Codimension-Two Geodesic Active Contours

SUBMISSION #352

Abstract

Curve evolution schemes for segmentation, implemented with level set methods, have become an important approach in computer vision. Previous work has modeled evolving contours which are curves in 2D or surfaces in 3D. A limitation inherent in the underlying mathematical theory was that the evolving manifold be a hypersurface. Our objective is to explore more recent mathematical work enabling the evolution of manifolds of higher co-dimension. In particular, we consider 1D curves in 3D (codimension two) for the application of automatically segmenting blood vessels in volumetric magnetic resonance angiography (MRA) images. Blood vessels are modeled as curves in 3D which can branch arbitrarily. We have thus developed the CURVES *system which segments MRA images by evolving an initial estimate toward the true structures in the image using the codimension two regularization force. This novel regularization force enables the segmentation of very thin structures which would not be possible with a traditional codimension one evolution. This paper describes CURVES and its theoretical foundations, then provides MRA segmentation results compared against segmentations obtained interactively by a neurosurgeon. Segmentations of bronchi in lung CT scans are also presented.*

1 Introduction

Curve-shortening flow is the evolution of a curve over time to minimize some distance metric. When this distance metric is based on image properties, it can be used for segmentation. The idea of *geodesic active contours* is to define the metric so that indicators of the object boundary, such as large intensity gradients, have a very small "distance" [5, 12, 3]. The minimization will attract the curve to such image areas, thereby segmenting the image, while preserving properties of the curve such as smoothness and connectivity. Geodesic active contours can also be viewed as a more mathematically sophisticated variant of classical snakes [11]. Further, they are implemented with level set methods [24] which are based on recent results in differential geometry [8, 6]. The method can be extended to evolve surfaces in 3D, for the segmentation of 3D imagery such as medical datasets [5, 26].

A limitation of the method was that it and its theoretical foundations applied only to hypersurfaces. More recent work in differential geometry, however, developed the equations necessary to evolve arbitrary dimensional manifolds in arbitrary dimensional space [1]. Subsequent work developed and analyzed a diffusion-generated motion scheme for codimension-two curves [22]. We have developed the first application of geodesic active contours in 3D, based on [1]. Our system, CURVES, evolves 1D curves in 3D images and has been applied to automatic segmentation of blood vessels in medical images [16]. This paper reviews CURVES and its underlying theory, then presents novel experiments and validation comparisons.

2 Vascular Segmentation in Medical Images

In neurological surgery the vasculature is of utmost importance. Direct visualization of images acquired with current imaging modalities, however, cannot provide a spatial representation of small vessels. These vessels, and their branches which show considerable variations, are most important in planning and performing neurosurgical procedures. In planning they provide information on where the lesion draws its blood supply, and where it drains. This is of particular interest in vascular malformations. During surgery the vessels serve as landmarks, and guidelines to the lesion. The more minute the information is, the more precise are navigation and localization of computer guided procedures.

For these reasons we consider the segmentation of volumetric vasculature images, such as the magnetic resonance angiography (MRA) image pictured in section 5, with a focus on segmenting the small vessels. Areas of flowing blood appear bright in this imaging modality. The images are displayed in *maximum intensity projection* in which the stack of slices is collapsed into a single image for viewing by performing a projection through the stack that assigns to each pixel in the projection the brightest voxel over all slices. The approach of simply thresholding the raw data is commonly used for segmentation but incorrectly labels bright noise regions as vessel and cannot recover

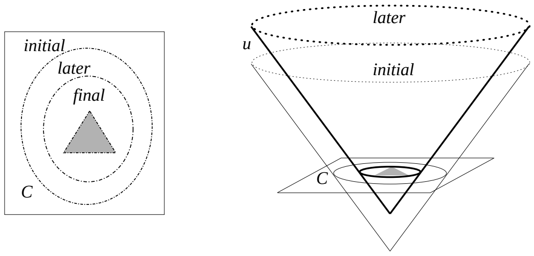

(a) (b)

Figure 1: Simple segmentation example: (a) Evolving curve. (b) Level set implementation of curve evolution.

very small vessels which may not appear connected in the volumetric image.

Several more sophisticated approaches have been explored. Multiscale approaches have involved convolving the image with gaussian filters at multiple scales and analyzing the eigenvalues of the hessian matrix at each voxel in the image to determine the local shape of the structures in the image [23, 19, 17, 14, 15]. Anisotropic diffusion has also been applied [13, 20]. Another multiscale approach based on medial axes uses the fact that the centerlines of the vessels appear brightest to detect these centerlines as intensity ridges of the image [2]. A differential geometry approach treats the 3D MRA image as a hypersurface of 4D space in which extrema of curvature correspond to vessel centerlines [21]. A statistical approach based on the expectation maximization (EM) algorithm has been proposed as a fully automatic MRA segmentation method [25, 7]. Deformable models have also been applied to 3D vascular segmentation [5, 18].

3 Background

The CURVES algorithm is an extension of geodesic active contours research, also using a level set implementation. The extension was enabled by theoretical work on level set methods for arbitrary co-dimensional manifolds.

3.1 Geodesic Active Contours

The task of finding the curve that best fits the object boundary is posed [5, 3, 12] as a minimization problem over all closed planar curves $C(p) : [0,1] \to \mathbb{R}^2$. The objective function is

$$\oint_0^1 g(|\nabla I(C(p))|)|C'(p)|dp$$

where $I : [0,a] \times [0,b] \to [0,\infty)$ is the image and $g : [0,\infty) \to \mathbb{R}^+$ is a strictly decreasing function such that $g(r) \to 0$ as $r \to \infty$, such as $g(|\nabla I|) = \frac{1}{1+|\nabla I|^2}$.

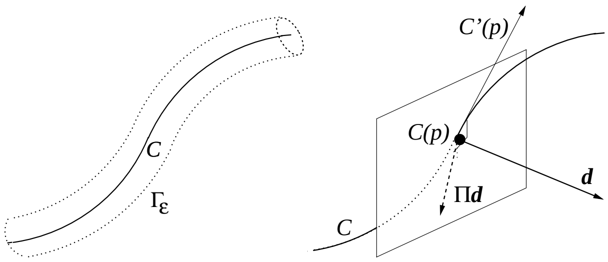

(a) (b)

Figure 2: Codimension-two curve: (a) Tubular isolevel set Γ_ε of C. (b) The tangent to C at p, the normal plane, the external vector $\vec{d}$, and its projection onto the normal plane.

To minimize this objective function by steepest descent, consider C to be a function of time t as well as spatial parameter p. The Euler-Lagrange equations yield the curve evolution equation

$$\vec{C}_t = g\kappa\vec{N} - (\nabla g \cdot \vec{N})\vec{N} \qquad (1)$$

where κ is the Euclidean curvature and $\vec{N}$ is the unit inward normal. In the absence of image gradients, this equation causes the curve to shrink according to its curvature; the presence of image gradients causes the curve to stop on the object boundary (Figure 1a).

3.2 Level Set Method for Hypersurfaces

Level set methods increase the dimensionality of the problem from the dimensionality of the evolving manifold to the dimensionality of the embedding space [24]. For the example of planar curves, instead of evolving the one-dimensional curve, the method evolves a two-dimensional surface. Let $u : \mathbb{R}^2 \to \mathbb{R}$ be the signed distance function to curve C as in Figure 1b; C is thus the zero level-set of u, and u is an implicit representation of C. Let C_0 be the initial curve. It is shown in [8, 6] that evolving C according to

$$\vec{C}_t = \beta\vec{N}$$

with initial condition $C(\cdot,0) = C_0(\cdot)$ for any function β, is equivalent to evolving u according to

$$u_t = \beta|\nabla u|$$

with initial condition $u(\cdot,0) = u_0(\cdot)$ and $u_0(C_0) = 0$ in the sense that the zero level set of u is identical to the evolving curve for all time. Choosing $\beta = g\kappa - (\nabla g \cdot \vec{N})$ as in Equation 1 gives the behavior illustrated in Figure 1b according to the update equation

$$u_t = g\kappa|\nabla u| + \nabla g \cdot \nabla u.$$

The extension to surfaces in 3D is straightforward and is called *minimal surfaces* [5]. The advantages of the level set representation are that it is intrinsic (independent of parameterization) and that it is topologically flexible since different topologies of C are represented by the constant topology of u.

3.3 Level Set Method for Curves in Higher Codimension

For the task of evolving one-dimensional curves in three-dimensional space, however, the above level set relation does not hold. It is applicable only to hypersurfaces, that is, surfaces whose *co-dimension* is one: the *co-dimension* of a manifold is the difference between the dimension of the evolving space and the dimension of the manifold. The examples of a planar curve and a 3D surface have co-dimension one, but 1D curves in three dimensions have co-dimension two. Intuition for why the level set method above no longer holds is that there is not an "inside" and an "outside" to a manifold with co-dimension larger than one, so one cannot create the embedding surface u in the same fashion as for planar curves; a distance function must be everywhere positive, and is thus singular on the curve itself. Recently, however, more general level set equations were found for curvature-based evolution [1]. It is these equations that motivated the development of CURVES, which uses image information to create the auxiliary vector field.

Let $v : \mathbb{R}^3 \to [0, \infty)$ be an auxiliary function whose zero level set is identically $C(p) : [0,1] \to \mathbb{R}^3$, that is smooth near C, and such that ∇v is non-zero outside C. For a nonzero vector $\mathbf{q} \in \mathbb{R}^\mathbf{n}$, define

$$P_\mathbf{q} = I - \frac{\mathbf{q}\mathbf{q}^\mathbf{T}}{|\ \mathbf{q}\ |^\mathbf{2}}$$

where I is the identity matrix as the projector onto the plane normal to $\mathbf{q}$. Further define $\lambda(\nabla v(x,t), \nabla^2 v(x,t))$ as the smaller nonzero eigenvalue of $P_{\nabla v} \nabla^2 v P_{\nabla v}$. The level set evolution equation is then [1]

$$v_t = \lambda(\nabla v(x,t), \nabla^2 v(x,t)).$$

That is, this evolution is equivalent to evolving C according to $\vec{C}_t = \kappa \vec{N}$ in the sense that C is the zero level set of v throughout the evolution. For intuition, let C be some curve in $\mathbb{R}^3$ and v the distance function to C. Consider then an isolevel set $\Gamma_\varepsilon = \{x | v(x) = \varepsilon\}$ of v where ε is small and positive, so Γ_ε is a thin tube around C (Figure 2a). The nonzero eigenvalues of $P_{\nabla v} \nabla^2 v P_{\nabla v}$ are equal to the principal curvatures of this tube. The larger principal curvature depends on ε while the smaller is related to the geometry of C. It is according to C that we want the evolution to proceed; thus, the smaller principal curvature is chosen.

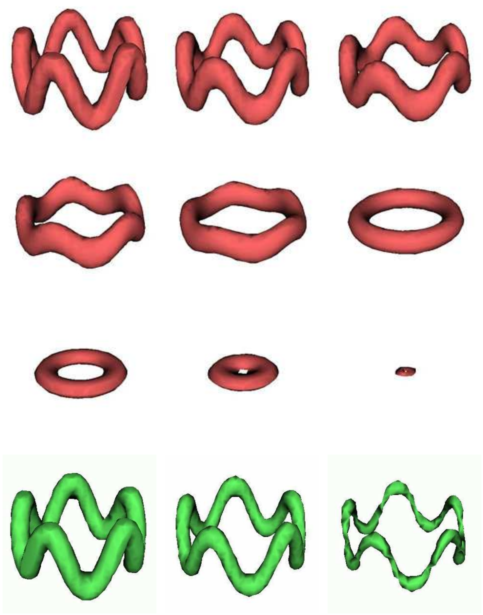

Figure 3: First three rows demonstrate the tubular object evolving to smooth the underlying curve, as in CURVES. Notice the bumps are first smoothed out until the shape approximates a torus, then the torus shrinks to a point. The final row demonstrates the shape evolving under codimension-one flow. The high curvatures curresponding to the small radius of the tube cause the shape to become thinner until it disappears without perceptibly affecting the geometry of the underlying curve.

The first rows of Figure 3 demonstrate the behavior of a shape undergoing this motion, where the smoothing force correspondes to the curvature of the underlying 1D curve. The final row then compares this behavior to that of traditional (codimension-one) mean curvature flow in which the regularization is based on the mean curvature of the surface.

Now assume there is an underlying vector field driving the evolution, so the desired evolution equation is

$$\vec{C}_t = \kappa \vec{N} - \Pi \vec{d},$$

where Π is the projection operator onto the normal space of C (which is a vector space of dimension 2) and $\vec{d}$ is a given vector field in $\mathbb{R}^3$, (Figure 2). The evolution equation for the embedding space then becomes [1]

$$v_t = \lambda(\nabla v, \nabla^2 v) + \nabla v \cdot \vec{d}.$$

4 CURVES

The evolution equation we use follows directly from an energy-minimization problem statement. Beyond that equation, several additional features of the program are incorporated for numerical and application-specific reasons.

4.1 Evolution Equation

For the case of 1D structures in 3D images, we wish to minimize

$$\oint_0^1 g(|\nabla I(C(p))|)|C'(p)|dp$$

where $C(p) : [0,1] \to \mathbb{R}^3$ is the 1D curve, $I : [0,a] \times [0,b] \times [0,c] \to [0,\infty)$ is the image, and $g : [0,\infty) \to \mathbb{R}^+$ is a strictly decreasing function such that $g(r) \to 0$ as $r \to \infty$ (analogous to [5]). For our current implementation, we use $g(r) = \exp(-r)$ because it works well in practice. The Euler-Lagrange equations give

$$\vec{C}_t = \kappa \vec{N} - \frac{g'}{g}\Pi(\mathbf{H}\frac{\nabla \mathbf{I}}{|\nabla \mathbf{I}|}),$$

where $\mathbf{H}$ is the Hessian of the intensity function. The auxiliary vector field in the above equation is thus

$$\vec{d} = \frac{g'}{g}\mathbf{H}\frac{\nabla \mathbf{I}}{|\nabla \mathbf{I}|},$$

so the equation for the embedding space is

$$v_t = \lambda(\nabla v(x,t), \nabla^2 v(x,t)) + \frac{g'}{g}\nabla v(x,t) \cdot \mathbf{H}\frac{\nabla \mathbf{I}}{|\nabla \mathbf{I}|}$$

4.2 Features

Initial experiments required that the evolving volume be a distance function to the underlying curve; however, it was not clear how to robustly extract the zero level set or even evolve those points since the distance function was singular exactly there. For this reason, we developed the *ε-Level Set Method* which defines a thin tube of radius ε around the initial curve, then evolves that tube instead of the curve. ε does not denote a fixed value here, but means only that the evolving shape is a "tubular" surface of some unspecified and variable nonzero width. Thus, we are now evolving surfaces similar to [5], but that follow the motion of the underlying curve so they do not regularize against the high curvatures found in thin cylindrical structures such as blood vessels and bronchi. In addition to being more robust, this method better captures the geometry of such structures, which have nonzero diameter.

Next, we modified the update equation for the MRA segmentation application. To control the trade-off between fitting the surface to the image data and enforcing the smoothness constraint on the surface, we incorporate an image weighting term ρ which is set by the user or is pre-set to a default value. Second, because vessels in MRA and bronchi in CT appear brighter than the background, we weight the image term by the cosine of the angle between the normal to the surface and the gradient in the image. This cosine is given by the dot product of the respective gradients of v and I, so the update equation becomes

$$v_t = \lambda(\nabla v, \nabla^2 v) + \rho(\nabla v \cdot \nabla I)\frac{g'}{g}\nabla v \cdot \mathbf{H}\frac{\nabla \mathbf{I}}{|\nabla \mathbf{I}|} \quad (2)$$

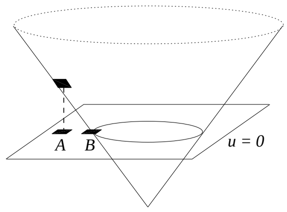

Figure 4: To evolve a point on the distance function, CURVES chooses image information from A instead of B.

A third aspect of our system is that we do not propagate the image information off the current object boundary to obtain the values near the boundary, as is customary in level set methods [5, 3, 12, 9]. Instead, we use directly the image information at each point in the image (Figure 4). Our choice has the advantage of enabling attraction to image gradients not on the current boundary, thereby reducing sensitivity to initialization, at the expense of requiring more frequent reinitializations of the distance function and losing the equivalence to the Lagrangian method (updating C directly instead of v) off the boundary.

In short, the CURVES system takes in a 3D image that contain thin curvilinear structures, such as an MRA image. An initial segmentation estimate is generated by thresholding the image. That estimate is used to generate an initial distance function, which is iteratively updated according to Equation 2. Convergence is detected automatically when volumetric change in the segmentation is very small over some number of iterations. Further detail was published in [16].

5 Results

We have run CURVES on over 20 medical datasets, primarily phase contrast magnetic resonance angiography (PC-MRA), of various resolutions and scanner types. We provide images of several representative segmentations.

We first show successive boundary estimates in a segmentation of a cerebral MRA image to demonstrate the behavior of the algorithm over time, until convergence is reached. We then show CURVES segmentations of more cerebral MRA images compared to those obtained with a manual segmentation technique used clinically at our institution. Finally, we illustrate the advantage of our system compared to codimension-one surface evolution with an experiment involving the segmentation of bronchi in a computed tomography (CT) image of lung. The codimension-one result was obtained using the mean curvature of the surface as the regularization force, as in previous level set segmentation schemes [5]; otherwise, all parameter settings were identical to those used in the CURVES experiment.

5.1 Example Evolution

Figure 5 illustrates the behavior of our system over time on a phase contrast MRA image of cerebral vessels. The initial surface is obtained by thresholding the raw dataset, then CURVES evolution produces the subsequent images.

5.2 Segmentations of Cerebral Vasculature

One specific practical motivation for our work is the use of surface models of cerebral vasculature as an aid in neurosurgical planning and procedure, especially in the context of the image-guided surgery program at our institution [10]. Currently the vessel models are obtained manually as follows. A neurosurgeon interactively chooses a threshold that is used to binarize the MRA dataset: all voxels brighter than that threshold are labeled as vessel, while all others are discarded. A "connectivity" program then partitions the set of labeled voxels into connected components. Each connected component appears in a distinct color on the user interface. The surgeon looks at individual slices and clicks on colored regions that correspond to vasculature. All connected components so chosen are stored as the final manual segmentation. The first drawback of this method is the expert user-interaction required, the second is that the thresholding step implies that all regions of image "noise" which adjoin vasculature are incorrectly labeled as vessel and small thin vessels which may appear broken or disconnected from larger structures will often be omitted. Thus, our goal is to reduce user interaction while increasing the ability to segment thin vessels.

Figure 6 shows CURVES segmentations (red) compared to segmentations acquired using the manual procedure just described (blue). The dataset shown here is PC-MRA acquired on a 1.5T scanner with voxel resolution of $1.171875 \times 1.171875 \times 0.8\text{mm}^3$ and a size of $256 \times 256 \times 84$ voxels. The raw dataset is shown in maximum intensity projection from three orthogonal viewpoints, as are the two segmentations. Notice that CURVES is able to capture more of the thin vessels than the manual procedure which is based on simple thresholding. Figure 7 shows CURVES results on three more datasets for which we also have manual segmentations. The first dataset was acquired using the same protocol as that in Figure 6, and the second and third were acquired on a 0.5T scanner with a size of $256 \times 256 \times 60$ voxels and with the same protocol otherwise.

5.3 Bronchi Segmentation, Comparison to Codimension-One

For comparison purposes, we have created a version of the CURVES program which uses the codimension-one regularization force. Experimental results have demonstrated that the codimension-two regularization force does indeed allow the segmentation of more thin structures than does the codimension-one force. Figure 8 shows the CURVES segmentation (red) of bronchi in a lung CT dataset compared to the codimension-one segmentation (green), for the same parameter settings.

Acknowledgments

Blinded for review process.

References

[1] Ambrosio, L. and Soner, H.M.: Level set approach to mean curvature flow in arbitrary codimension. Journal of Differential Geometry **43** (1996) 693–737

[2] Aylward, S., Pizer, S.M., Bullitt, E., and Eberly, D.: Intensity ridge and widths for 3D object segmentation and description. IEEE Proc. Workshop Mathematical Models Biomedical Image Analysis 96TB100056 (1996) 131–138

[3] Caselles, V., Catte, F., Coll, T. and Dibos, F.: A geometric model for active contours. Numerische Mathematik **66** (1993) 1–31

[4] Caselles, V., Morel, J.M., Sapiro, G. and Tannenbaum, A.: Introduction to the special issue on partial differential equations and geometry-driven diffusion in image processing and analysis. IEEE Transactions on Image Processing **7(3)** (1998) 269–273

[5] Caselles, V., Kimmel, R. and Sapiro, G.: Geodesic active contours. Int'l Journal of Computer Vision **22(1)** (1997) 61–79

[6] Chen, Y.G., Giga, Y. and Goto, S.: Uniqueness and existence of viscosity solutions of generalized mean curvature flow equations. Journal of Differential Geometry **33** (1991) 749–786

[7] Chung, A., and Noble, J.A.: Statistical 3D vessel segmentation using a Rician distribution. Proc. Int'l Conf. Medical Image Computing Computer-Assisted Intervention. Lect. Notes Comp. Sci. 1679 (1999) 82–89

[8] Evans, L.C. and Spruck, J.: Motion of level sets by mean curvature: I. Journal of Differential Geometry **33** (1991) 635–681

[9] Reference blinded for anonymity of authors.

[10] Reference blinded for anonymity of authors.

[11] Kass, M., Witkin, A., and Terzopoulos, D.: Snakes: active contour models. Int'l Journal Computer Vision **1(4)** (1988) 321–331

[12] Kichenassamy, A., Kumar, A., Olver, P., Tannenbaum, A. and Yezzi, A.: Gradient flows and geometric active contour models. In Proc. IEEE Int'l Conf. Computer Vision (1995) 810–815

[13] Krissian, K., Malandain, G., and Ayache, N.: Directional anisotropic diffusion applied to segmentation of vessels in 3D images. Proc. Int'l Conf. Scale-Space (1997) 345–348

[14] Krissian, K., Malandain, G., and Ayache, N.: Model based detection of tubular structures in 3D images. INRIA Technical Report **3736** (1999)

[15] Krissian, K., Malandain, G., Ayache, N., Vaillant, R., Trousset, Y.: Model based multiscale detection of 3D vessels. Proc. IEEE Int'l Conf. Computer Vision Pattern Recognition (1998) 722–727

[16] Reference blinded for anonymity of authors.

[17] Lorenz, C., Carlsen, I.-C., Buzug, T.M., Fassnacht, C., and Weese, J.: A multi-scale line filter with automatic scale selection based on the hessian matrix for medical image segmentation. Proc. First Int'l Conf. Scale-Space Theory in Computer Vision. Lect. Notes Comp. Sci. 1252 (1997) 152–163

[18] McInerney, T., and Terzopoulos, D.: Topology adaptive snakes. Medical Image Analysis, *In Press* (1999)

[19] Frangi, A.F., Niessen, W.J., Vincken, K.L., and Viergever, M.A.: Multiscale vessel enhancement filtering. Proc. Int'l Conf. Medical Image Computing Computer-Assisted Intervention. Lect. Notes Comp. Sci. 1496 (1998) 130–137

[20] Perona, P., and Malik, J.: Scale-space and edge detection using anisotropic diffusion. IEEE Trans. Pattern Analysis Machine Intelligence **12(7)** (1990) 629–639

[21] Prinet, V., Monga, O., Ge, C., Sheng, L.X., and Ma, S.D.: Thin network extraction in 3D images: application to medical angiograms. Int'l Conf. Pattern Recognition (1996) 3:386–390

[22] Ruuth, S.J., Merriman, B, Xin, J., and Osher, S.: Diffusion-generated motion by mean curvature for filaments. UCLA Computational and Applied Mathematics Report 98–47 (1998)

[23] Sato, Y., Nakajima, S., Shiraga, N., Atsumi, H., Yoshida, S., Koller, T., Gerig, G., and Kikinis, R.: Three-dimensional multi-scale line filter for segmentation and visualization of curvilinear structures in medical images. Medical Image Analysis **2(2)** (1998) 143–168

[24] Sethian, J.A.: Level Set Methods. Cambridge University Press (1996)

[25] Wilson, D., and Noble, J.A.: Segmentation of cerebral vessels and aneurysms from MR angiography data. Proc. Int'l Conf. Information Procession Medical Imaging. Lect. Notes Comp. Sci. 1230 (1997) 423–428

[26] Zeng, X., Staib, L.H., Schultz, R.T., and Duncan, J.S.: Segmentation and measurement of the cortex from 3D MR images. Proc. Int'l Conf. Medical Image Computing Computer-Assisted Intervention. Lect. Notes Comp. Sci. 1496 (1998) 519–530

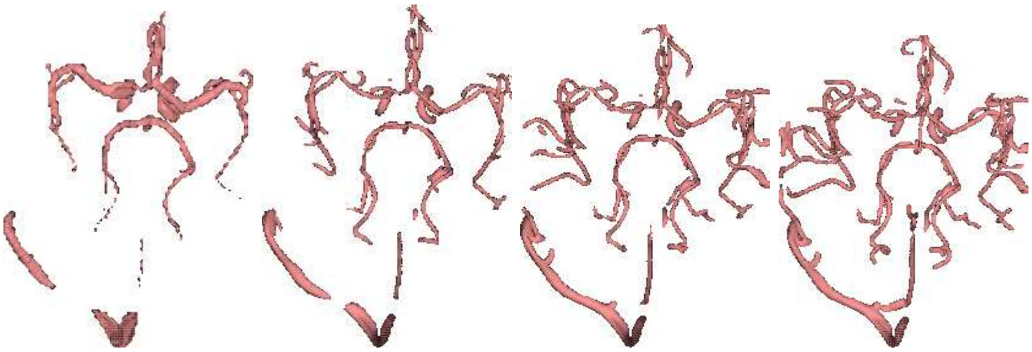

Figure 5: Surface evolution over time: initialization, followed by successive boundary estimates.

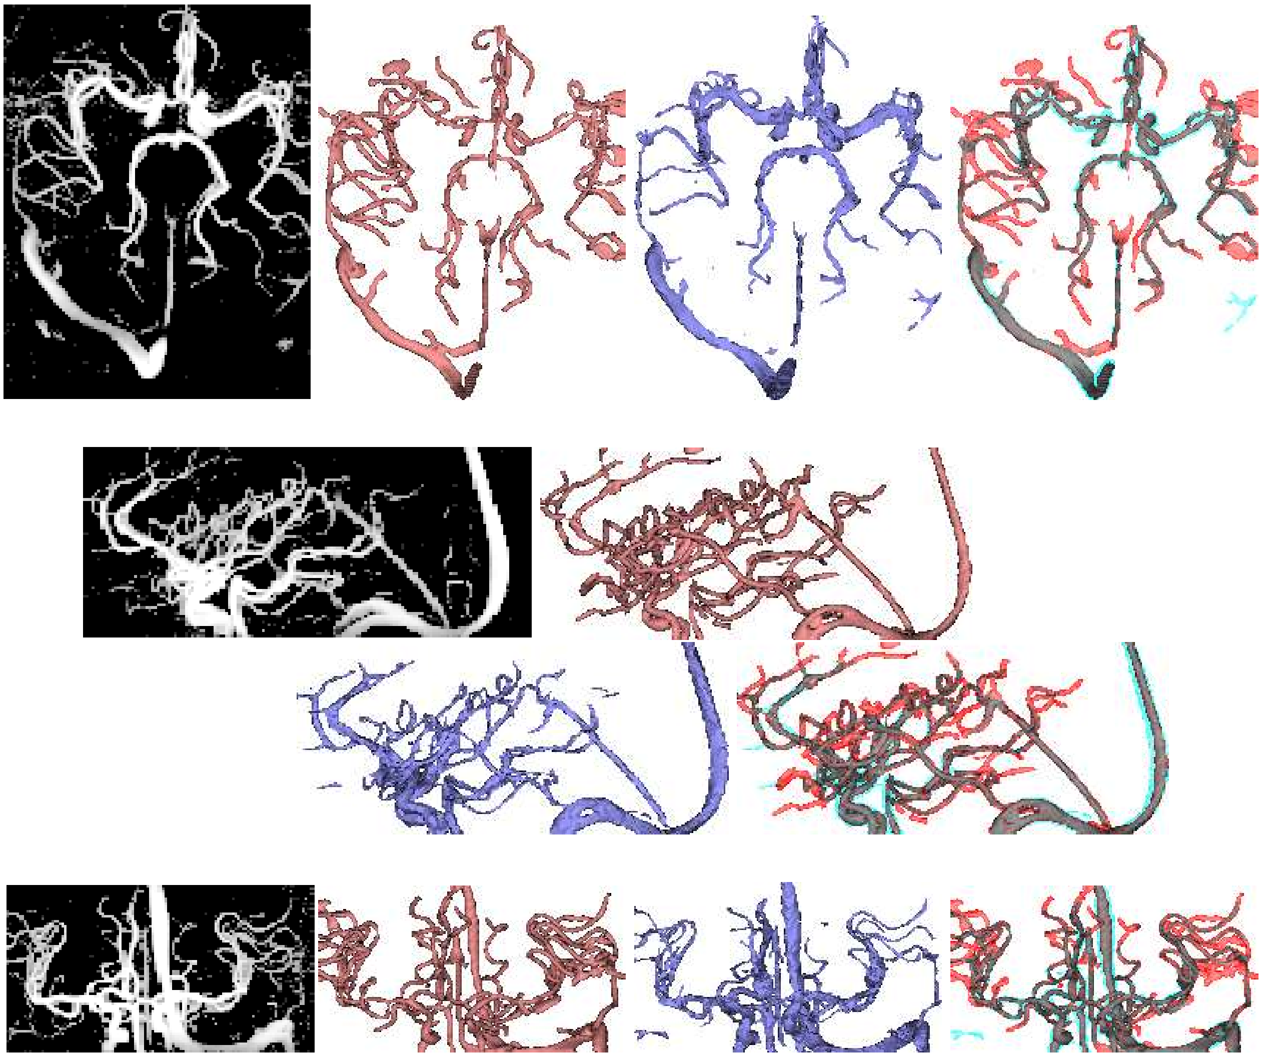

Figure 6: The same MRA dataset is shown from three orthogonal viewpoints. For each viewpoint, the maximum intensity projection of the raw data is shown, followed by the CURVES segmentation (red), the manual segmentation (blue), and a combination image showing the differences between the segmentations.

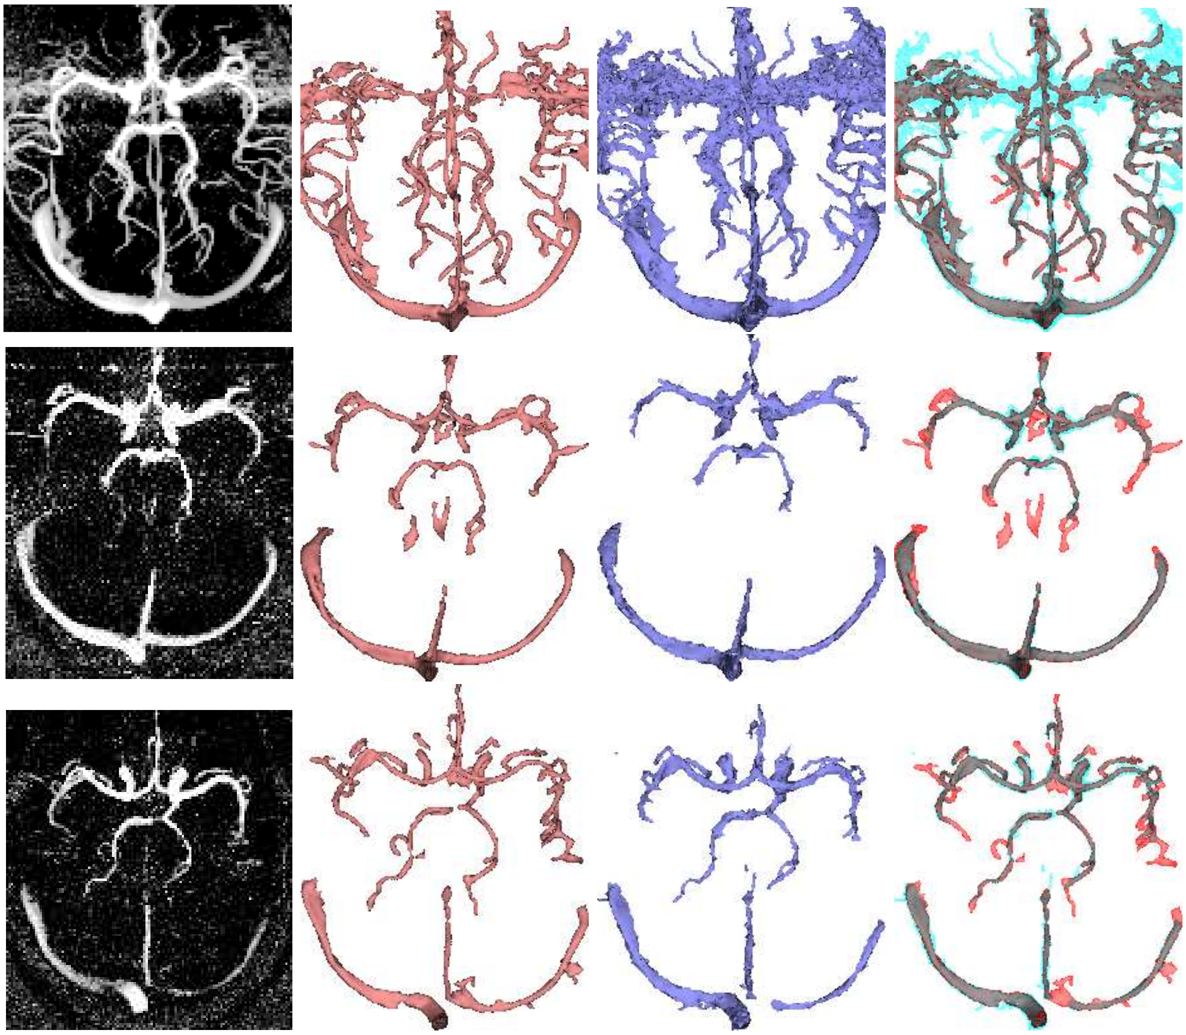

Figure 7: Each row shows a different dataset. Left to right: maximum intensity projection of raw data, CURVES segmentation, manual segmentation, combination image showing the differences between the segmentations.

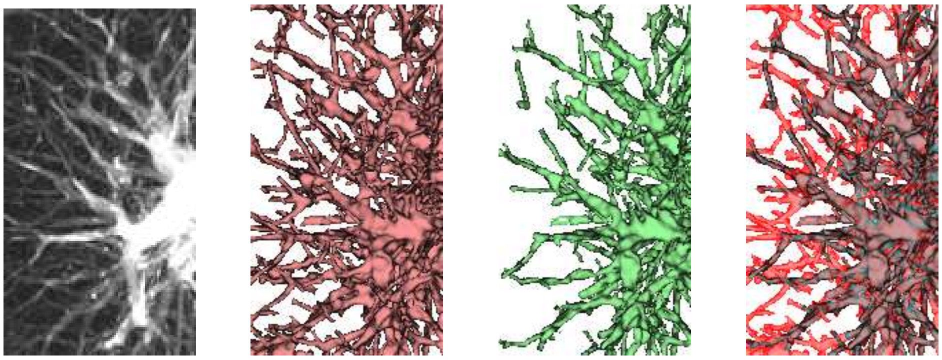

Figure 8: Segmentation of bronchi in CT lung scan: maximum intensity projection of CT dataset, CURVES segmentation, codimension-one segmentation, combination image showing the differences between the segmentations. Notice that CURVES is better able to capture thin bronchi than is the codimension-one method due to the new regularization force.