Decoupled Active Surface for Volumetric Image Segmentation

A. Mishra, Member, IEEE, P. W. Fieguth, Member, IEEE
D. A. Clausi, Senior Member, IEEE

Abstract

Finding the surface of a volumetric 3D object is a fundamental problem in computer vision. Energy minimizing splines, such as active surfaces, have been used to carry out such tasks, evolving under the influence of internal and external energies until the model converges to a desired surface. The present deformable model based surface extraction techniques are computationally expensive and are generally unreliable in identifying the surfaces of noisy, high-curvature and cluttered 3D objects.

This paper proposes a novel decoupled active surface (DAS) for identifying the surface of volumetric 3D objects. The proposed DAS introduces two novel aspects which leads to robust, efficient and accurate convergence. First, rather than a parameterized surface, which leads to difficulties with complex shapes and parameter singularities, the DAS uses a conforming triangular mesh to represent the surface. Second, motivated by earlier successes in two-dimensional segmentation, the DAS treats the two energy components separately and uses novel solution techniques to efficiently minimize the two energy terms separately. The performance of DAS in segmenting static 3D objects is presented using several natural and synthetic volumetric images, with excellent convergence results.

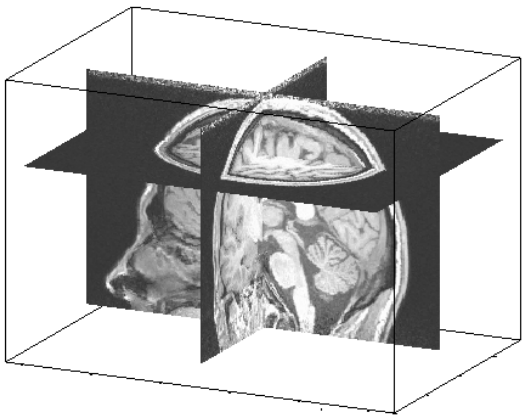

Figure 1. MRI of the brain [4]: producing a 3D segmentation of such data has countless applications in diagnosis and treatment.

1 Introduction

The job of identifying the surface of a three dimensional object in the presence of noise and background clutter has applications throughout medical image analysis [2, 14], surface reconstruction [6, 12], visual tracking [15] and motion estimation [17]. As shown in Fig. 1, a surface extraction technique is essential to segment and quantify important structures of a volumetric image for successful diagnosis. Active surfaces, the extension of 2D active contours to 3D, are now primarily used to perform such surface extraction tasks. The key idea in active surfaces is to evolve a deformable model [5, 8, 16] under the influence of internal (prior) and external (measurement) energies to capture the surface of a 3D object. Typically, the internal energy that maintains surface smoothness is defined using elastic / thin-plate constraints, whereas the external energy that pulls the deformable model towards the object boundary is usually defined using an image gradient (or, possibly, other image metrics of interest). Many approaches based on the original active surface idea have been proposed, broadly falling into two categories: parametric [8, 13, 16], and non-parametric [1, 3].

All of these active-surface methods are direct extension of their 2D active-contours counterparts to the 3D context. For parametric approaches, the 3D surfaces are explicitly represented [8, 16] using a 2D regular grid. There are two primary problems with such a representation scheme. First to have a unique mapping from 3D surface to regular 2D grid for all geometries is not feasible. Second, even for regular geometries the mapping is not unique and is not evenly spaced. Most parametric surface approaches are not able to capture

978-0-7695-4040-5/10 $26.00 © 2010 IEEE
DOI 10.1109/CRV.2010.45

high-curvature surfaces in the presence of noise and background clutter, without parameter tuning. On the other hand, non-parametric active surfaces are robust to initial solution and are capable of finding multiple high curvature surfaces, however these methods are extremely slow and are not able to extract the discontinuous surface of a 3D object.

Many of the above problems — lack of robust convergence in the presence of clutter, and high computational complexity — are also present in two-dimensional active contours. To address these problems in 2D, Mishra et al. [11] implemented a decoupled active contour (DAC) which successfully identifies the boundary of a single object in the presence of noise and background clutter. The novelty of the DAC was to separately address the external energy (measurement) from the internal energy (prior). The method uses a Viterbi search algorithm on the external energy to locate high-gradient boundary like features. Then, as a separate step, the internal energy representing the non-stationary prior is enforced on the measured boundary using a Bayesian linear least square estimator.

Next, a fast decoupled active contour [10] (FDAC) was developed, specifically anticipating the computational limitations of regular active contours, including the DAC, in being applied to large, three-dimensional problems. The FDAC replaced the Viterbi gradient-finding step of DAC with an iterative, quasi-random search, and the assertion of the prior model, accomplished as a Bayesian least-squares problem, was implemented in a sparse conjugate gradient algorithm.

This paper extends the concepts of DAC [11] and FDAC [10] from 2D to 3D and proposes a decoupled active surface (DAS) for identifying the surface of a volumetric 3D object. Unlike traditional active surface based methods, instead of representing the 3D surface using a regular 2D grid, DAS uses a conforming triangular mesh to represent the 3D surface and employs a decoupling co-ordinate descent concept to find an intermediate surface in the vicinity of the initial solution. The segmentation accuracy and computational advantage of DAS compared to one existing parametric active surface based method are demonstrated using several synthetic and natural 3D volumetric images.

The rest of the paper is organized as follows. We begin Section 2 with a brief discussion on the theory of parametric active surfaces, followed by identifying the limitations of traditional approaches. The theory and implementation of DAS are provided in Section 3. Finally, the experimental results and the validation of DAS claims are provided in Section 4.

2 Parametric Active Surface

An active surface is an energy-minimizing 3D spline, represented using a regular grid $v(s,r) = [x(s,r), y(s,r), z(s,r)]$, whose total energy is expressed as

$$E = \int_0^1 \int_0^1 \left(E_{\text{int}}(v(s,r)) + E_{ext}(v(s,r))\right) dsdr. \quad (1)$$

$E_{\text{int}}(v(s,r))$ and $E_{ext}(v(s,r))$ are the internal and external energies, respectively, of the active surface. The internal energy ensures the the prior or inherent constraints on the surface, typically penalizing slope and/or curvature terms:

$$E_{\text{int}}(v(s,r)) = \left(\frac{dv}{ds}\right)^2 + \left(\frac{dv}{dr}\right)^2 + 2\left(\frac{d^2v}{dsdr}\right)^2 + \left(\frac{d^2v}{ds^2}\right)^2 + \left(\frac{d^2v}{dr^2}\right)^2. \quad (2)$$

The external energy creates an attractive force towards the desired boundary, normally a region of high gradient:

$$E_{ext}(v(s,r)) = -\left(\nabla G_\sigma * I_3\right)^2 (v(s,r)), \quad (3)$$

Here ∇ is a derivative operator and G_σ is a Gaussian kernel of bandwidth σ. By considering the active surface to be deforming with respect to time, the solution to (1) can be presented using a set of linear equations [5, 8] as shown:

$$A\mathbf{x} + F_x = 0, \quad A\mathbf{y} + F_y = 0, \quad A\mathbf{z} + F_z = 0. \quad (4)$$

The iterative solution [5, 8] to (4) using an Euler-Lagrangian formulation is found to be

$$\mathbf{x}_t = \left(I + AA^T\right)^{-1}\left(\mathbf{x}_{t-1} - \gamma F_{\mathbf{x}_{t-1}}\right) \quad (5)$$

and similar equations hold for $\mathbf{y}_t$ and $\mathbf{z}_t$. The Matrix A is a banded matrix that contains the internal energy and I is an identity matrix. The convergence of the traditional active surface using (5) is demonstrated in Fig. 2(a) using a synthetic volumetric cube (red) with a spherical initial solution (black). Boundary condition enforcement is a complicated task in the traditional representation schemes. As shown in Fig. 2(b), (c) and (d), different boundary conditions lead to substantially different solutions. Using free boundary and free pole conditions in (5), we get a broken surface (Fig. 2(b)), similarly free pole condition generates holes at both poles of the converged surface (Fig. 2(c)). On the other

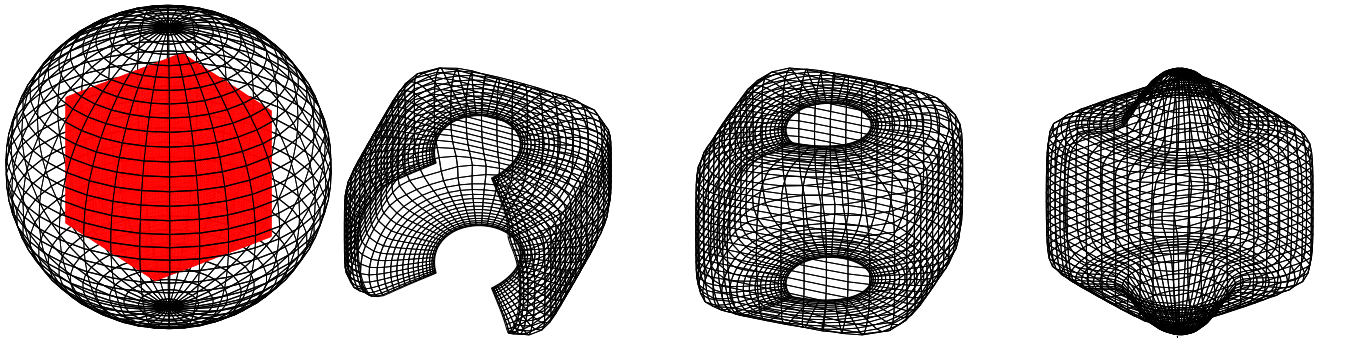

(a) Initial boundary (b) Free pole and boundary (c) Constrained boundary and free pole (d)Constrained boundary and pole

Figure 2. Three examples (b, c and d) illustrating the limitations of traditional, parameterized active-surface schemes. We start with an initial, parameterized spherical surface (a), seeking to converge to the red synthetic cube. The partially-converged active surfaces after 500 iterations are shown in panels (b), (c) and (d). Parametric models have difficulties with boundaries (b) and singularities at poles (c,d).

hand, constrained boundary and poles lead to a biased internal force which creates a non uniform vertex motion; that the vertices near the poles have a lower velocity compared to other vertices can be observed from (Fig. 2(d)).

In summary, the following three problems are prevalent with the conventional active surface based representation scheme and solution technique:

- The conventional 3D surface representation scheme $v(s,r) = [x(s,r), y(s,r), z(s,r)]$ is not able to handle complex geometries. Even for a simple geometry, such as for a sphere as shown in Fig. 2(a), there exist one-to-many mappings from the 3D surface to the 2D grid. Such a one-to-many mapping scheme creates a biased internal force (prior), resulting in non-uniform motion of the vertex of the 3D surface. As shown in Fig. 2(d), such a bias creates a weak attractive force near the pole compared to other regions, resulting in non-uniform vertex motion.
- The conventional active surface is a direct extension of 2D active contour to 3D and uses iterative gradient descent technique to solve (1). Such an iterative solver is very slow and prone to local minima.
- There exists a delicate relationship between the parameters in (1).

Most of the present parametric active surface based segmentation methods, such as gradient vector flow (GVFS) [16] and vector field convolution [8], have attempted to solve some of these problems by increasing the capture range of the traditional active surface by diffusing the traditional gradient-based external force. However, these modifications are not able to overcome the problems of local minima and convergence speed. A description of our proposed decoupled active surface, which seeks to deal with these problems, follows next.

3 Decoupled Active Surface

A decoupled active surface is a random field (surface spline) (v) represented using closed triangular meshes. As shown in Figs.3and 4, the triangular meshes are represented using a set of surface nodes $v_i, i \in [1, q]$ and triangular faces $f_j, j \in [1, f]$, where $v_i = [x_i, y_i, z_i]$. We use a conforming triangular mesh, where an edge of a triangle is shared exactly by two triangles and a node is connected to its first order neighbors. These constraints are necessary to keep the prior computation step simple. The number of neighboring triangles of a node is equal to the number of its first-order neighbors. The number of neighboring nodes of a particular node varies to facilitate the surface spline to be able to accommodate to a wide variety of shapes. Given the above surface spline representation scheme and an observed volumetric image I_3, we aim to find an optimal surface spline $\mathbf{v}^*$ by maximizing the posterior probability of the random field $\mathbf{v}$:

$$\mathbf{v}^* = \arg\max \left(p\left(I_3|\mathbf{v}\right) p\left(\mathbf{v}\right) \right), \tag{6}$$

where $p(I_3|\mathbf{v})$ and $p(\mathbf{v})$ are the likelihood and prior probability of $\mathbf{v}$.

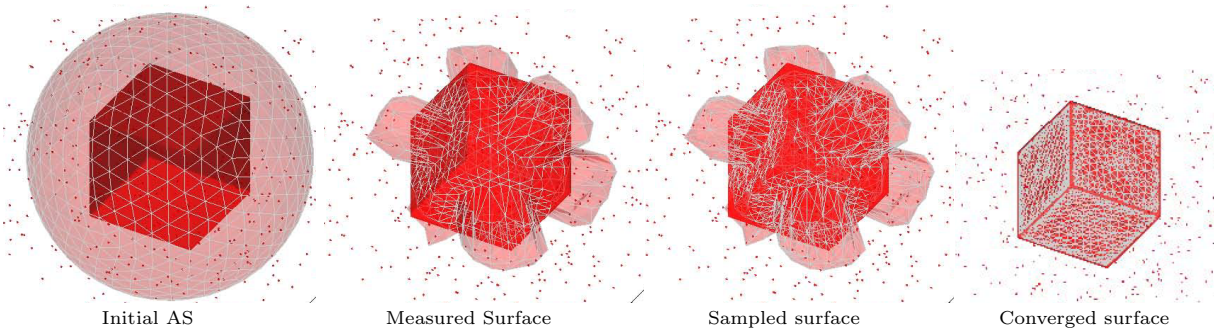

Figure 3. The three important steps of DAS: for a given initial surface (left), gradients are measured, and the resulting surface resampled on the basis of curvature. Despite the presence of noise in this example, the proposed iterative approach converged well to the desired cube, right.

In practice, we seek to obtain a discrete surface $\mathbf{v}^*$ using a non-iterative method that maximizes (6). However, no analytic method is available to solve (6). Therefore, instead of trying to solve the maximum a posteriori (MAP) (6) problem directly, DAS solves (6) in three steps using concepts analogous to those of the decoupled active contour [11, 9]. First, an approximate surface $\mathbf{v^m}$ is extracted from the likelihood $(p(I_3|\mathbf{v}))$ using an iterative quasi-random search (IQRS) algorithm [10]. Second, a non-stationary prior is secured to capture high curvature boundary by relaxing the prior penalty near regions of high curvature. Algorithmically, DAS achieves the effect of a non-stationary prior by employing a curvature-based mesh re-sampling scheme. Third, the prior knowledge about the surface, $p(\mathbf{v})$, is enforced with the measured boundary $\mathbf{v^m}$ using a Bayesian linear least square estimator.

To always find an ideal converged discrete surface $\mathbf{v}^*$ in a single iteration is not possible. Therefore, DAS takes an iterative approach, constraining the search space at time t based on the solution $\mathbf{v}_{t-1}$ from the previous time step.

The three steps involved in DAS are illustrated in Fig. 3. First, an initial spherical surface using a conforming triangular mesh is initialized around the object of interest (Fig. 3 (a)). Second, DAS uses an iterative quasi random search (IQRS) method [10] to find an optimal surface without the prior, within a specified search space (Fig. 3 (b)). The search space is defined using a set of normals, as shown in Fig. 4. Third, DAS uses an adaptive mesh re-sampling scheme to generate a non-stationary prior to capture high-curvature surfaces. The adaptive re-sampling scheme is described in Section 3.1. Fourth, similar to the DAC [11], DAS uses a Bayesian linear least square estimator to fuse the prior with the measurements.

Fundamentally, the surface nodes can be identified iteratively using IQRS without updating. However, it is not possible to incorporate complex shape priors and noise variances into IQRS, both of which are essential for reliable convergence, especially for noisy images.

3.1 Adaptive Mesh Re-sampling

The computational intractability of active surface based methods is well documented [10, 8, 3] and it is known that the computational complexity [10] of active surface based approaches is proportional to the number of active surface nodes q. Therefore, the practicality of active surface based methods can be realized by exploring methods that can help in reducing the number of surface nodes without sacrificing accuracy.

To reduce the number of active contour nodes and to capture high curvature regions of 2D objects, Mishra et al. [11, 9, 10] used a curvature-based importance sampling scheme to create more samples near high curvature regions than uniform regions. DAS extends the similar concept of DAC [11] from 2D to 3D for capturing high curvature surface and at the same time substantially reducing the total number of surface nodes. The curvature-based mesh re-sampling scheme follows.

The curvature k of a parametric 2D curve $v(s) = [x(s), y(s)]$, is defined as the rate of change of tangent angle with respect the arclength s and can be computed as:

$$\kappa = \frac{x_s y_{ss} - y_s x_{ss}}{(x_s^2 + y_s^2)^{3/2}}. \tag{7}$$

However, for a 3D surface, the curvature at a particular point varies as the plane through the normal at that point changes. Therefore, there is no unique definition of curvature for a 3D surface, so 3D curvature

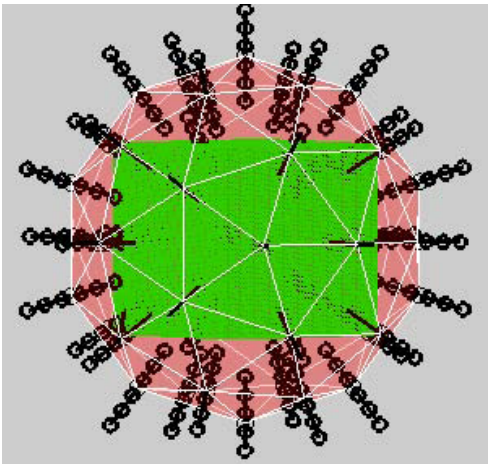

Figure 4. The 3D search space defined by a set of normals (black) to an initial surface (sphere). The sphere is triangulated using conforming triangles. The goal is to converge to the cube (green) inside the sphere.

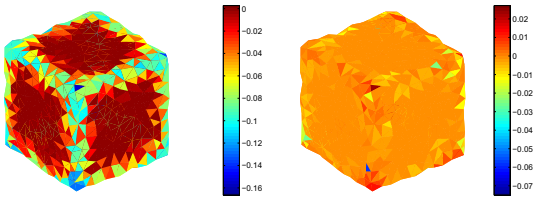

Figure 5. The principal and Gaussian curvatures of a cubic surface

is usually defined using principal and Gaussian curvatures. The principal and Gaussian curvatures of a cubic surface are shown in 5. In the context of our application, we do not need an exact definition of curvature, rather we want a technique that can measure the perpendicular distance of a surface node from its first-order neighboring nodes. To achieve this goal, we first fit a plane with the first-order neighbors of a surface node using a least-squares fitting algorithm, then we compute the perpendicular distance of the corresponding surface node from the fitted plane and we consider that distance as a notion of curvature. We repeat this procedure for all surface nodes to compute the curvature of a given 3D surface. The curvature profile obtained using the above approach for a cubic surface is presented in Fig. 6(b).

We then re-sample the 3D surface based on the computed curvature, such that the density of samples at any point increases with the curvature of the surface. The re-sampled cubic surface using the proposed re-sampling scheme is shown in Fig. 6c). More samples near the corners and edges of the cube can be observed in Fig. 6(c).

Table 1. Comparison table showing execution time in seconds for the surfaces in Fig. 7, comparing DAS and VFC [8]. The proposed DAS is more than an order of magnitude faster than the competing approach.

	DAS	VFC
(A)	**323**	12853
(B)	**579**	13423

4 Testing and Discussions

The capabilities of DAC [11] and FDAC [10] to segment 2D objects in the presence of noise and background clutter with fast convergence rates are presented in [11, 9, 10]. This paper generalizes the DAC/FDAC concepts to 3D and validates the claims of DAS using natural and synthetic 3D volumetric images. We compare the proposed DAS to the vector field convolution [8] active surface. The performance of DAS is validated on five natural and synthetic 3D data sets. The five volumetric images depict a wide range of characteristics including high curvature and noise.

We implemented part of the code for VFC [8] and downloaded rest of the code from [7]. All experiments are performed on a P4 Intel 2.4Ghz processor, 1Gb RAM using MatLab. The DAS's ability to identify the surface of 3D volumetric images is presented in Figs. 7 and 8.

As shown in Fig. 7, VFC could not identify the surfaces of both volumetric images (A and B). The unbiased internal force created due to the conventional representation scheme used in VFC smoothed the high curvature portion of the U-shaped object. In contrast, the DAS successfully identified the high curvature surface of both the objects (A and B), using separated measurement and prior steps, and also by enforcing a lower force in high curvature regions.

The convergence time of DAS compared to VFC is presented in Table 1. The convergence time of DAS is clearly significantly lower than that of VFC. The ability of DAS to identify the surface of noisy synthetic 3D images is shown in Fig. 7 (second row) and in Fig. 8, and in a real measured human tomogram in the top half of Fig. 8.

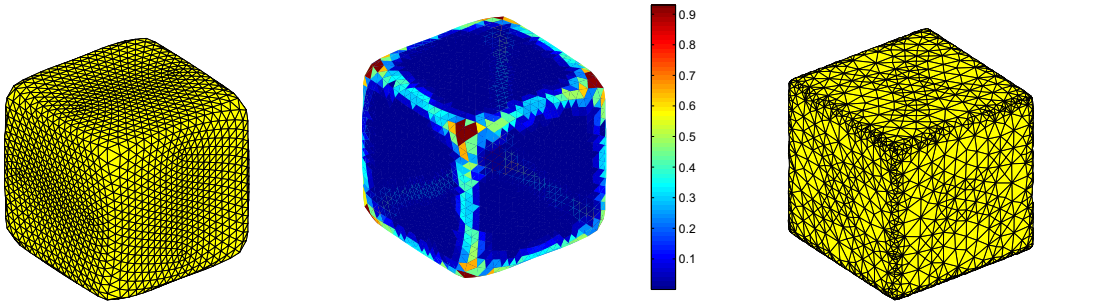

(a) Extracted surface without re-sampling (b) Curvature of the cube (c) Extracted surface with re-sampling

Figure 6. The re-sampling scheme. The cube, left is iterated without re-sampling, but ends up with an excess of grid points in flat regions, and possibly a deficiency in regions of high curvature. The right cube shows the mesh after resampling, based on the curvature as shown in the middle.

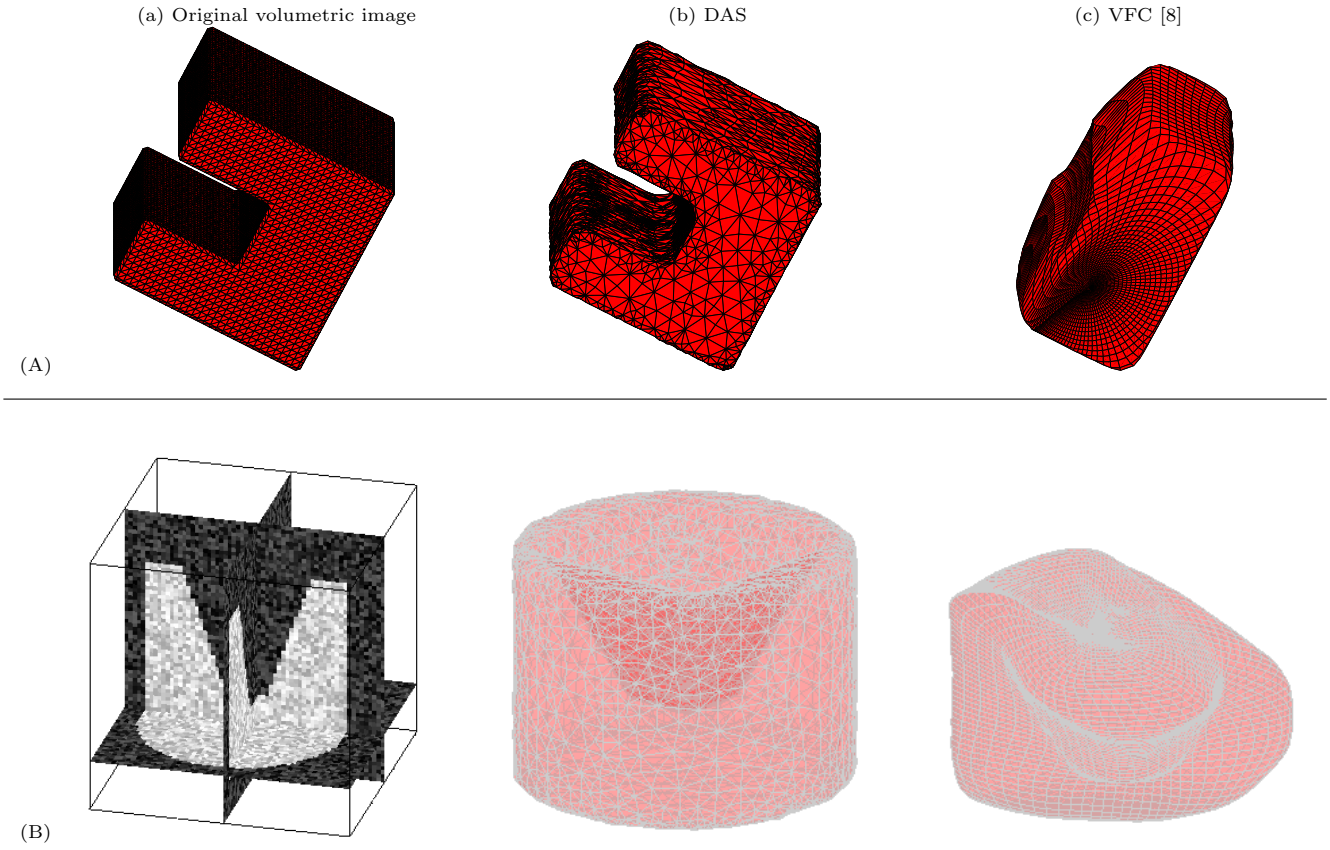

Figure 7. Boundary identification accuracy of DAS and VFC. We test on two surfaces: a U-shape (A), top, and a noisy concentric hollow cylindrical volume (B), bottom. The proposed DAS (middle) segments both objects successfully, however VFC (right) is ineffective for both objects.

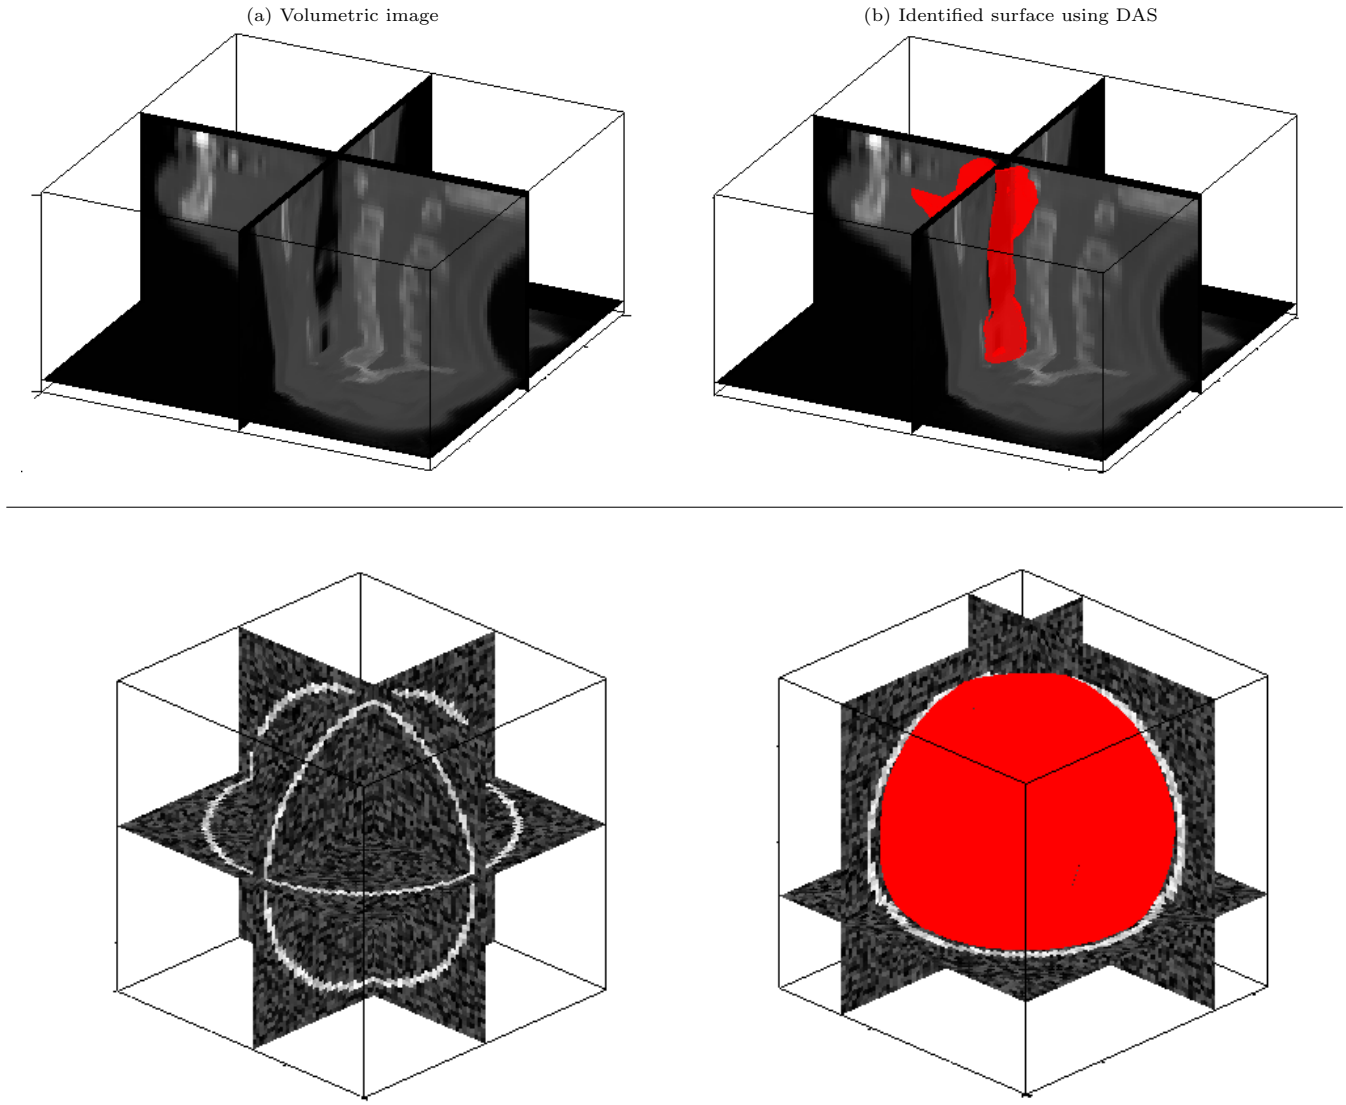

Figure 8. DAS successfully identified the throat using a human CT image (top), and the surface of a hollow sphere in the presence 40 percent additive Gaussian noise (bottom).

5 Conclusions

A decoupled active surface, an extension of DAC and FDAC from 2D to 3D, for identifying the surface of static 3D objects has been proposed and implemented. The dramatic computational gain and superior segmentation performance of DAS compared to one published parametric active surface based method have been demonstrated in this paper. The DAS's computational advantage over other parametric methods is mainly due the use of the iterative quasi random search as the non-iterative solver and conjugate gradient as the Bayesian least square solver. The superior segmentation performance of DAS in the presence of noise and high curvature is due to the use of the novel concept of separating the prior from measurement and the incorporation of a non-stationary prior using a mesh re-sampling scheme.

In the future, we seek to use DAS to find the trajectories of moving objects in video, and to validate DAS on a wide range of 3D data sets and comparing the results with other published parametric and non parametric methods.

Acknowledgment

This research has been sponsored in part by the Natural Sciences and Engineering Research Council (NSERC) of Canada through individual Discovery Grants as well as GEOIDE (GEOmatics for Informed Decisions) which is a Network of Centres of Excellence under NSERC and Vision and Image Processing Lab, Systems Design Engineering, University of Waterloo. The CT medical data was kindly made available by The Christie Hospital, Withington, Manchester, UK. We would like to thank Justin Eichel for allowing us to use his Brain MRI data.

References

[1] T. Chan and L. Vese. Active contours without edges. *IEEE Transactions on Image Processing*, 10(2):266–277, 2001.

[2] C. Davatzikos and J. Prince. An active contour model for mapping the cortex. *IEEE Transactions on Medical Imaging*, 14(1):65–80, March 1995.

[3] E. Debreuve, M. Barlaud, G. Aubert, I. Laurette, and J. Darcourt. Space-time segmentation using level set active contours applied to myocardial gated spect. *IEEE Transactions on Medical Imaging*, 20(7):643–659, July 2001.

[4] J. Eichel. Justin eichel - mri brain data, 2009. http://www.eng.uwaterloo.ca/ jaeichel/brain/mridata.html.

[5] M. Kass, A. Witkin, and D. Terzopoulos. Snakes: Active contour models. *International Journal of Computer Vision*, 1(4):321–331, 1988.

[6] G. Knopf and A. Sangole. Freeform surface reconstruction from scattered points using a deformable spherical map. *International Journal of Image and Graphics*, 6(3):341–356, July 2006.

[7] B. Li and A. Acton. Verginia image and video analysis, vector field convolution, 2009. http://viva.ee.virginia.edu/researchvfc/.

[8] B. Li and T. Acton. Active contour external force using vector field convolution for image segmentation. *IEEE Transactions on Image Processing*, 16(8):2096–2106, 2007.

[9] A. Mishra, P. Fieguth, and D. Clausi. Robust snake convergence based on dynamic programming. In *International conference on image processing*, pages 1092–1095, Florida, USA, September 2008.

[10] A. Mishra, P. Fieguth, and D. Clausi. Fast decoupled active contour (FDAC) for boundary detection. *IEEE Transactions on Image Processing*, Will be submitted, February, 2010.

[11] A. Mishra, P. Fieguth, and D. Clausi. Decoupled active contour (DAC) for boundary detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, accepted, Februray, 2010.

[12] J. Pons and J. Boissonnat. Delaunay deformable models: Topology-adaptive meshes based on the restricted delaunay triangulation. In *IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pages 1–8, 2007.

[13] G. Slabaugh and G. Unal. Active polyhedron: Surface evolution theory applied to deformable meshes. In *IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pages II: 84–91, 2005.

[14] G. Tsechpenakis, J. Wang, B. Mayer, and D. Metaxas. Coupling crfs and deformable models for 3d medical image segmentation. In *Workshop on Mathematical Methods in Biomedical Image Analysis*, pages 1–8, 2007.

[15] K. Varanasi, A. Zaharescu, E. Boyer, and R. Horaud. Temporal surface tracking using mesh evolution. In *European Conference on Computer Vision*, pages II: 30–43, 2008.

[16] C. Xu and J. Prince. Snakes, shapes, and gradient vector flow. *IEEE Transactions on Image Processing*, 7(3):359–369, 1998.

[17] J. Zhang, J. Gao, and W. Liu. Image sequence segmentation using 3d structure tensor and curve evolution. *IEEE Transactions on Circuits and Systems for Video Technology*, 11:626–6411, 2001.