

# Flexible And Topologically Localized Segmentation

Gunnar Johansson<sup>1</sup>, Ken Museth<sup>1</sup> and Hamish Carr<sup>2</sup>

<sup>1</sup>Graphics Group, Digital Media, Linköping University, Sweden

<sup>2</sup>School of Computer Science and Informatics, University College Dublin, Ireland

---

## Abstract

*One of the most common visualization tasks is the extraction of significant boundaries, often performed with isosurfaces or level set segmentation. Isosurface extraction is simple and can be guided by geometric and topological analysis, yet frequently does not extract the desired boundary. Level set segmentation is better at boundary extraction, but either leads to global segmentation without edges, [CV01], that scales unfavorably in 3D or requires an initial estimate of the boundary from which to locally solve segmentation with edges. We propose a hybrid system in which topological analysis is used for semi-automatic initialization of a level set segmentation, and geometric information bounded topologically is used to guide and accelerate an iterative segmentation algorithm that combines several state-of-the-art level set terms. We thus combine and improve both the flexible isosurface interface and level set segmentation without edges.*

Categories and Subject Descriptors (according to ACM CCS): I.3.5 [Computer Graphics]: Computational Geometry and Object Modeling I.3.6 [Computer Graphics]: Methodology and Techniques I.4.6 [Image Processing and Computer Vision]: Segmentation

---

## 1. Introduction


A common task in medical imaging is locating and segmenting objects such as tumors and internal organs. Two principal methods for segmenting such objects are isosurface extraction and level set segmentation, but each has drawbacks. *Isosurfaces*, surfaces of the form  $\{x : f(x) = h\}$  for an *isovalue*  $h$ , are cheap to extract and can be explored interactively, but often do not coincide neatly with a single object's boundary, due to non-uniform scanning, scan artifacts, aliasing, natural tissue variation, or multiple objects sharing a single isovalue. *Level set segmentation*, does not depend on a global measure of importance, but iteratively relaxes an estimated boundary until energy-minimization terms are met. This produces superior boundary segmentation, but is computationally expensive, requires an initial boundary to be specified, and can depend on global properties, which may not be fundamentally related to the object being extracted.

These approaches have complementary strengths and weaknesses. Isosurfaces can be explored visually as entire objects and analysed with contextual geometric and topological information, but do not always produce good boundaries. In comparison, level set methods can produce good boundaries by minimizing functionals defined with differen-

tial properties of the boundary and medical data. However, global level set minimizations are typically computationally intensive and local minimizations are highly dependent on the initial conditions, *i.e.* starting boundary.

In this paper, we combine these approaches, using the flexible isosurface [CS03] with a low-resolution version of the data for rapid identification of an initial surface of interest, followed by a localized, but full-resolution, level set refinement driven by several state-of-the-art segmentation terms, most importantly [CV01]. For the level set implementation we furthermore utilize highly optimized data structures and algorithms [NM06]. We thus use the strengths of several methods to remedy the weaknesses of each other, reducing both computational costs, memory footprints and to some extent user input. In addition, the topological neighborhoods defined by contour topology can be used to address two weaknesses of the original segmentation without edges, [CV01]. Since this method is based on the solution of large linear systems of equations, involving global intensity variance, the resulting segmentation is both relatively slow and sensitive to (irrelevant) background intensities.

We next review the relevant work on topological analysis and on level set segmentation. Section 4 then describes our



**Figure 1:** A small 3-D data set shown as a sequence a-f of isosurfaces on the left and as a contour tree on the right. Note the 1-1 correspondence between individual contours on the left and points in the tree on the right.

modifications to Chan-Vese level set segmentation using narrow band methods for efficiency and local geometric properties, instead of global properties, for guidance. We then show how to augment the original global segmentation without edges with several local regularization and edge-based terms for improved flexibility. Section 6 shows some results, while Section 7 presents our conclusions and future work.

## 2. Contour Trees and Flexible Isosurfaces

In a scalar field  $f : \mathbb{R}^3 \rightarrow \mathbb{R}$ , the *isosurface* for an *isovalue*  $h$  is the set  $f^{-1}(h) = \{x \in \mathbb{R}^3 : f(x) = h\}$ . This set can be called a *level set*, but we use *isosurface* to avoid confusion, as *level set* is also used for the deforming surface representation described in Section 3.

A *contour* is a connected component of an isosurface. As  $h$  increases, contours appear at local minima, join or split at saddles, and disappear at local maxima. Shrinking each contour to a point gives the *contour tree*, which tracks this evolution, as illustrated in Figure 1. Because each contour is a single point in the contour tree, it gives an abstract representation of all possible contours in a data set, with efficient access to individual contours and the ability to annotate contours with additional geometric and topological information.

The contour tree has been used for fast isosurface extraction [vKvOB\*97, CS03], abstract representation of scalar fields [BPS97, ZBB04], manipulation of individual contours [CS03, CSvdP04], and transfer function design [TFT04, WDC\*], among other purposes. Efficient algorithms are known to compute the contour tree for simplicial meshes [vKvOB\*97, CSA03], for trilinear inter-

polants [PCM03] and for digital 4/8 connectivity [TIS\*95, TFT04]: since level set methods use a discrete voxel representation, we use 6/16 connectivity.

In particular, the contour tree can be used for interactive exploration of a data set, as it gives a fast, efficient method of manipulating individual contours using *flexible isosurfaces* [CS03]. Moreover, geometric properties can be pre-computed for contours and stored in the contour tree, allowing exploration to be guided with quantitative measurements as well as visual feedback [CSvdP04].

Unfortunately, the contour tree only represents boundaries that can be defined in terms of a single isovalue. In many situations, particularly for experimentally-acquired biological data, the boundaries of interest are not isovalued contours, but need more sophisticated segmentation methods such as *level set segmentation*. Since these methods often need to be initialized with an approximate surface of interest, the flexible isosurface can be used to advantage in a hybrid interface.

## 3. Segmentation and Level Sets

Isosurfaces are not the only method available to segment boundaries between regions of interest and the background. Other segmentation methods balance the desire to detect sharp edges against the desire to extract smooth curves, often by minimizing a variational functional of a deforming contour or surface, with terms that may include first derivatives (for edge detection) and second derivatives (for smoothness).

An early example of this approach is the “snakes” of Kass et al. [KWT88] that iteratively deform an *active contour* represented by parametric Lagrangian curves until the curves

converge to a stable solution. Limitations of this model, however, include aliasing and difficulty handling changes of topology or self-intersections [MT95]. Worse, these limitations get worse as the dimensionality of the data goes up.

Osher & Sethian [OS88] gave an elegant and robust multi-dimensional alternative based on an implicit *Eulerian* model, called *level set segmentation*. In this model, a boundary (or interface) is treated as the zero contour of a time-dependent Euclidean distance function,  $\phi$ , embedded in the data set, *i.e.*  $\{\vec{x}(t) \mid \phi(\vec{x}(t), t) = 0\}$ . By convention,  $\phi$  is negative inside the contour and positive outside. The fundamental (Hamilton-Jacobi) equations of motion are easily derived by applying the chain-rule to  $\phi(\vec{x}(t), t) = 0$ . This essentially recasts the problem of arbitrary interface deformations into the problem of solving the following partial differential equations (PDE),

$$\frac{\partial \phi}{\partial t} = -\frac{d\vec{x}}{dt} \cdot \vec{\nabla} \phi \quad (1a)$$

$$= -\mathcal{F}(\vec{x}, \vec{n}, \phi, \dots) |\vec{\nabla} \phi|. \quad (1b)$$

where Eq. (1a) describes advection of  $\phi$  in the velocity vector field,  $d\vec{x}(t)/dt$ , and Eq. (1b) describes propagation of the interface,  $\phi$ , in its local normal direction,  $\vec{n} = \vec{\nabla} \phi / |\vec{\nabla} \phi|$ .  $\mathcal{F}$  denotes the scalar velocity of this normal motion, *i.e.*  $\mathcal{F} \equiv \vec{n} \cdot d\vec{x}(t)/dt$ . Moreover, the mean curvature of the interface,  $\phi$ , can be computed as  $\vec{\nabla} \cdot \vec{n}$  [MBW\*05]. Although Eq. (1a) and Eq. (1b) are formally equivalent, their numerical properties are distinctly different, and great care must be taken to obtain stable numerical discretizations, [OF03].

Medical segmentation is sometimes referred to as an ill-posed problem in the sense that it requires additional constraints (or information) to define a unique solution. These constraints are often user-defined and imposed through the regularization of the underlying computational model. As such it comes as no surprise that a large body of work exist on different level set methods for medical segmentation, see [Kim03] for a summary. While several people have come very close, there is still no “silver bullet” for automatic level set segmentation. However, one of the most successful is the recent level set method for “segmentation without edges” by [CV01]. It essentially defines the boundary as a level set,  $\phi$ , that partitions the image,  $I$ , into approximately piecewise-constant regions in a least square sense [MS89]. This can be expressed as a minimization problem of the intensity variance of  $I$  computed respectively inside and outside of  $\phi$ , which in turn leads to a set of Euler-Lagrange equations [ZCMO96] that are solved using semi-implicit integration. All in all this results in a system of globally defined linear equations that are solved iteratively till convergence. Due to its global nature this technique is known to be relatively invariant to the initial condition (*i.e.* the initial guess of  $\phi$ ). However, this robustness comes at the price of having to solve a linear system of equations globally. The consequence is that the original method of [CV01] suffers from two fundamental limitations; the numerical complex-

ity scales with the embedding space,  $I$ , as opposed to the size of the boundary, causing a computational bottleneck in 3D, and the solutions are sensitive to variations in  $I$  that might be irrelevant for the segmentation (*e.g.* background intensities). As will be explained in next section we address this issue using a localized narrow band formulation. Finally, it is a well known fact that sometimes boundaries are best defined from differential properties of  $I$ , *i.e.* using various edge detectors. For this reason we propose to use a combination of several proven segmentation techniques that can be summarized in the following level set equation

$$\frac{\partial \phi}{\partial t} = \left[ w_1 (I - C_1)^2 - w_2 (I - C_2)^2 \right] |\vec{\nabla} \phi| \quad (2a)$$

$$+ w_3 f(\vec{x}) |\vec{\nabla} \phi| \quad (2b)$$

$$+ w_4 \kappa |\vec{\nabla} \phi| \quad (2c)$$

$$+ w_5 \nabla g(\vec{x}) \cdot \nabla \phi \quad (2d)$$

$$+ w_6 \text{sign}(\nabla I \cdot \nabla \phi) \Delta I |\vec{\nabla} \phi| \quad (2e)$$

Eq. (2a) constitutes the terms for segmentation without edges, [CV01] where respectively the inside ( $C_1$ ) and outside ( $C_2$ ) image variances are given by


$$C_1 = \int_{\Omega_I} I(\vec{x}) H(-\phi(\vec{x})) d\vec{x} / \int_{\Omega_I} H(-\phi(\vec{x})) d\vec{x} \quad (3a)$$

$$C_2 = \int_{\Omega_I} I(\vec{x}) H(\phi(\vec{x})) d\vec{x} / \int_{\Omega_I} H(\phi(\vec{x})) d\vec{x} \quad (3b)$$

where  $\Omega_I$  denotes the global domain of the image and  $H(x)$  is a Heaviside function (*i.e.*  $H(x) = 1$  if  $x \geq 0$  otherwise  $H(x) = 0$ ). Eq. (2b) is a “weighted region” regularization term which minimizes a quantity given by the scalar function  $f(\vec{x})$  inside the boundary. A simple example is  $f(\vec{x}) = 1$  for which the volume (*i.e.* interior) of the boundary is minimized. Eq. (2c) is another regularization term that essentially minimizes the area (*i.e.* size) of the boundary through mean curvature ( $\kappa$ ) based diffusion. Eq. (2d) is the “Geodesic Active Contour” model of [CKS95] which is essentially an edge-based attractions term, where  $g(\vec{x})$  denotes an inverse edge detector typically defined as  $1/(1 + |\nabla G * I|^2)$  where  $G * I$  denotes a Gaussian convolution of the image data to suppress noise. Finally Eq. (2e) is the “Robust Alignment” term that aligns the boundary normals with the gradient field of the image. To allow for a flexible weighting of these different state-of-the-art segmentation terms we introduce the scaling parameters  $\{w_i, i = 1, \dots, 6\}$ . We finally note that in [CV01]  $w_1 = w_2 = 1$  and  $w_3 = w_5 = w_6 = 0$  whereas  $w_4$  is varied to control smoothness of the segmentation, but is always non-negative to avoid instabilities.

#### 4. Topologically Local Chan-Vese Segmentation

One drawback of the Chan-Vese method [CV01] is that the segmentation relies on global properties through Eq. (3). Although the global convergence is often cited as the major



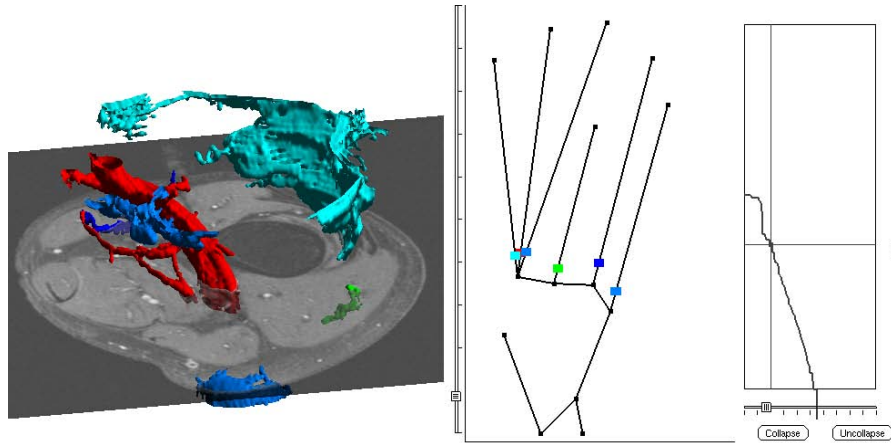
**Figure 2:** Illustrates the global behavior of the Chan-Vese method and the implications of using a narrow-band scheme initialized from the contour shown in (1). The upper row (a) and the lower row (b) differ only in the intensity of the background (black (a) vs. very dark grey (b)). Although semantically this should not affect the final result, the outside image variance term is affected by it, leading to two different results (2a) and (2b) depending on the background intensity. Narrow-band methods improve the result as shown in (3a) and (3b), but the intensity difference still causes different results. Our methods, which use topology to localize the Chan-Vese computation, are more robust to these differences, extracting essentially the same contour in each case (4a) and (4b).

strength of this method, for some applications in medical segmentation it is also the major weakness, as it fails to represent local properties. For our work, it is important to localize the behavior for two reasons. First, computational efficiency is required: for 3D medical datasets, which often exceed  $512^3$  in size, the Chan-Vese method leads to linear systems of dimensions ( $> 100$  million) that are simply not feasible to solve. Secondly, the global solution is not very meaningful for applications where only one object is of interest. Recall that the original Chan-Vese method is automatic in the sense that it does not allow for any direct control of the result - it always extracts all boundaries separating constant variance of the image intensities.

Figure 2 illustrates the drawbacks associated with Chan-Vese segmentation of a single 2D slice, initialized by the contour shown in Figure 2(1), by comparing the results obtained with two very slightly different background intensities (a) - black vs. (b) - very dark grey. Since Chan-Vese depends on global terms, these small intensity changes over large areas disturb the computation significantly, leading to two distinct segmentations (2a) and (2b). Moreover, the original Chan-Vese method used in these two images leads to the result being cluttered with unrelated additional structures.

In comparison, (3a) and (3b) show the result of the more efficient narrow band structures - while this suppresses the additional structures quite effectively, the images still converge to different solutions. Global convergence and the abil-

ity to detect inner contours are results of solving the level set equation implicitly on the full domain without re-initializing the function (*i.e.* solve  $|\nabla\phi| = 1$ ). Unfortunately, narrow-band schemes rely on re-initialization to rebuild and propagate the narrow band, thus preventing the Chan-Vese method from detecting inner contours. Therefore, using the Chan-Vese model with narrow-bands has significant implications on its behavior, a consequence which has not, to our knowledge, been noted in previous work. Figure 2(3a) shows the result of solving the level set equation in a narrow-band while computing the properties, Eq. (3), on the full domain. Since the computation is contained within the narrow-band, the solution does not propagate far from the object of interest. However, it is still very sensitive to changes in regions completely unrelated to the segmentation as shown in Figure 2(3b). To localize the behavior completely, we propose a simple but surprisingly robust approach which is well adapted to current narrow-band data structures; we limit all our computations to the narrow-band and the inside of the level set which we know contain the object of interest. This means *i.e.* in Eq. (3)  $\Omega_f$  is limited to the domain of  $I$  for which  $\phi < \gamma$  assuming  $\phi$  is a signed distance function with value  $-\gamma$  inside the narrow band. Given the initialization based on the contour tree, the segmentation is then based on topologically distinct regions of data, as selected by the user. Using this approach gives a stable convergence to the object of interest, as shown in Figure 2(4a) and Figure 2(4b).



**Figure 3:** Screen shot of the flexible isosurface and segmentation interface. From left to right; the isosurface visualization, the contour tree with the selected contours and the amount of simplification (collapsing).

## 5. Implementation

For our application, we use a reference implementation of the contour tree algorithms [CSA00] and the flexible isosurface interface [CS03], depicted in Figure 3. Using this interface one can effectively visualize and evolve contours individually by means of the contour tree. To deal with the problem of noisy datasets, it is possible to simplify the tree by collapsing contours based on certain geometric properties such as height or volume (right). When a target object has been selected, the segmentation is started and the parameters of the segmentation model can be tuned in real time.

Since the level set equation, Eq. (2), is a time-dependent Eulerian PDE, it defines an initial value problem. Consequently we need an efficient and robust procedure to convert the flexible isosurface into a level set so it can serve as the initial condition. To efficiently accomplish this we have developed the following procedure; employing the continuation method of [WMW86] we offset the contour by the selected isovalue of the flexible isosurface. Next to compute the signed distance we solve the Eikonal equation,  $|\nabla\phi| = 1$ , by means of the Fast Sweeping Method of [Zha04].

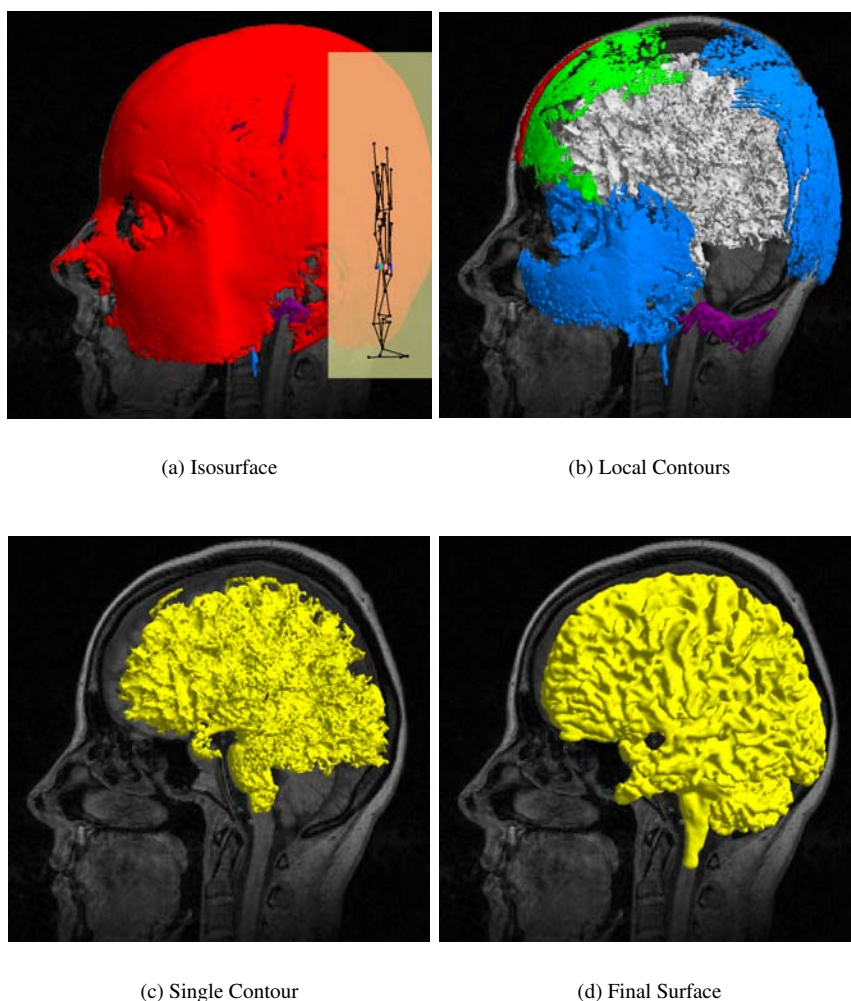
For the level set segmentation we use a reference implementation of the very compact and efficient DT-Grid level set data structure of [NM06]. This allows us to effectively represent and deform level sets of resolutions exceeding  $2048^3$  on a desktop PC with only 1GB of RAM. To solve Eq. (2), we use a first order forward Euler discretization in time and the variable 3rd-5th order WENO scheme, [LOC94], for the spatial discretization of the hyperbolic terms Eq. (2). For more general information on numerical finite difference schemes for solving level set equations we refer the reader to [OF03, Set99, MBW\*05].

As a final optimization we have developed a simple but efficient algorithm that localizes the computations of the in-

side ( $C_1$ ) and outside ( $C_2$ ) image variances in Eq. (3) as the boundary deforms. This is simply done by explicitly tracking the grid points that are respectively added and removed from the propagating narrow band and then incrementally update Eq. (3). As a result the computational complexity of the overall level set algorithm now scales with the size of the boundary as opposed to its enclosing domain.

## 6. Results

In Figure 4, Figure 5 and Figure 6, we show some sample segmentations obtained with our system. For Figure 4, we took the well-known UNC head MRI data set, computed the contour tree at full ( $256 \times 256 \times 109$ ) resolution, then applied the flexible isosurface interface (a) - (c) to select the contour shown in Figure 4(c) to initialize the level set segmentation. We then ran the modified Chan-Vese segmentation to produce the final surface shown in Figure 4(d). As we see, traditional isosurfaces, shown in Figure 4(a), produce large surfaces relating to the skull as well as the desired surface representing the brain. We then apply *local contour segmentation* [MHS\*96] based on a simplified contour tree to show the set of maximal topologically distinct large objects in the data in Figure 4(b). Although the brain is more clearly identifiable in this image, when we display it independently of the other surfaces, we note that the surface is incompletely developed, with numerous sharp edges instead of the rounded folds we would expect. This artifact occurs because the actual folds of the brain do not correspond to sharp isovalued edge features, and because, below the isovalue shown, a spurious narrow spur connects the brain to the skull fragments shown in Figure 4(b), resulting in an incorrect segmentation. Applying the modified Chan-Vese method, however, results in an appropriate segmentation in Figure 4(d).




**Figure 4:** *Interactive Brain Segmentation. In (a), an isosurface has been chosen, but the skull surface occludes the brain. In (b), local contours are used to display the largest topologically distinct surfaces. In (c), one of these contours has been selected visually as the initialization surface for level set segmentation. This contour fails to segment the brain properly since at lower isovalues a spurious connection is formed with the skull fragments. In (d), the final level set segmentation is shown. Using a level set segmentation instead of the single-isovalue contour of (c) allows us to produce a markedly superior result.*


In Figure 5 we show a segmentation based on MRI scans of a knee joint ( $512 \times 512 \times 28$ ), in which the clinical goal is to perform quantitative tests on the synovial capsule. As we see in Figure 5(a), isosurfaces generate multiple surfaces, only one of which (shown in (b)) is relevant. Furthermore, for isovalues lower than the one chosen, the contour representing the synovial capsule starts merging with these other contours into a single surface. While this single contour is not ideal, it is the best available using isosurface techniques alone, and we therefore choose it using the flexible isosurface interface as the initializing surface for our localized Chan-Vese segmentation. As we see from the magnified

view in (c), the level set methods have significantly improved the extracted capsule, filling in many of the visible holes and smoothing the surface significantly. Moreover, when we compared the segmentation with each slice of the image, it matched closely to the visual boundaries that demarcate the capsule. This segmentation completed in 6 minutes on an AMD Athlon 2.1 GHz machine.

Finally, Figure 6 shows a segmentation of a blood vessel in a  $112 \times 512 \times 416$  dataset. As we see in (a), ordinary isosurfaces generate multiple surfaces, which may not be helpful in initializing level-set methods. In (b), however,



**Figure 5:** *Synovial Capsule Segmentation.* In this data set, the clinical goal is to segment the synovial capsule for further processing. (a) shows a coarse-resolution isosurface with a reference slice of the data set. (b) shows a single contour selected from the isosurface using the flexible isosurface interface. (c) shows the final segmented boundary initialized from the contour in (b).



**Figure 6:** This figure illustrates that our segmentation method is capable of segmenting thin structures, such as blood vessels, given a rough initial surface. Due to the relatively large dataset ( $112 \times 512 \times 416$ ) the initial contour was extracted at quarter scale resolution to achieve rapid feedback when exploring the data.

we have used the flexible isosurface interface with a coarse  $28 \times 128 \times 104$  version of the data to select a single contour, then used that contour to initialize our localized Chan-Vese segmentation, resulting in the surface shown in (c). As we can see, the localized Chan-Vese segmentation can extract even complex branching structures effectively based on a coarse approximate contour chosen from a visual interface.

## 7. Conclusions and Future Work

We have shown that topological analysis, in the form of the contour tree, and boundary segmentation, in the form of level set methods, can be combined to produce a segmentation method that is less vulnerable than contours to unevenness

of isovalue, but allows intuitive exploration of contours to initialize level set segmentation methods.

In the future, we would like to explore methods of reducing the memory footprint of the contour tree computation, especially since much of the memory is currently used to represent small noisy features of the data set. We would also like to explore whether topological analysis can be performed for specific level set segmentation methods, allowing visual exploration of the space of possible segmentations. Furthermore, we are interested in solving the level set PDEs using semi-implicit schemes for faster convergence. Finally, we acknowledge that we still need to explore the full parameter space of our combined segmentation model given in Eq. (2), and validate the results quantitatively.

## 8. Acknowledgements

Acknowledgements are due to University College Dublin and to Science Foundation Ireland for research grant funding and to Drs. Patrick Brennan and John Ryan of the School of Diagnostic Imaging, University College Dublin, for making the knee dataset available for testing. We would also like to thank Michael Nielsen and other members of the Graphics Group at Linköping University for allowing us to use their software - in particular the DT-Grid implementation.

## References

- [BPS97] BAJAJ C. L., PASCUCCI V., SCHIKORE D. R.: The contour spectrum. In *Proceedings of Visualization 1997* (1997), pp. 167–173.
- [CKS95] CASELLES V., KIMMEL R., SAPIRO G.: Geodesic active contours. In *ICCV* (1995), pp. 694–699.
- [CS03] CARR H., SNOEYINK J.: Path seeds and flexible isosurfaces: Using topology for exploratory visualization. In *Proceedings of Eurographics Visualization Symposium 2003* (2003), pp. 49–58, 285.
- [CSA00] CARR H., SNOEYINK J., AXEN U.: Computing contour trees in all dimensions. In *Proceedings of the 11th Annual ACM-SIAM Symposium on Discrete Algorithms* (2000), pp. 918–926.
- [CSA03] CARR H., SNOEYINK J., AXEN U.: Computing contour trees in all dimensions. *Computational Geometry: Theory and Applications* 24, 2 (2003), 75–94.
- [CSvdP04] CARR H., SNOEYINK J., VAN DE PANNE M.: Simplifying flexible isosurfaces with local geometric measures. In *Proceedings of Visualization 2004* (2004), pp. 497–504.
- [CV01] CHAN T., VESE L.: Active contours without edges. *IEEE Trans. Image Processing* 10, 2 (2001), 266–277.
- [Kim03] KIMMEL R.: Fast edge integration. In *Geometric Level Set Methods in Imaging, Vision and Graphics*, Osher S., Paragios N., (Eds.). Springer Verlag, 2003.
- [KWT88] KASS M., WITKIN A., TREZOPOULOS D.: Snakes: Active contour models. *Int. J. Comput. Vis.* 1 (1988), 321–331.
- [LOC94] LIU X.-D., OSHER S., CHAN T.: Weighted essentially nonoscillatory schemes. *J. Comput. Phys.* 115 (1994), 200–212.
- [MBW\*05] MUSETH K., BREEN D., WHITAKER R., MAUCH S., JOHNSON D.: Algorithms for interactive editing of level set models. *Computer Graphics Forum* 24, 4 (2005), 821–841.
- [MHS\*96] MANDERS E. M. M., HOEBE R., STRACKEE J., VOSSEPOEL A., ATEN J.: Largest contour segmentation: A tool for the localization of spots in confocal images. *Cytometry* 23 (1996), 15–21.
- [MS89] MUMFORD D., SHAH J.: Optimal approximation by piecewise smooth functions and associated variational problems. *Commun. Pure Applied Mathematics* 42 (1989), 577–685.
- [MT95] MCINERNEY T., TERZOPOULOS D.: Topologically adaptable snakes. In *ICCV* (1995), pp. 840–845.
- [NM06] NIELSEN M. B., MUSETH K.: Dynamic tubular grid: An efficient data structure and algorithms for high resolution level sets. *Journal of Scientific Computing* (February 2006), 1–39.
- [OF03] OSHER S., FEDKIW R.: *Level Set and Dynamic Implicit Surfaces*. Springer-Verlag New York Inc., 2003.
- [OS88] OSHER S., SETHIAN J.: Fronts propagating with curvature-dependent speed: Algorithms based on Hamilton-Jacobi formulations. *Journal of Computational Physics* 79 (1988), 12–49.
- [PCM03] PASCUCCI V., COLE-MCLAUGHLIN K.: Parallel computation of the topology of level sets. *Algorithmica* 38, 2 (2003), 249–268.
- [Set99] SETHIAN J. A.: *Level Set Methods and Fast Marching Methods*, second ed. Cambridge University Press, Cambridge, UK, 1999.
- [TFT04] TAKAHASHI S., FUJISHIRO I., TAKESHIMA Y.: Topological volume skeletonization and its application to transfer function design. *Graphical Models* 66, 1 (2004), 24–49.
- [TIS\*95] TAKAHASHI S., IKEDA T., SHINAGAWA Y., KUNII T. L., UEDA M.: Algorithms for extracting correct critical points and constructing topological graphs from discrete geographical elevation data. *Computer Graphics Forum* 14, 3 (1995), C–181–C–192.
- [vKvOB\*97] VAN KREVELD M., VAN OOSTRUM R., BAJAJ C. L., PASCUCCI V., SCHIKORE D. R.: Contour trees and small seed sets for isosurface traversal. In *Proceedings of the 13th ACM Symposium on Computational Geometry* (1997), pp. 212–220.
- [WDC\*] WEBER G., DILLARD S., CARR H., PASCUCCI V., HAMANN B.: Topology-controlled volume rendering. To appear in *IEEE Transactions on Visualization and Computer Graphics*.
- [WMW86] WYVILL G., MCPHEETERS C., WYVILL B.: Data structure for soft objects. *Visual Computer* 2 (1986), 227–234.
- [ZBB04] ZHANG X., BAJAJ C. L., BAKER N.: *Fast Matching of Volumetric Functions Using Multi-resolution Dual Contour Trees*. Tech. rep., Texas Institute for Computational and Applied Mathematics, Austin, Texas, 2004.
- [ZCMO96] ZHAO H.-K., CHAN T., MERRIMAN B., OSHER S.: A variational level set approach to multiphase motion. *Jrnl. of Comp. Phys.* 127 (1996), 179–195.
- [Zha04] ZHAO H.: Fast sweeping method for eikonal equations. *Mathematics of Computation* 74 (2004), 603–627.