tools provide a set of techniques that are not as direct as simple thresholding or volume rendering but are more powerful than the "hand-contouring" that is currently the state-of-the-art in many applications, such as the dendrite example in Figure 1.

4 Segmentation Stages

Our level set segmentation process has two major stages, initialization and level set surface deformation, as seen in Figure 6. Each stage is equally important for generating a correct segmentation. Within our framework a variety of operations are available in each stage. A user must "mix-and-match" these operations in order to produce the desired result.

4.1 Initialization

Because the deformable models move using gradient descent, they seek *local solutions*, and therefore the results are strongly dependent on the initialization, i.e., the starting position of the surface. Thus, one controls the nature of the solution by specifying an initial model from which the surface deformation process proceeds. We have implemented both computational (i.e. "semi-automated") and manual/interactive initialization schemes; each offers distinct advantages in different situations.

4.1.1 Computational Initialization

Because the level set modeling technology is based on the isosurfaces of volumes, we can, for many different types of problems, computationally construct reasonable initial estimates directly from the input data. We do this by combining a variety of techniques.

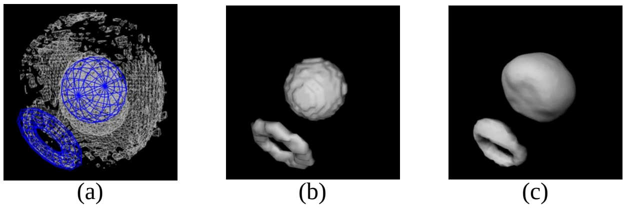
(a) (b) (c)

Figure 7: a) *[color]* Interactively positioning a CSG model relative to a Marching Cubes mesh. b) Isosurface of a binary scan conversion of the initialization CSG model. c) Final internal embryo structures.

Linear filtering: We can filter the input data with a low-pass filter (e.g. Gaussian kernel) to blur the data and thereby reduce noise. This tends to distort shapes, but the initialization need only be approximate.

Voxel classification: We can classify pixels based on the filtered values of the input data. For greyscale images, such as those used in this paper, the classification is equivalent to high and low thresholding operations. These operations are usually accurate to only voxel resolution (see [12] for alternatives), but the deformation process will achieve sub-voxel results.

Topological/logical operations: This is the set of basic voxel operations that takes into account position and connectivity. It includes unions or intersections of voxel sets to create better initializations. These logical operations can also incorporate user-defined primitives. Topological operations consist of connected-component analyses (e.g. flood fill) to remove small pieces or holes from objects.

Morphological filtering: This includes binary and greyscale morphological operators on the initial voxel set. For the results in the paper we implement openings and closings using *morphological propagators* [45, 46] implemented with level set surface models. This involves defining offset surfaces of ϕ by expanding/contracting a surface according to the following PDE,

$$\frac{\partial \phi}{\partial t} = \pm |\nabla \phi|, \qquad (5)$$

up to a certain time t. The value of t controls the offset distance from the original surface of $\phi(t = 0)$. A dilation of size α, D_α, corresponds to the solution of Equation 5 at $t = \alpha$ using the positive sign, and likewise erosion, E_α, uses the negative sign. One can now define a morphological opening operator O_α by first applying an erosion followed by a dilation of ϕ, i.e. $\mathrm{O}_\alpha \phi = \mathrm{D}_\alpha \circ \mathrm{E}_\alpha \phi$, which removes small pieces or thin appendages. A closing is defined as $\mathrm{C}_\alpha \phi = \mathrm{E}_\alpha \circ \mathrm{D}_\alpha \phi$, and closes small gaps or holes within objects. Both operations have the qualitative effect of low-pass filtering the isosurfaces in ϕ—an opening by removing material and a closing by adding material. Both operations tend to distort the shapes of the surfaces on which they operate, which is acceptable for the initialization because it will be followed by a surface deformation.

4.1.2 Interactive Initialization

Computational initialization may not always produce a reasonable starting model that deforms into an acceptable final result. Such is the case with the frog-embryo data shown in Figure 4. For volumes that do not allow one to automatically generate an initial model, it is desirable and easier for the user to interactively specify the initial model which is then deformed to fit to the input data. The interactive initialization process has four steps and is presented in Figures 4(b), and 7(a-c). First, the user generates a Marching Cubes mesh from the input volume. This gives some indication of the structures present in the data (Figure 4(b)). The user then creates a Constructive Solid Geometry (CSG) model which defines the shape of the initial surface. The CSG model in blue is interactively positioned relative to the Marching Cubes mesh (Figure 7(a)). The CSG model is scan-converted into a binary volume, with voxels simply marked as inside (1) or outside (0), using standard CSG evaluation techniques [47] within our modeling system [48]. An isosurface of the initialization volume dataset generated from the torus and sphere in Figure 7(a) is presented in Figure 7(b), isovalue = 0.5. This volume dataset is then used as the starting model for the level set deformation stage, which produces the final result seen in Figure 7(c).

4.2 Level Set Surface Deformation

The initialization should position the model near the desired solution while retaining certain properties such as smoothness, connectivity, etc. Given a rough initial estimate, the surface deformation process moves the surface model toward specific features in the data. One must choose those properties of the input data to which the model will be attracted and what role the shape of the model will have in the deformation process. Typically, the deformation process combines a data term with a smoothing term, which prevents the solution from fitting too closely to noise-corrupted data. There are a variety of surface-motion terms that can be used in succession or simultaneously, in a linear combination to form $F(x)$ in Equation 4.

Curvature: This is the smoothing term. For the work presented here we use the mean curvature of the isosurface H to form a vector in the direction of the surface normal n given by

$$F(x) = Hn = \left(\nabla \cdot \frac{\nabla\phi}{|\nabla\phi|} \right) \frac{\nabla\phi}{|\nabla\phi|}. \qquad (6)$$

The mean curvature is also the normal variation of the surface area (i.e., minimal surface area). There are a variety of options for second-order smoothing terms [41], and the question of efficient, effective higher-order smoothing terms is the subject of on-going research [7]. For the work in this paper, we combine mean curvature with one of the following three terms, weighting it by a factor β, which is tuned to each specific application.

Edges: Conventional edge detectors from the image processing literature produce sets of "edge" voxels that are associated with areas of high contrast. For this work we use a gradient magnitude threshold combined with non-maximal suppression, which is a 3D generalization of the method of Canny [16]. The edge operator typically requires a scale parameter and a gradient threshold. For the scale, we use small, Gaussian kernels with standard deviation $\sigma = [0.5 - -1.0]$ voxel units. The threshold depends on the contrast of the volume. The distance transform on this edge map produces a volume that has minima at those edges. The gradient of this volume produces a field that attracts the model to these edges. The edges are limited to voxel resolution because of the mechanism by which they are detected. Although this fitting is not sub-voxel accurate, it has the advantage that it can pull models toward edges from significant distances, and thus inaccurate initial estimates can be brought into close alignment with high-contrast regions, i.e. edges, in the input data. If $\mathscr{E}$ is the set of edges, and $D_{\mathscr{E}}(x)$ is the distance transform to those edges, then the movement of the surface model is given by

$$F(x) = \nabla D_{\mathscr{E}}(x). \qquad (7)$$

Greyscale features—gradient magnitude: Surface models can also be attracted to certain greyscale features in the input data. For instance, the gradient magnitude indicates areas of high contrast in volumes. By following the gradient of such greyscale features, surface models are drawn to minimum or maximum values of that feature. Typically greyscale features, such as the gradient magnitude are computed with a scale operator, e.g., a derivative-of-Gaussian kernel. If models are properly initialized, they can move according to the *gradient of the gradient magnitude* and settle onto the edges of an object at a resolution that is finer than the original volume.

If $G(x)$ is some greyscale feature, for instance $G(x) = |\nabla I(x)|$, where $I(x)$ is the input data (appropriately filtered—we use Gaussian kernels with $\sigma \approx 0.5$), then

$$F(x) = \pm\nabla G(x), \qquad (8)$$

where a positive sign moves surfaces towards maxima and the negative sign towards minima.

Isosurface: Surface models can also expand or contract to conform to isosurfaces in the input data. To a first order approximation, the distance from a point $x \in U$ to the k-level surface of I is given by $(I(x) - k)/|\nabla I|$. If we let $g(\alpha)$ be a fuzzy threshold, e.g., $g(\alpha) = \alpha/\sqrt{1+\alpha^2}$, then

$$F(x) = \frac{\nabla\phi}{|\nabla\phi|} g\left(\frac{I(x) - k}{|\nabla I|} \right) \qquad (9)$$

causes the surfaces of ϕ to expand or contract to match the k isosurface of I. This term combined with curvature or one of the other fitting terms can create "quasi-isosurfaces" that also include other considerations, such as smoothness or edge strength.

5 Results

This section describes how our approach may be used to extract structures from the data described in Section 2. We present surface renderings of the resulting models and detail the specific methods needed to construct each model.

Figure 10 shows 3D renderings of the sequence of steps performed on the ET dendrite data from Figure 1. The first two are the initialization steps, generating a smoothed isosurface and filling gaps with topological and morphological operations. The second two are surface deformation steps, first fitting to discrete edges and then to the gradient magnitude. Figure 8 shows a slice with the boundary of the solution drawn in red, that confirms the accuracy of the results—the red boundary is only an indicator of the solution because it is limited to voxel resolution while the level set model has sub-voxel resolution. This figure also shows the same result for a smoothed isosurface—which is significantly affected by density fluctuations in the data. Figures 9 and 11 show the results of the proposed method compared to the results of a manual segmentation, which took approximately 10 hours of slice-by-slice hand contouring. The manual method suffers from slice-wise artifacts, and, because of the size and complexity of the dataset, the manual segmentation is unable to capture the level of detail that we obtain with the surface-fitting results. Manual segmentation can, however, form connections that are not well supported by the data in order to complete the "spines" that cover this dendrite. These types of

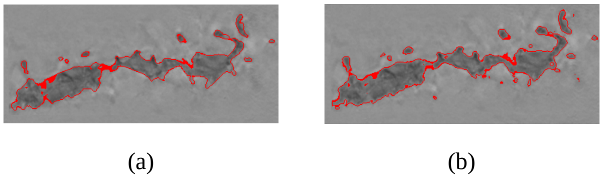
(a) (b)

Figure 8: *[color]* a) Voxel-resolution contours of a dendrite using our level set approach. b) Voxel-resolution contours of the isosurface of the smoothed sampling of the same dendrite.

"judgments" that humans make when they perform such tasks by hand are a mixed blessing. Humans can use high-level knowledge about the problem to fill in where the data is weak, but the expectations of a trained operator can interfere with seeing unexpected or unusual features in the data. Our future work will attempt to incorporate user input to guide the surface-fitting results to obtain a better blend of user expectations and data-driven modeling.

Figure 12 shows the results of fitting a surface model to the MR head data shown in Figure 2. Figure 12(a) is a rendering of the initial model which is the result of smoothing the data, using a flood fill on the exterior to remove isolated holes or bubbles within the head, and treating the model with a closing, $C_{5.0}$. Figure 12(b) shows the results of fitting to the isosurface with a curvature term to ensure smoothness. Some detail is lost around the lips and ears, but overall the fidelity is good and the smoother, simpler surface model suits our application quite well [20].

Figure 13 presents 3D renderings of the sequence of steps performed on the mouse MR data from Figure 3. The first step is the initialization, and the second two are the surface deformation, first fitting to discrete edges and then to the gradient magnitude. This is a significant improvement over the result in Figure 3(b) which suffers from noise and misclassifications. Figure 13(d) presents several other structures that were segmented from the mouse embryo dataset. The skin (grey) and the liver (blue) were isolated using computational initialization. The brain ventricles (red) and the eyes (green) were segmented with interactive initialization.

Figure 14 presents models from four samples of the MR series of the developing frog embryo. The top left image (Hour 9) shows the first evident structure, the blastocoel, in blue, surrounded by the outside casing of the embryo in grey. The top right image (Hour 16) demonstrates the expansion of the blastocoel and the development of the blastoporal lip in red. In the bottom left image (Hour 20) the blastoporal lip has collapsed, the blastocoel has contracted, and the archenteron in green has developed. In the bottom right image (Hour 30) the blastocoel has collapsed and only the archenteron is present. As can been seen from Figure 4(b) that it may be difficult to isolate structures using only their voxel values. We therefore used our interactive techniques to isolate (during initialization) most of the structures in the frog embryo samples.

Table 1 describes for each dataset the specific techniques and parameters we used for the results in this paper. These parameters were obtained by first making a sensible guess based on the contrasts and sizes of features in the data and then using trial and error to obtain acceptable results. Each dataset was processed between 4 and 8 times to achieve these results. More tuning could improve things further, and once these parameters are set, they work moderately well for similar modalities with similar subjects. The method is iterative, but the update times are proportional to the surface area. On an SGI 180MHz MIPS 10000 machine, the smaller mouse MR dataset required approximately 10 minutes of CPU time, and the dendrite dataset ran for approximately 45 minutes. Most of this time was spent in the initialization (which requires several complete passes through the data) and in the edge detection. The frog

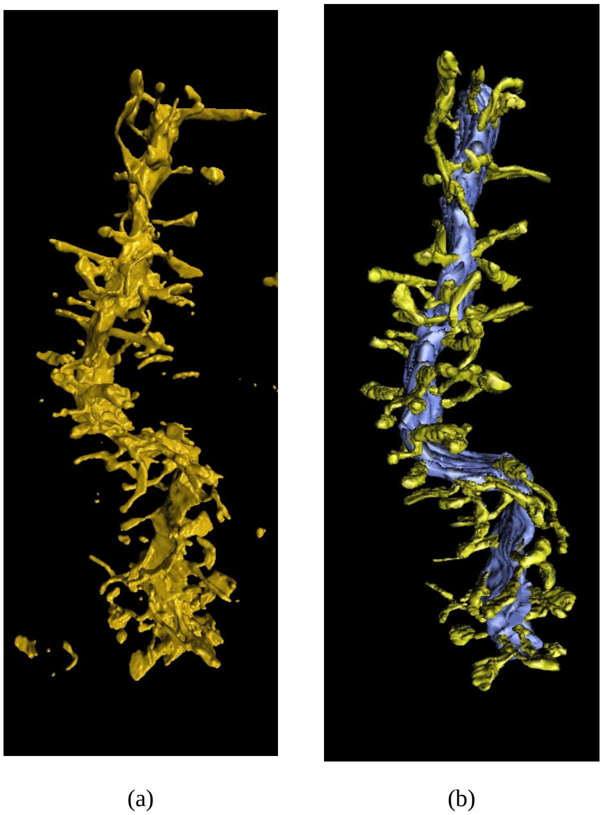
(a) (b)

Figure 9: *[color]* a) Rendering of a dendrite segmented using our the proposed method. b) Rendering of the same dendrite, but this time segmented manually.

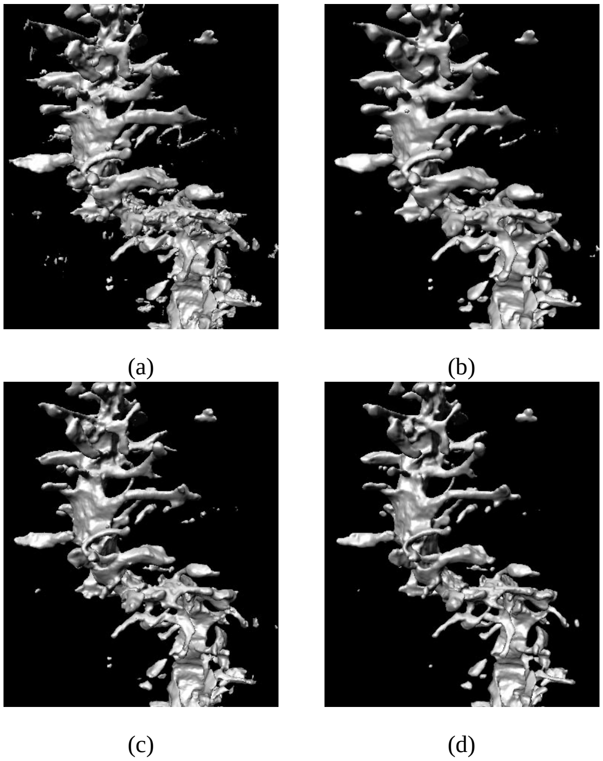

Figure 10: a) The steps in the surface fitting process: An isosurface of smoothed data. b) Morphological operators fill in gaps and remove smaller, disconnected pieces. c) Fitting to edges brings the model closer to high-contrast regions in the data. d) Fitting to maximal gradient magnitude gives more detail.

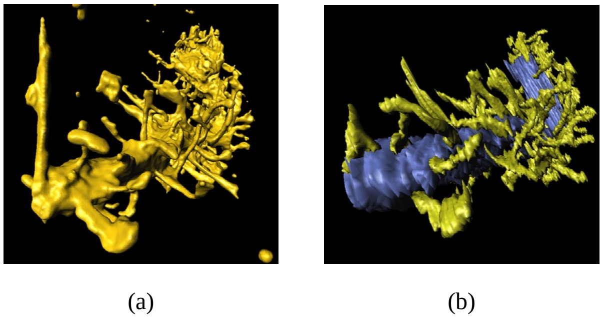

Figure 11: *[color]* a) Close-up view of the final result of the dendrite rendering using our scheme – note the level of details. b) Close-up view of the manual segmentation – note the lack of detail compared to the proposed method.

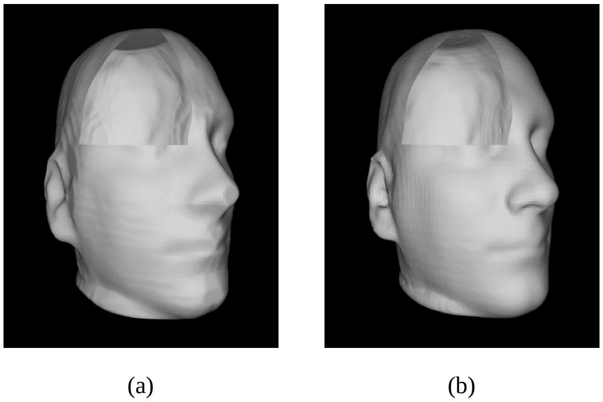

Figure 12: a) This image shows the rough initialization surface used for a level set segmentation of an MR scanned head. A small section of the surface has been removed to show that it does not contain internal structures. b) The final result is smoothed (almost no aliasing from the scanning slices), but with good fidelity.

embryo datasets needed only a few minutes of processing time, because they did not require computational initialization and are significantly smaller than the other example datatsets.

6 Conclusions

This paper describes a system that uses level set surface models in conjunction with a suite of initialization techniques to segment structures in volume data. Level set surface modeling is a technology that allows one to manipulate or deform the isosurfaces of a volume toward interesting features in the input data. Because the technology is volumetric, it provides opportunities to combine voxel-based techniques, such as filtering, classification, and morphology with surface-fitting methods based on deformable models. We have shown that combining level set methods with a variety of initialization techniques produces a powerful framework capable of segmenting many different types of volume datasets. In the case of the ET dendrite data, our approach offers significant advantages in both time and quality over hand-contoured segmentations, which are currently the state-of-the-art.

Currently there are two significant drawbacks of the proposed method. First is the choice of parameters. There are a number of parameters that must be tuned, and their settings affect the final solution. The second drawback is the computation time, which is quite long for large datasets. The second problem aggravates the first, because exploring the parameter space by trial and error is a potentially lengthy process. Future work will focus on increasing the update rates by parallelizing the computation. This is feasible because the numerical methods lend themselves to a spatial decomposition of the model domain. If the updates were sufficiently fast, users can explore the parameter space interactively by turning various knobs and evaluating the quality of the results. This would greatly increase the effectiveness of the method.

Acknowledgements

We would like to thank Dr. Alan Barr for his support of this project, and Dr. Cyrus Papan, Dr. Russ Jacobs, Dr. Eric Ahrens, Dr. Mark Ellisman, Dr. Maryanne Martone and Mr. David Weinstein for their assistance and cooperation. This work was supported

Dataset	Initialization	Surface Fitting
Dendrite	1. Gaussian blur $\sigma = 0.5$ 2. Threshold: $I < 127$ 3. Fill isolated holes 4. Morphology: $\mathrm{O}_{0.5} \circ \mathrm{C}_{1.5}$	1. Edge fitting: $\sigma = 0.75$, threshold = 6, $\beta = 0.1$ 2. Gradient magnitude fitting: $\sigma = 0.5$, $\beta = 1.0$
Head	1. Gaussian blur $\sigma = 1.0$ 2. Threshold: $I > 30$ 3. Fill isolated holes 4. Morphology: $\mathrm{C}_{5.0}$	1. Isosurface fitting: $\beta = 5.0$, $k = 30$.
Mouse	1. Gaussian blur $\sigma = 0.5$ 2. Threshold: $I > 3$, $I < 60$ 3. Fill isolated holes 4. Morphology: $\mathrm{O}_{2.0} \circ \mathrm{C}_{3.0}$	1. Edge fitting: $\sigma = 0.75$, threshold = 20, $\beta = 2$ 2. Gradient magnitude fitting: $\sigma = 0.5$, $\beta = 16.0$
Frog	1. Interactive	1. Gradient magnitude fitting: $\sigma = 1.25$, $\beta = 1.0$

Table 1: Parameters for processing example datasets.

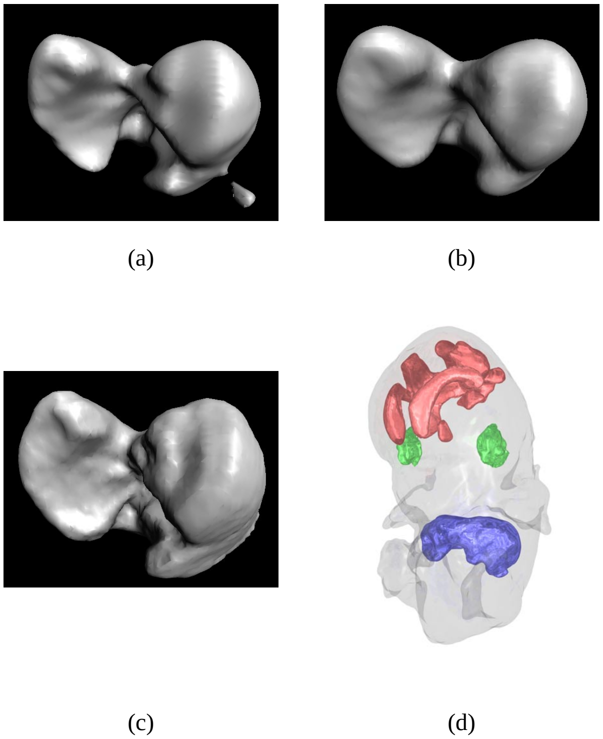

Figure 13: a) The initialization of a mouse liver dataset using morpholog to remove small pieces and holes. b) Surface fitting to discrete edges. c) The final fit to maxima of gradient magnitude. d) *[color]* Final mouse embryo model with skin (grey), liver (blue), brain ventricles (red), and eyes (green).

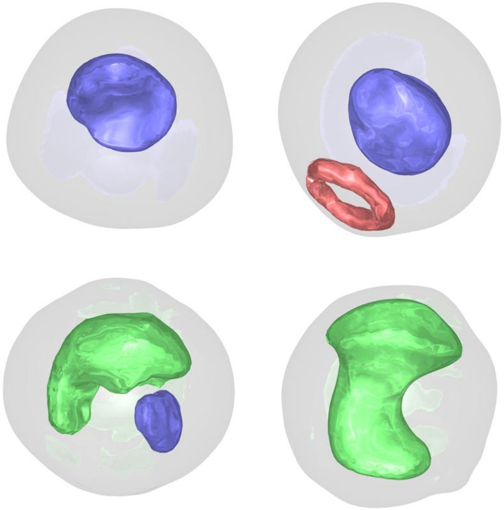

Figure 14: *[color]* Geometric structures extracted from MRI scans of a developing frog embryo, with blastocoel (blue), blastoporal lip (red), and archenteron (green). Hour 9 (top left). Hour 16 (top right). Hour 20 (bottom left). Hour 30 (bottom right).

by National Science Foundation grants #ASC-89-20219, #ACI-99-82273, and #ACI-00-89915, Office of Naval Research grant #N00014-01-10033, the National Institute on Drug Abuse, the National Institute of Mental Health and the NSF, as part of the Human Brain Project, the National Library of Medicine "Insight" Project #N01-LM-0-3503, and the Caltech SURF Program. The datasets are courtesy of the National Center for Microscopy and Imaging Research (funded by NIH grant #P41-RR04050), at UC San Diego (Figure 1), the University of Utah SCI Institute (Figure 2), and the Caltech Biological Imaging Center (Figures 3 and 4).

References

[1] R. A. Drebin, L. Carpenter, and P. Hanrahan, "Volume rendering," in *SIGGRAPH '88 Proceedings*, pp. 65–74, August 1988.

[2] M. Levoy, "Display of surfaces from volume data," *IEEE Computer Graphics and Applications*, vol. 9, no. 3, pp. 245–261, 1990.

[3] D. Laur and P. Hanrahan, "Hierarchical splatting: A progressive refinement algorithm for volume rendering," in *SIGGRAPH '91 Proceedings* (T. W. Sederberg, ed.), pp. 285–288, July 1991.

[4] S. Parker, M. Parker, Y. Livnat, P. Sloan, C. Hansen, and P. Shirley, "Interactive ray tracing for volume visualization," *IEEE Trans. on Viz. and Comp. Graph.*, vol. 5, no. 3, pp. 238–250, 1999.

[5] M. Leventon, O. Faugeraus, W. Grimson, and W. Wells III, "Level set based segmentation with intensity and curvature priors," in *Workshop on Mathematical Methods in Biomedical Image Analysis Proceedings*, pp. 4–11, June 2000.

[6] R. Malladi, J. A. Sethian, and B. C. Vemuri, "Shape modeling with front propagation: A level set approach," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 17, no. 2, pp. 158–175, 1995.

[7] J. Sethian, *Level Set Methods and Fast Marching Methods*. Cambridge: Cambridge University Press, second ed., 1999.

[8] L. Staib, X. Zeng, R. Schultz, and J. Duncan, "Shape constraints in deformable models," in *Handbook of Medical Imaging* (I. Bankman, ed.), ch. 9, pp. 147–157, Academic Press, 2000.

[9] Z. Wu, H.-W. Chung, and F. W. Wehrli, "A Bayesian approach to subvoxel tissue classification in NMR microscopic images of trabecular bone," *Journal of Computer Assisted Tomography*, vol. 12, no. 1, pp. 1–9, 1988.

[10] Y.-H. Kao, J. A. Sorenson, and S. S. Winkler, "MR image segmentation using vector decompsition and probability techniques: A general model and its application to dual-echo images," *Magnetic Resonance in Medicine*, vol. 35, pp. 114–125, 1996.

[11] H. E. Cline, W. E. Lorensen, R. Kikinis, and F. Jolesz, "Three-dimensional segmentation of MR images of the head using probability and connectivity," *Journal of Computer Assisted Tomography*, vol. 14, pp. 1037–1045, Nov., Dec. 1990.

[12] D. H. Laidlaw, K. W. Fleischer, and A. H. Barr, "Partial-volume Bayesian classification of material mixtures in MR volume data using voxel histograms," *IEEE Transactions on Medical Imaging*, vol. 17, pp. 74–86, feb 1998.

[13] V. E. Johnson, "A framework for incorporating structural prior information into the estimation of medical images," in *Information Processing in Medical Imaging (IPMI'93)* (H. H. Barrett and A. F. Gmitro, eds.), no. 687 in Lecture Notes in Computer Science, pp. 307–321, Springer-Verlag, 1993.

[14] D. Marr and E. Hildreth, "Theory of edge detection," *Proceedings of the Royal Society of London*, vol. B, no. 207, pp. 187–217, 1980.

[15] D. Marr, *Vision*. San Francisco: Freeman, 1982.

[16] J. Canny, "A computational approach to edge detection," *IEEE Trans. on Pat. Anal. and Mach. Intel.*, vol. 8, no. 6, pp. 679–698, 1986.

[17] T. Cootes, A. Hill, C. Taylor, and J. Haslam, "The use of active shape models for locating structures in medical images," in *Information Processing in Medical Imaging (IPMI'93)* (H. H. Barrett and A. F. Gmitro, eds.), no. 687 in Lecture Notes in Computer Science, pp. 33–47, Springer-Verlag, 1993.

[18] G. Stetten and S. Pizer, "Medial node models to identify and measure objects in real-time 3d echocardiography," *IEEE Transactions on Medical Imaging*, vol. 18, no. 10, pp. 1025–1034, 1999.

[19] W. Lorenson and H. Cline, "Marching Cubes: A high resolution 3D surface construction algorithm," *Computer Graphics*, vol. 21, no. 4, pp. 163–169, 1982.

[20] Z. Wood, M.Desbrun, P. Schröder, and D. Breen, "Semi-regular mesh extraction from volumes," in *Proceedings of Visualization 2000*, pp. 275–282, 2000.

[21] J. Miller, D. Breen, W. Lorensen, R. O'Bara, and M. Wozny, "Geometrically deformed models: A method for extracting closed geometric models from volume data," in *SIGGRAPH '91 Proceedings*, pp. 217–226, July 1991.

[22] A. P. Pentland, "Perceptual organization and the representation of natural form," *Artificial Intelligence*, vol. 28, pp. 293–331, 1986.

[23] D. Terzopoulos and D. Metaxas, "Dynamic 3D models with local and global deformations: Deformable superquadrics," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 13, no. 7, pp. 703–714, 1991.

[24] A. Gupta and R. Bajcsy, "Volumetric segmentation of range images of 3D objects using superquadric models," *CVGIP: Image Understanding*, vol. 58, no. 3, pp. 302–326, 1993.

[25] S. Muraki, "Volumetric shape description of range data using "blobby model"," in *SIGGRAPH '91 Proceedings* (T. W. Sederberg, ed.), pp. 227–235, July 1991.

[26] R. Szeliski, D. Tonnesen, and D. Terzopoulos, "Modeling surfaces of arbitrary topology with dynamic particles," in *Proc. Fourth Int. Conf. on Comp. Vision (ICCV'93)*, (Berlin, Germany), pp. 82–87, IEEE Computer Society Press, May 1993.

[27] T. McInerney and D. Terzopoulos, "A dynamic finite element surface model for segmentation and tracking in multidimensional medical images with application to cardiac 4d image analysis," *Computerized Medical Imaging and Graphics*, vol. 19, no. 1, pp. 69–83, 1995.

[28] J. Park, D. Metaxas, A. A. Young, and L. Axel, "Deformable models with parameter functions for cardiac motion analysis from tagged MRI data," *IEEE Transactions on Medical Imaging*, vol. 15, pp. 278–289, June 1996.

[29] D. DeCarlo and D. Metaxas, "Shape evolution with structural and topological changes using blending," *IEEE Trans. on Pat. Anal. and Mach. Intel.*, vol. 20, pp. 1186–1205, November 1998.

[30] R. Ramamoorthi and J. Arvo, "Creating generative models from range images," in *SIGGRAPH '99 Proceedings*, pp. 195–204, August 1999.

[31] S. Osher and J. Sethian, "Fronts propagating with curvature-dependent speed: Algorithms based on Hamilton-Jacobi formulations," *Journal of Computational Physics*, vol. 79, pp. 12–49, 1988.

[32] J. Sethian, "A fast marching level set method for monotonically advancing fronts," in *Proceedings of the National Academy of Science*, vol. 93 of 4, pp. 1591–1595, 1996.

[33] J. Tsitsiklis, "Efficient algorithms for globally optimal trajectories," *IEEE Transactions on Automatic Control*, vol. 40, no. 9, pp. 1528–1538, 1995.

[34] R. Fedkiw, R. Aslam, R. Merriman, and S. Osher, "A non-oscillatory eulerian approach to interfaces in multimaterial flows (the ghost fluid method)," *Journal of Computational Physics*, vol. 152, pp. 457–492, 1999.

[35] L. Alvarez and J.-M. Morel, "A morphological approach to multiscale analysis: From principles to equations," in *Geometry-Driven Diffusion in Computer Vision* (B. M. ter Haar Romeny, ed.), pp. 4–21, Kluwer Academic Publishers, 1994.

[36] V. Caselles, R. Kimmel, and G. Sapiro, "Geodesic active contours," in *Fifth Int. Conf. on Comp. Vision*, pp. 694–699, IEEE, IEEE Computer Society Press, 1995.

[37] B. B. Kimia and S. W. Zucker, "Exploring the shape manifold: the role of conservation laws.," in *Shape in Picture: the mathematical description of shape in greylevel images* (Y.-L. O, A. Toet, H. Heijmans, D. H. Foster, and P. Meer, eds.), Springer-Verlag, 1992.

[38] R. T. Whitaker and D. T. Chen, "Embedded active surfaces for volume visualization," in *SPIE Medical Imaging 1994*, (Newport Beach, California), 1994.

[39] R. Whitaker and D. Breen, "Level-set models for the deformation of solid objects," in *The Third International Workshop on Implicit Surfaces*, pp. 19–35, Eurographics, 1998.

[40] R. T. Whitaker, "A level-set approach to 3D reconstruction from range data," *Int. Jrnl. of Comp. Vision*, vol. October, no. 3, pp. 203–231, 1998.

[41] R. T. Whitaker, "Volumetric deformable models: Active blobs," in *Visualization In Biomedical Computing 1994* (R. A. Robb, ed.), (Mayo Clinic, Rochester, Minnesota), pp. 122–134, SPIE, 1994.

[42] R. T. Whitaker, "Algorithms for implicit deformable models," in *Fifth Intern. Conf. on Comp. Vision*, IEEE, IEEE Computer Society Press, 1995.

[43] A. Yezzi, S. Kichenassamy, A. Kumar, P. Olver, and A. Tannenbaum, "A geometric snake model for segmentation of medical imagery," *IEEE Transactions on Medical Imaging*, vol. 16, pp. 199–209, April 1997.

[44] L. Lorigo, O. Faugeraus, W. Grimson, R. Keriven, and R. Kikinis, "Segmentation of bone in clinical knee MRI using texture-based geodesic active contours," in *Medical Image Computing and Computer-Assisted Intervention (MICCAI '98)* (W. Wells, A. Colchester, and S. Delp, eds.), pp. 1195–1204, October 1998.

[45] R. van den Boomgaard and A. W. M. Smeulders, "The morphological structure of images, the differential equations of morphological scale-space," *IEEE Trans. on Pat. Anal. and Mach. Intel.*, vol. 16, no. 11, pp. 1101–1113, 1994.

[46] P. Maragos, "Differential morphology and image processing," *IEEE Transactions on Image Processing*, vol. 5, pp. 922–937, June 1996.

[47] A. Requicha and H. Voelcker, "Boolean operations in solid modeling: Boundary evaluation and merging algorithms," *Proceedings of the IEEE*, vol. 73, no. 1, pp. 30–44, 1985.

[48] P. Getto and D. Breen, "An object-oriented architecture for a computer animation system," *The Visual Computer*, vol. 6, pp. 79–92, March 1990.