

A New Deformable Model for 3D Image Segmentation

Zixin Zhang, Michael Braun and Phillip Abbott

Department of Applied Physics, University of Technology, Sydney,
P.O.Box 123, Broadway, NSW 2007, Australia

Abstract. A fully 3D active surface model is presented with self-inflation and self-deflation forces. The model makes full use of 3D image information, deforms locally and allows strong deformation. The self-inflation and self-deflation forces enable the active surface to travel a long distance without the help from any external forces. We introduce a method of adapting model parameters, which enables our model to bypass some noise and irrelevant edge points. The model is tested with synthetic and real images. Accurate segmentation results are obtained in the presence of image noise and imperfect image data. Importantly, the model is capable of converging to the correct boundary even if the initial estimate is not close. Computational efficiency of segmentation with our model is addressed.

1 Introduction

Increasing availability of 3D image data, e.g. clinical CT and MR, has spurred growing interest in 3D image segmentation. By making full use of 3D image data, a segmentation algorithm is more accurate and more robust to noise and imperfect data than a 2D slice-by-slice approach.

Active surfaces [11] are a 3D generalization of the 2D active contour models (snakes) [8]. They are generalized splines with piecewise smoothness and have a diverse coverage of region shapes. They segment a 3D image by fitting the surface to desired image features, e.g. edge points. Once initialized in the vicinity of the desired image features, the surface evolves under the influence of internal and external forces until the balance of all forces is reached. The process of segmentation is formulated as the minimization of the energy cost functional of the active surface [11, 3]

$$E = \iint_{\Omega} [w_1 E_{\text{int}}(s(u, v)) + w_2 E_{\text{ext}}(s(u, v))] dudv, \quad (1)$$

where $s(u, v) = (x(u, v), y(u, v), z(u, v))$ defines the Ω surface parameterized by u and v ; E_{int} and E_{ext} are the internal and external energy, respectively; w_1 and w_2 are the corresponding weights. The internal energy enables the active surface to bridge gaps in edge data and to bypass image noise. The external energy serves to attract the active surface to the desired image features.

Polygonal closed meshes are commonly used to represent an active surface. Triangular and simplex [4] meshes can deform locally and thus are able to fit

various shapes accurately. A surface can also be represented by nodes (vertices of mesh elements). The number of surface nodes can be constant or variable. With a constant number of nodes, the inter-nodal distance (or the size of the mesh), and thus the surface resolution will generally change during surface deformation. To maintain adequate surface resolution, it is generally necessary to vary the number of nodes as the surface deforms. Surface nodes are typically restricted to move within a plane. In fully 3D deformable models, surface nodes are free to move in any direction, which enables full use of 3D image data.

Of the several active surface models proposed [12, 3, 6, 4, 1, 13], most are not fully 3D. For example, Cohen & Cohen's model [3] allows surface node to move only within a plane. Whitaker *et al* [13] use 2D active contours but introduce interslice energy generated by extrapolating slice contours across slices. Some 3D models are excessively rigid due to the fixed number of surface nodes [12, 3, 4]. Huang *et al* [6] and Bulpitt & Efford [1] propose fully 3D surface models with triangular meshes and a variable number of surface nodes. Huang's model is designed specifically for the reconstruction of the left ventricle motion. Bulpitt & Efford's model aims at maintaining constant computational cost while setting adequate surface resolution by allocating more nodes to the high curvature parts of the surface.

Current models lack appropriate long distance attraction forces [9]. In particular, gradient forces operate over a short range. Some models overcome the problem by introducing spring forces [8, 4, 6, 1] which attract surface nodes to desired image features, e.g. edges. However, finding a node's corresponding edges is difficult and computationally expensive [6]. In the balloon model [2], a force normal to the surface inflates or deflates the surface to help drive it towards desired image features. It is an artificially added external force which, as implemented, is not adaptive.

Another limitation, crucial to successful implementation, is the shortage of practical methods for adapting model parameters to image data [9]. It is often desired to vary the values of model parameters from one image to another, from one iteration to another, and from one surface node to another. Among other major limitations, active surface models tend to shrink under the influence of internal force. They are also prone to becoming trapped by noise and irrelevant edge points. Though computationally efficient compared with region-based segmentation methods, computational cost for 3D image segmentation is significant due to the large data set.

2 Surface representation

Closed triangular meshes are used to represent the surface in our model (Figure 1). Each edge of a triangle is shared by exactly two triangles. This is the so called conforming triangular mesh [10]. Non-conforming triangular mesh will complicate the calculations of continuity and smoothness energies. The mesh elements are planar. The surface can also be described by surface nodes. Each node is directly connected to its (first order) neighbour nodes, and is shared by

its (first order) neighbour triangles. The number of the neighbour nodes equals the number of the neighbour triangles. This number can vary from node to node, enabling the model to deform locally, thus permitting a diverse range of surface shapes.




Fig. 1. Triangular meshes. A node j (solid square) is surrounded by its first order neighbour nodes (solid circles). Triangles with node j as a vertex are the first order neighbour triangles.

Our model allows surface nodes to move in all directions during surface deformation to make full use of the available image information. As a result of node movement, the inter-nodal distances will change. A large inter-nodal distance may lead to important image features being ignored. On the other hand, a small inter-nodal distance requires more surface nodes to represent a surface, thus increasing computational cost. To maintain adequate surface resolution and high computational efficiency, node-insertion and node-removal [5, 7, 1] are introduced. A new node is inserted somewhere between two directly connected nodes if their inter-nodal distance is larger than a threshold. A node is removed if the inter-nodal distance is smaller than another threshold (see [15] for details).

The thresholds for both node-insertion and node-removal are application dependent. Also, they can be made dependent on surface smoothness and time (iteration). Surface smoothness dependency is effected by raising the thresholds where the surface is smooth. This will make dense distribution of nodes on parts of the surface with high curvature and sparse distribution on parts with low curvature. The time dependency is implemented by starting with and maintaining a small number of surface nodes during surface evolution, and adding more nodes in the final stages of the evolution. This control over node density can help speed up the evolution while providing the final segmentation result with adequate surface resolution (see section 5 for more discussion).

3 Energies

In this section, we describe the internal and external energies of our model as well as the self-inflation/deflation forces.

3.1 Surface energy

In our model, only smoothness energy is used for the internal energy, and the external energy simply consists of the edge energy. Kass' continuity energy [8] is redundant because our model maintains inter-nodal distance by node insertion and removal. Furthermore, continuity energy tends to pull surface nodes closer, leading to undesirable surface shrinking. The discrete form of our energy function is expressed as a sum over M surface nodes,

$$E = \sum_{j=1}^M (w_j^s E_j^s + w_j^e E_j^e), \quad \text{where } E_j^e = -|\nabla I_j|^2; \quad (2)$$

at each node j , E_j^s is the smoothness energy, E_j^e the edge energy; w_j^s and w_j^e the corresponding weights, and ∇I_j denotes the gradient at node j of image intensity I . Since the method for calculation of the smoothness energy of a 3D active surface cannot be readily generalized from that of the 2D active contour, we designed our smoothness measures based on the angles formed by the surface normal at a node and the edges connecting the node to its neighbour nodes (see [15] for details).

3.2 Self-inflation and self-deflation forces

The self-inflation/deflation forces arise from the manipulation of the smoothness energy. The inflation/deflation switch and the adjustment of the force strength can be done merely by tuning a model parameter. The self-inflation/deflation forces enable the active surface to travel a long distance without any help from external forces (see [15] for details).

3.3 Energy minimization

Energy minimization is performed iteratively by the greedy algorithm [14]. A search neighbourhood consists of six closest voxels in 3D. A surface node is free to move to one of its neighbour voxel positions, or stay at its current position, depending on the location of the energy minimum. The size of the neighbourhood affects the number of iterations needed for surface evolution. Also, it affects the computational cost for each iteration. The larger the size of the neighbourhood, the higher the computational cost for each iteration.

4 Adaptation of model parameters

Noise and irrelevant edge points give rise to difficulties in image segmentation. Active surface nodes often stick to such points and fail to reach true edges. To reduce these effects, edge maps are commonly thresholded. However, thresholding eliminates only some noise and irrelevant edge points. We propose an adaptation of the edge energy weight w_j^e (Eq. 2) during surface evolution to help the model bypass noise and isolated irrelevant edge points.

The algorithm for adaptation of the edge energy weight is based on the assumption that noise and irrelevant edge points are isolated compared with true edge points. During surface evolution, once a surface node attaches to a true edge point, most of its neighbour nodes are expected to adhere to other true edge points in the next few iterations. If a surface node has stuck on a high gradient point and, after a few iterations, few of its neighbour nodes also stick to some other high gradient points, the point is regarded as an isolated noise or irrelevant edge point. By relaxing w_j^e (Eq. 2) at that surface node, active surface can bypass the noise or isolated irrelevant edge point under the action of the smoothness force. The procedure is as follows:

1. test if the edge strength at a surface node is larger than a certain threshold; if not, test the next surface node;
2. count the neighbour nodes whose edge strength also exceeds the threshold;
3. if the counter is less than, say 30%, in several consecutive iterations, say 5, relax the weight of edge energy at that surface node.

In generally, this algorithm is less effective for 2D active contour models because each contour node has only two neighbour nodes. This manifests one of the advantages of full 3D segmentation over 2D.

5 Computational efficiency

Though active surface models are computationally efficient relative to region-based segmentation methods, computational cost for 3D image segmentation is significant due to the large image matrix. Computational cost of our model with greedy algorithm is proportional to the number of iterations needed for the active surface to converge to region boundaries. For each iteration, the computational cost is proportional to the number of surface nodes and the size of the search neighbourhood for each node.

A large surface mesh size suggests a small number of surface nodes needed for surface representation, reducing the computational cost for surface evolution. It also suggests low surface resolution. To minimize computational cost while obtaining adequate surface resolution, a large surface mesh size is applied at early stages of surface evolution, whereas a small size of surface mesh is applied at final stages. The size of the surface mesh is controlled by the inter-nodal distance thresholds for surface node insertion and removal. Therefore, by manipulating the inter-nodal distance thresholds, computational cost can be reduced.

The active surface model, with the greedy algorithm, lends itself well to a parallel implementation because the measures of both smoothness and edge energies are local. We implement the parallelisation by splitting the linked list, which connects surface nodes, into sections and assigning each section to a different processor. The algorithm runs on a multi-processor architecture (Sun Ultra-2 with two processors). We are investigating the computational gain of this and other possible approaches to parallelisation.

6 Experiments and results

Two 3D synthetic images and a 3D CT image were used. The synthetic images were of size $128 \times 128 \times 128$ with added Gaussian noise of signal to noise ratio of 1. A simple image contains a single ellipsoid and a compound image contains two partly merged ellipsoids. The CT image, from the Visible Human Set at the National Library of Medicine, Bethesda, MD, USA, was interpolated to 1mm interslice spacing. A Sun SPARC-10 workstation was used in our experiments.

Experiment 1: To demonstrate the deformability of the active surface model with the self-inflation/deflation forces, the compound image was segmented with an initial spherical surface placed within the object (Figure 2). The segmentation was successful even though the object splits into disjoint regions in some slices.

Experiment 2: To demonstrate the ability to bypass isolated irrelevant edge points by parameter adaptation, a bar and a cluster of high gradient points were added to the edge map of the simple image (Figure 3). Normally, the active surface sticks to the high gradient points (Figure 3c). With the adaptation of edge energy weight w_j^e (Eq. 2), the active surface bypassed the high gradient points (Figure 3d).

Experiment 3: To investigate the computational efficiency, the simple synthetic image was segmented with and without manipulating the number of nodes during surface evolution. An initial spherical surface was placed within the object. Controlling the number of surface nodes resulted in about 40% reduction of computational time (from 5m 35s to 3m 15s).

Experiment 4: To test the model incorporating the self-inflation/deflation forces and the adaptation of model parameters, a kidney in the 3D CT image was segmented (Figure 4). The initial surface was a sphere totally contained within the kidney. Examination of each cross section of the segmentation result by eye revealed that the segmentation was accurate.

7 Conclusion

We present a fully 3D active surface model with the self-inflation and self-deflation forces. Our surface model has the following features: surface nodes are free to move in all directions, thus enabling the model to make full use of 3D image information; the model is deformable locally, thus providing coverage for a diverse range of surfaces; the model is capable of strong deformations, thus reducing its sensitivity to initial shape and position. The self-inflation/deflation forces permit the surface to travel a long distance without the aid of external forces. They are simple and easy to control. We introduce an adaptation of model parameters. By adapting the weight of edge energy at each surface node, the model is capable of bypassing some noise and isolated irrelevant edge points. By adapting the inter-nodal distance thresholds for node insertion and node removal, the model is capable of adjusting the number of surface nodes during surface evolution so that the computational cost is minimum compatible with adequate surface resolution of segmentation.




Fig. 2. The representative parallel cross sections of the compound synthetic image showing the initial (top row) and the final surface (bottom row).




Fig. 3. A central cross section of the simple synthetic image showing the initial surface (a), edge map with added high gradient points (b), and the final surface segmented without (c), and with (d), parameter adaptation.




Fig. 4. A central cross section of the CT image with the initial (a) and final (b) surface. A perspective view of the final surface mesh is shown in (c).

The model is tested with two synthetic images and a 3D CT image. Accurate segmentation results were obtained. The full 3D active surface model produced correct segmentation where an object splits into disjoint regions in some slices. By combining the self-inflation/deflation forces and the adaptation of model parameters, our model is computationally efficient, capable of converging to the correct boundary even if the initial estimate is not close, and appears to be robust to noise and imperfect image data.

References

1. Bulpitt A.J. and Efford N.D.: An effective 3d deformable model with a self-optimising mesh. *Image Vision Computing* **Vol.14** pp 573–580, August 1996.
2. Cohen L.D.: On active contour models and balloons. *Computer Graphics and Image Processing: Image Understanding* **Vol.53, No.2**, pp 211–218, 1991
3. Cohen L.D. & Cohen I.: Finite–element methods for active contour models and balloons for 2–D and 3–D images. *IEEE Trans. Pattern Analysis and Machine Intelligence* **Vol.PAMI–15, No.11**, pp 1131–1147, 1993
4. Delingette H.: Adaptive and deformable models based on simplex meshes. In *IEEE Computer Vision and Pattern Recognition CVPR'94*, pp 152–157 1994
5. Hoppe H. *et al*: Piecewise smooth surface reconstruction. *Computer Graphics Proceedings, Annual Conference Series, 1994 SIGGRAPH 94*, pp 295–302 1994
6. Huang W.C. & Goldgof D.B.: Left ventricle motion modeling and analysis by adaptive–size physically–based models. *SPIE—Biomedical Image Processing and Three–Dimensional Microscopy* **Vol.1660**, pp 299–310 1992
7. Gupta A., O'Donnell T. & Singh A.: A 3–D deformable model for segmentation and tracking of anisotropic cine cardiac MR images. *SPIE—Image Processing* **Vol.2167**, pp 77–86, 1994
8. Kass M., Witkin A. & Terzopoulos D.: Snakes: active contour models. *International Journal of Computer Vision* **Vol.1**, pp 321–331, 1988
9. Leymarie F. & Levine M.D.: Tracking deformable objects in the plane using an active contour model. *IEEE Trans. Pattern Analysis and Machine Intelligence* **Vol.PAMI–15, No.6**, pp 617–634, 1993
10. Rivara M.C.: Algorithms for refining triangular grids suitable for adaptive and multigrid techniques. *International Journal for mathematical methods in engineering* **Vol.20**, pp 745–756, 1984
11. Terzopoulos D., Witkin A. & Kass M.: Constraints on deformable models: recovering 3D shape and nonrigid motion. *Artificial Intelligence* **Vol.36**, pp 91–123, 1988
12. Terzopoulos D. & Vasilescu M.: Adaptive surface reconstruction. *SPIE—Sensor fusion III: 3–D perception and reconstruction* **Vol.1383**, pp 257–264, 1990
13. Whitaker J.M. & Braun M.: Three dimensional image segmentation using active contours with interslice energy. *Proc. APRS/CBT Image Segmentation Workshop*, Sydney, Australia pp 47–51, December 1996
14. Williams D.J. & Shah M.: A fast algorithm for active contours and curvature estimation. *Computer Graphics and Image Processing: Image Understanding* **Vol.55, No.1**, pp 14–26, 1992
15. Zhang Z. & Braun M.: Fully 3D active surface models with self-inflation and self-deflation forces. *Proc. IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'97)* San Juan, Puerto Rico, June 1997