Published in the *Proc. CVRMed'97*, Grenoble, France, March, 1997.

Medical Image Segmentation Using Topologically Adaptable Surfaces

Tim McInerney and Demetri Terzopoulos

Dept. of Computer Science, University of Toronto, Toronto, ON, Canada

Abstract. Efficient and powerful topologically adaptable deformable surfaces can be created by embedding and defining discrete deformable surface models in terms of an *Affine Cell Decomposition* (ACD) framework. The ACD framework, combined with a novel and original reparameterization algorithm, creates a simple but elegant mechanism for multiresolution deformable curve, surface, and solid models to "flow" or "grow" into objects with complex geometries and topologies, and adapt their shape to recover the object boundaries. ACD-based models maintain the traditional parametric physics-based formulation of deformable models, allowing them to incorporate *a priori* knowledge in the form of energy and force-based constraints, and provide intuitive interactive capabilities. This paper describes ACD-based deformable surfaces and demonstrates their potential for extracting and reconstructing some of the most complex biological structures from medical image volumes.

1 Introduction

Segmenting and reconstructing structures from medical image volumes into a more manageable analytic form is hindered by the sheer size of the data sets and the complexity and variability of the anatomic shapes of interest. The shortcomings typical of sampled data, such as sampling artifacts, spatial aliasing, and noise, often cause the boundaries of structures to be indistinct and disconnected. The challenge is to extract boundary elements belonging to the same structure and integrate these elements into a coherent and consistent model of the structure. In recent years, the use of deformable surface models [14, 11, 2, 9] has been proposed to meet this challenge.

Deformable models offer an attractive approach to the medical image segmentation problem because they combine many desirable features such as compact and analytical representations of object shape, inherent connectivity and smoothness that compensates for noise and boundary irregularities, the ability to incorporate anatomic knowledge, and the support of intuitive interaction mechanisms. Since the initial development of this successful boundary integration and feature extraction technique, many useful variants of these models [5, 3, 13, 15, 6] have been developed in an attempt to overcome their most significant limitations. In particular, deformable models must usually be placed close to the boundaries of the target object to guarantee good performance. The internal energy constraints of deformable models can limit their geometric flexibility and prevent

them from representing long tube-like shapes or shapes with significant protrusions. Furthermore, the topology of the object of interest must be known in advance since traditional deformable models are parametric and are incapable of topological transformations without additional machinery. Finally, the energy function of most deformable models depends on its parameterization and is not directly related to the geometry of the object. New parameterizations can change the model energy in arbitrary ways.

Our goal is the development of a unified, complementary framework that overcomes the limitations of classical deformable models and increases their degree of automation, while retaining all of their traditional strengths. In [10], we introduced a new class of 2D deformable models known as *topologically adaptable snakes (T-snakes)*. These models exploit an *Affine Cell Decomposition* (ACD) of the image domain, creating a theoretically sound framework that significantly extends the abilities of classical deformable models. Embedding deformable models in an ACD framework allows the models to extract and reconstruct even the most complex biological structures. The ACD framework, combined with a novel reparameterization algorithm, creates a simple but elegant and powerful mechanism for multiresolution deformable curve, surface, and solid models to "flow" or "grow" into objects with complex geometries and topologies, adapting their shape to conform to the object boundaries. The ACD framework enables the models to maintain the traditional properties associated with classical deformable models, such as user interaction and the incorporation of constraints through energy functions or force functions, while overcoming many of their limitations. The framework also provides a convenient mechanism for the incorporation of "hard" geometric, topological and global shape constraints [8]. In this paper, we will review the 3D ACD framework and describe the formulation of topologically adaptable deformable surfaces (*T-surfaces*). We will then provide several segmentation examples using these models in order to demonstrate their potential for efficient, accurate and reproducible extraction and analysis of anatomic structures from medical image volumes.

2 ACD-based Deformable Models

ACD-based deformable models, T-snakes and T-surfaces, are parametric models that have the power of an implicit formulation by using a superposed affine cell grid to quickly and efficiently reparameterize the models during their evolution. In particular, a T-surface is defined as an elastic triangular mesh. This closed surface model is a discrete approximation to traditional multidimensional snakes models and retains many of the snakes properties. As the T-surface moves under the influence of external and internal forces, we reparameterize it with a new set of nodes and triangles by efficiently computing the intersection points of the model with the superposed grid. We also keep track of the interior region of the model by "turning on" any grid vertices the T-surface passes over during its motion. By reparameterizing the T-surface at each iteration of the evolutionary process, we create a simple, elegant and automatic model subdivision technique and the grid provides a framework for robust topological transformations. This

allows the model to be relatively independent of its initial placement and "flow" into complex shapes with complex topologies in a stable manner. Conversion to a traditional parametric deformable surface representation is simply a matter of disabling the grid at any time during the evolutionary process. By providing a boundary representation as well as a representation of the interior region of an object, this hybrid deformable model combines the space partitioning, intrinsic parameterization and topological flexibility properties of an implicit formulation, such as those described in [1, 7], with the boundary properties of a parametric model. The advantage of retaining a parametric model formulation is that any "soft" constraint expressed as an energy function or force function can be easily incorporated into the physics-based framework.

2.1 Model Description

We define a T-surface as a closed oriented triangular mesh. The vertices of these triangles, the model nodes, act as a dynamic particle system where the particles are interconnected by discrete spring units. That is, we associate with each node i a time varying position $\mathbf{x}_i(t) = [x_i(t), y_i(t), z_i(t)]$, along with internal "tension" forces $\boldsymbol{\alpha}_i(t)$ and "rigidity" forces $\boldsymbol{\beta}_i(t)$, external inflation forces $\boldsymbol{\rho}_i(t)$, and external image edge forces $\mathbf{f}_i(t)$ [8]. The behavior of the dynamic node-spring system is governed by a simplified form of the discrete Lagrange equations of motion

$$\gamma_i \dot{\mathbf{x}}_i + \boldsymbol{\alpha}_i + \boldsymbol{\beta}_i = \boldsymbol{\rho}_i + \mathbf{f}_i, \tag{1}$$

where $\dot{\mathbf{x}}_i$ is the velocity of node i, γ_i is a damping coefficient that controls the rate of dissipation of the kinetic energy of the nodes, and $\mathbf{f}_i$ are the external image forces (typically scaled by a weight p) attracting the model towards image edges. To drive the model towards object boundaries we use an external "inflation" force, $\boldsymbol{\rho}_i = q\mathbf{n}_i$, where q is the force amplitude or weight and $\mathbf{n}_i$ is the unit normal at node i. The weight q is set to a negative value at node i if the image intensity falls below a threshold value at this point. The strengths of the internal forces, $\boldsymbol{\alpha}_i(t)$ and $\boldsymbol{\beta}_i(t)$, are controlled by weights a_i and b_i, respectively. We integrate the first-order dynamic system (1) forward through time using an explicit Euler method.

2.2 Simplicial Cell Decomposition

The grid of discrete cells used to reparameterize our deformable surface is an example of affine cell decomposition using simplicial cells. In a simplicial decomposition, space is partitioned into cells defined by open simplices, where an n-simplex is the simplest geometrical object of dimension n: a triangle in 2D and a tetrahedron in 3D. The simplest triangulation of Euclidean space $\mathcal{R}^n$ is the Coxeter-Freudenthal triangulation. It is constructed by dividing space using a uniform cubic grid and the triangulation is obtained by subdividing each cube into $n!$ simplices. In $\mathcal{R}^3$, one such subdivision produces 6 tetrahedra per cubic cell (Fig. 1).

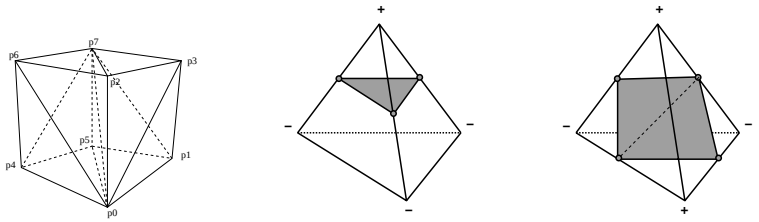

(a) (b) (c)

Fig. 1. (a) Cubic cell division into 6 tetrahedra: $\rho_0 = (p_0, p_1, p_3, p_7)$, $\rho_1 = (p_0, p_1, p_5, p_7)$, $\rho_2 = (p_0, p_2, p_3, p_7)$, $\rho_3 = (p_0, p_2, p_6, p_7)$, $\rho_4 = (p_0, p_4, p_5, p_7)$, $\rho_5 = (p_0, p_4, p_6, p_7)$. (b) Intersection of object boundary with grid tetrahedra.

Simplicial decompositions provide an unambiguous framework for the creation of local polygonal approximations of a closed surface. A simplex can be classified in relation to the partitioning of space by a closed orientable surface by testing the "sign" of its vertices. If the signs are the same for all vertices, the simplex must be totally inside or outside the surface. If the signs are different, the surface must intersect the simplex (Fig. 1b,c). In an n-simplex, the negative vertices can always be separated from the positive vertices by a single plane; thus an unambiguous polygonalization of a simplex always exists. Furthermore, as long as neighboring cubical cells containing the simplices are decomposed so that they share common faces, a consistent polygonization of the entire surface will result.

The set of grid tetrahedra that intersect the surface (the boundary tetrahedra) form a three dimensional combinatorial manifold that has as its dual a two dimensional manifold that approximates the surface. The two dimensional manifold is constructed from the intersection of the true surface with the edges of each boundary tetrahedron. The intersection points result in one triangle or one quadrilateral (which can be subdivided into two triangles) approximating the surface inside each boundary tetrahedron (Fig. 1b,c), where each triangle or quadrilateral intersects a tetrahedron on three or four distinct edges, respectively. The triangle (or quadrilateral) separates the positive vertices from the negative vertices of the tetrahedron. The set of all these triangles constitute the combinatorial manifold that approximates the original surface.

2.3 Iterative Reparameterization

The time derivatives in (1) are approximated by finite differences. During k time steps of the numerical time integration (one deformation step), the model moves (expands or shrinks, but not both) from its current position to a new position. At the beginning of the deformation step, the model nodes are defined in terms of the edges of the grid boundary tetrahedra. At the end of the deformation step, the nodes have moved "off" of the tetrahedra edges (Fig. 2). We then reestablish the correspondence of the model with the grid by performing a local search and intersection test. A bounding box is formed around each model triangle at its new position and grid intersection points are computed and used as the new set of model nodes.

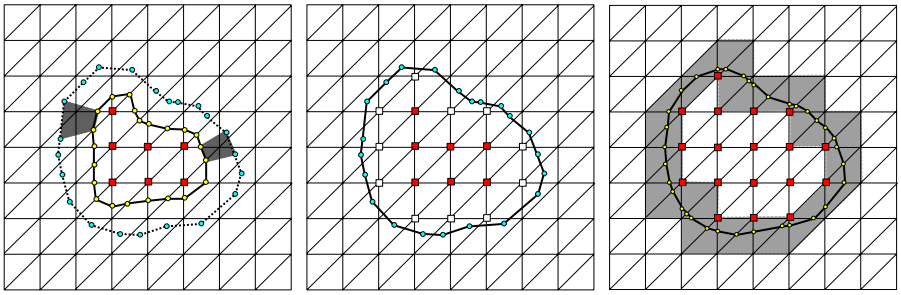

Fig. 2. Illustration of reparameterization process using 2D contour. (a) Shaded regions show examples of grid vertices that are "burned" by expanding contour, (b) new burned grid vertices (white) added to current burned vertices (dark), (c) new contour after one deformation step showing new grid intersections, burned grid vertices, and boundary grid simplices (gray shaded).

During the evolution of a T-surface, it "passes over" a set of grid vertices (Fig. 2). Using a flame propagation analogy, we say that these fixed grid points have been "burned". By identifying and recording these burned grid vertices during a deformation step, we are able to determine and track the interior region of the T-surface. Furthermore, throughout the evolution of the model the interior grid vertices act as an implicit function to unambiguously define the boundary of the model by partitioning the image domain into an inside region and an outside region. Determining the set of grid vertices that have been burned during a deformation step uses a simple and robust classification algorithm [8]. Each T-surface triangle may have passed over zero, one, or several grid vertices during a deformation step. For each triangle, we form a bounding box around its current position and new position, enabling us to quickly determine the subset of grid vertices that may have been burned. For each of the grid vertices in this subset, we partition the image domain into eight subspaces by forming three intersecting planes. The three planes are formed by joining the two nodes of each model triangle edge at their current position to the grid vertex in question. We then classify the three nodes of the model triangle at its new position into one of the eight subspaces. An efficient classification algorithm is then used to quickly determine whether the grid vertex in question has been burned.

2.4 Topological Transformations

When a T-surface collides with itself or with another T-surface, or splits into two or more parts, a topological transformation must take place. In order to effect consistent topological changes, consistent decisions must be made about disconnecting and reconnecting model nodes. The simplicial grid and the reparameterization process automatically and unambiguously performs these reconnections. By tracking the interior grid vertices (and hence the boundary grid tetrahedra) and reestablishing the correspondence of the surface model with the grid after every deformation step, we can always unambiguously determine the boundary

or "isosurface" of the new T-surface(s). We simply compute new T-surface triangles from the signs of the grid vertices in each boundary tetrahedron, and from the intersection points computed previously, such that the inside and outside grid vertices of these tetrahedra are separated by a single plane. Thus, by mimicing the evolving level set of an implicit function, the simplicial grid and the reparameterization process guarantees that topological transformations are handled automatically, consistently and efficiently.

2.5 The Final Algorithm

In contrast to the level set evolution techniques which accede control to a higher dimensional implicit function, we retain an explicit parametric model formulation. A parametric formulation allows us to track and control the evolution of the model. Consequently, reparameterizations can be performed very efficiently and constraints can be easily imposed on the model. The complete algorithm is as follows:

For each k time steps:

1. For each time step, compute all forces, integrate the equations of motion and update node positions.
2. Compute new grid intersection points for all model triangles.
3. For each model triangle, compute "burned" grid vertices and and determine new boundary grid tetrahedra and new model triangles.
4. From the set of current model triangles, identify valid triangles and discard invalid triangles. A model triangle is valid if its corresponding grid tetrahedron is still a boundary tetrahedron.

A T-surface is considered to have reached its equilibrium state when all of the model triangles have been inactive for a specified number of deformation steps. Model triangle activity or movement is measured via the grid again using a flame propagation analogy. Model triangles are assigned a "temperature" based on the number of deformation steps the triangle (and its corresponding boundary grid tetrahedron) has remained valid. A model triangle is considered inactive when its temperature falls below a preset "freezing point". Frozen model triangles can be removed from the computation. This adjustable mechanism allows the system to maintain a manageable computational burden for many segmentation scenarios. Once a T-surface has reached equilibrium the grid can be disabled, if desired, and the model run as a classical deformable surface.

3 Volume Image Segmentation with T-Surfaces

We have implemented a prototype system using T-surfaces. We have used the system to segment and reconstruct a wide range of anatomic structures with complex shapes and topologies. We initialize a T-surface using an implicitly defined superquadric function which can be quickly scaled, bent and tapered, if desired, and placed inside the object of interest (or scaled such that it contains the object). The superquadric is then converted into a T-surface and the model

will then expand (or shrink if it surrounds the object) and adapt to the object boundaries. Multiple T-surfaces can be dynamically created and placed throughout the object. They will then adapt and automatically merge to form one model of the object.

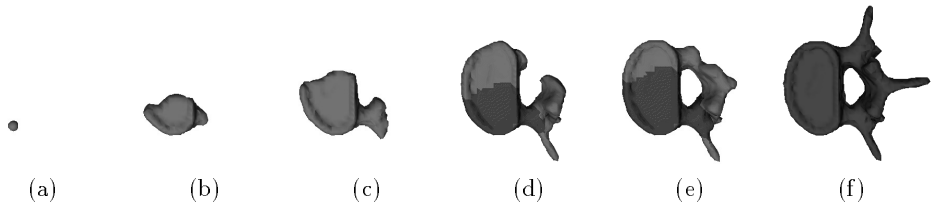

Fig. 3. T-surface segmenting vertebra phantom from CT image volume. The dark shaded regions are frozen and have been removed from the computation.

In the first example, we apply a T-surface to a $120 \times 128 \times 52$ CT image volume of a human vertebra phantom to demonstrate the topological adaptability of the T-surface (Figure 3). We use a $32 \times 30 \times 13$ cell grid (where each cubical cell is divided into 6 tetrahedra) with model parameters: $p = 20.0$, $q = 80.0$, $a = 20.0$, $b = 60.0$, and time step $\Delta t = 0.001$.

In the second example we segment and reconstruct the left-ventricular (LV) chamber and aorta from a CT image volume of a canine heart. The image volume dimensions are $128 \times 128 \times 118$ and a $20 \times 20 \times 20$ resolution grid was used (model parameters: $p = 61.0$, $q = 60.0$, $a = 20.0$, $b = 40.0$, $\Delta t = 0.002$).

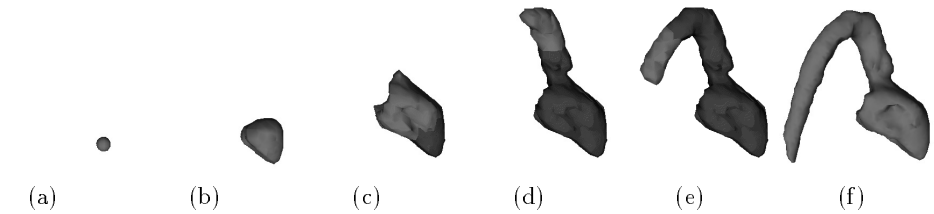

Fig. 4. T-surface segmenting LV and aorta from CT image volume.

We manually seed the LV with a small T-surface and the segmentation then proceeds automatically (Figure 4). The inflation force is weighted with LV region image intensity statistics to reinforce the image edge forces. The final result, after disabling the grid and converting to a classical deformable surface, is shown in Figure 4(f). While there are a few image slices that may require manual editing, most of the model fits very accurately and the entire process takes under 4 minutes on an SGI Indigo 2 workstation. Segmentation with T-surfaces is highly reproducible; as long as the T-surface is seeded within the bright region of the LV, it produces almost identical results.

In the third example, we have used a T-surface to segment and reconstruct the vascular system of the brain from an MRA image volume. The data con-

sist of a stack of 100 slices each with 512^2 pixels, 2 bytes per pixel (voxel size $0.4296875mm \times 0.4296875mm \times 0.7mm$). The image volume was interpolated to produce cubical voxels resulting in 162 slices. A T-surface is seeded at the root of the vessel tree and then flows into the vessels, automatically extracting the vascular system (Figure 5). A voxel resolution grid was used with model parameters: $p = 00.0$, $q = 80.0$, $a = 20.0$, $b = 60.0$, $\Delta t = 0.0005$. The segmentation takes about one hour on an SGI Indigo 2 workstation, primarily due to the disk thrashing resulting from the massive data set and limited memory of the workstation. When combined with image statistics-weighted inflation forces, a T-surface behaves as an active region growing model that is able to integrate edge information and filter out noise through the model smoothness constraints. Although MR angiography produces high contrast images, they can have a large

Fig. 5. T-surface segmentation of cerebral vasculature from MRA image volume.

deviation in the gray scale range. This large range makes it difficult to segment the images with simple thresholding techniques.

In this final example, we use a T-surface with a image statistics-weighted inflation force to segment the cerebral cortex from a preprocessed MR image volume (the skull has been manually removed from the image) (Figure 6). A $64 \times 64 \times 34$ resolution grid was used for the $256 \times 256 \times 136$ voxel image volume (model parameters: $p = 0.0$, $q = 48.0$, $a = 20.0$, $b = 40.0$, $\Delta t = 0.0008$). The T-surface was initialized to surround the cortex and then shrinks and conforms to the CSF/gray matter interface. The automatically subdividing T-surface is able to penetrate, with good success, into the narrow and deep cavities of the highly convoluted cortex. A more accurate result could be obtained with a finer resolution grid at the expense of increased computation time and a larger number of model triangles.

4 Discussion

A potential problem with the fluid-like T-surface is that the internal smoothness constraints and image edge forces may not be sufficient to prevent the model from "leaking" out of gaps or weakly defined edges in object boundaries. There

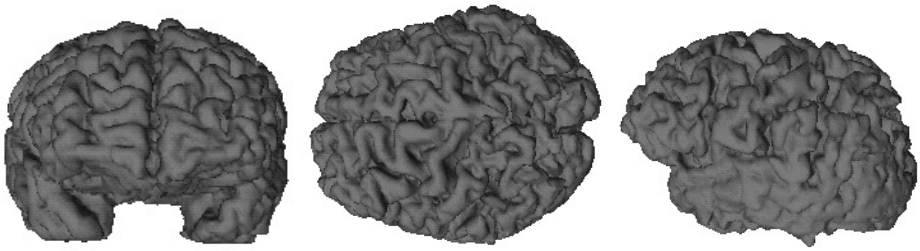

(a) (b) (c)

Fig. 6. T-surface segmentation of the cerebral cortex from MR image volume. (a) front view, (b) top view, (c) side view.

are several possible solutions to this problem. One solution that is used in several of the examples is to weight the inflation force with image intensity statistics of the object. This means that the T-surface will expand at a node only if the local image intensity is within some preset range of the mean gray value of the object. Global or region image statistics can be used. Another solution is to use a Chamfer distance map, which provides the distance to the nearest edge, to again weight the inflation force. Other gradient vector fields derived from image edge maps could also be used [12]. Finally, since ACD-based deformable models are naturally multi-resolution models, this feature can be used with multi-scale image preprocessing techniques [4, 14] to improve the robustness of the segmentation.

5 Conclusion

Deformable models offer a flexible and powerful approach to medical image analysis. Nevertheless, classical deformable surface models suffer from several limitations that prevent their application to the full range of medical image analysis problems and inhibits their potential degree of automation. We have proposed ACD-based deformable surfaces as a solution to many of these limitations. This new class of deformable models can be used to segment, reconstruct, and analyze complex anatomic structures from massive and noisy medical image volumes in a highly automated manner. When combined with various constraint mechanisms, the models have the capability of becoming completely automatic image analysis tools. Finally, the evolution of our discrete deformable model is implemented using a naturally parallel geometric algorithm, and we expect significant performance gains on parallel machines.

Acknowledgements

The MR brain image volume was provided courtesy of Christos Davatzikos of Johns Hopkins School of Medicine. The MRA image volume was provided courtesy of Allen R. Sanderson, Henry Bushwell, and Dennis Parker of the University of Utah, Dept. of Radiology and Dept. of Comp. Science. TM is grateful for the financial support of an NSERC postgraduate scholarship. DT is a Fellow of the Canadian Institute for Advanced Research. This work was made possible by the financial support of the Information Technologies Research Center of Ontario.

References

1. V. Caselles, R. Kimmel, and G. Sapiro. Geodesic active contours. In *Proc. Fifth International Conf. on Computer Vision (ICCV'95), Cambridge, MA, June, 1995*, pages 694–699, Los Alamitos, CA, 1995. IEEE Computer Society Press.
2. I. Cohen, L.D. Cohen, and N. Ayache. Using deformable surfaces to segment 3D images and infer differential structures. *CVGIP: Image Understanding*, 56(2):242–263, 1992.
3. H. Delingette, M. Hebert, and K. Ikeuchi. Shape representation and image segmentation using deformable surfaces. *Image and Vision Computing*, 10(3):132–144, April 1992.
4. M. Kass, A. Witkin, and D. Terzopoulos. Snakes: Active contour models. *International Journal of Computer Vision*, 1(4):321–331, 1988.
5. F. Leitner and P. Cinquin. Complex topology 3D objects segmentation. In *Model-Based Vision Development and Tools*, volume 1609 of *SPIE Proc.*, pages 16–26, Bellingham, WA, 1991. SPIE.
6. R. Malladi, R. Kimmel, D. Adalsteinsson, G. Sapiro, V. Caselles, and J.A. Sethian. A geometric approach to segmentation and analysis of 3D medical images. In *IEEE Workshop on Mathematical Methods in Biomedical Image Analysis, San Francisco, CA, June, 1996*, pages 244–252, Los Alamitos, CA, 1996. IEEE Computer Society Press.
7. R. Malladi, J. Sethian, and B.C. Vemuri. Shape modeling with front propagation: A level set approach. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 17(2):158–175, Feb. 1995.
8. T. McInerney. *Topologically Adaptable Deformable Models for Medcial Image Analysis*. PhD thesis, Department of Computer Science, University of Toronto, Toronto, ON, Canada, 1997. To be completed Jan., 1997.
9. 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*, 19(1):69–83, January 1995.
10. T. McInerney and D. Terzopoulos. Topologically adaptable snakes. In *Proc. Fifth International Conf. on Computer Vision (ICCV'95), Cambridge, MA, June, 1995*, pages 840–845, Los Alamitos, CA, 1995. IEEE Computer Society Press.
11. J.V. Miller, D.E. Breen, W.E. Lorensen, R.M. O'Bara, and M.J. Wozny. Geometrically deformed models: A method for extracting closed geometric models from volume data. In *Computer Graphics (Proc. SIGGRAPH'91 Conf., Las Vegas, NV, July, 1991)*, volume 25(4), pages 217–226, July 1991.
12. J.L. Prince and C. Xu. A new external force model for snakes. In *1996 Image and Multidimensional Signal Processing Workshop*, pages 30–31, 1996.
13. R. Szeliski, D. Tonnesen, and D. Terzopoulos. Modeling surfaces of arbitrary topology with dynamic particles. In *Proc. Conf. Computer Vision and Pattern Recognition (CVPR'93), New York, NY, June, 1993*, pages 82–87, Los Alamitos, CA, 1993. IEEE Computer Society Press.
14. D. Terzopoulos, A. Witkin, and M. Kass. Constraints on deformable models: Recovering 3D shape and nonrigid motion. *Artificial Intelligence*, 36(1):91–123, 1988.
15. R. Whitaker. Volumetric deformable models. In R.A. Robb, editor, *Proc. Third Conf. on Visualization in Biomedical Computing (VBC'94), Rochester, MN, October, 1994*, volume 2359 of *SPIE Proc.*, Bellingham, WA, 1994. SPIE.