Medical Image Analysis (????) volume ??, number ??, pp 1–14

CURVES: Curve Evolution for Vessel Segmentation

Liana M. Lorigo[1*], Olivier D. Faugeras[1,2], W. Eric L. Grimson[1], Renaud Keriven[3], Ron Kikinis[4], Arya Nabavi[4] and Carl-Fredrik Westin[4]

[1]Artificial Intelligence Laboratory, Massachusetts Institute of Technology, Cambridge MA, USA
[2]INRIA, Sophia Antipolis, France
[3]Cermics, ENPC, France
[4]Harvard Medical School, Brigham & Women's Hospital, Boston MA, USA

Abstract
The vasculature is of utmost importance in neurosurgery. 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. During surgery the vessels serve as landmarks and guidelines to the lesion. The more minute the information is, the more precise the navigation and localization of computer guided procedures. Beyond neurosurgery and neurological study, vascular information is also crucial in cardiovascular surgery, diagnosis, and research.
This paper addresses the problem of automatic segmentation of complicated curvilinear structures in three-dimensional imagery, with the primary application of segmenting vasculature in magnetic resonance angiography (MRA) images. The method presented is based on recent curve and surface evolution work in the computer vision community which models the object boundary as a manifold that evolves iteratively to minimize an energy criterion. This energy criterion is based both on intensity values in the image and on local smoothness properties of the object boundary, which is the vessel wall in this application. In particular, the method handles curves evolving in 3D, in contrast with previous work that has dealt with curves in 2D and surfaces in 3D. Results are presented on cerebral and aortic MRA data as well as lung computed tomography (CT) data.

Keywords: volumetric vascular segmentation, deformable models

Received February 11, 2001; revised ??; accepted ??

1. Introduction

The vasculature is of utmost importance in neurosurgery and neurological study. Elaborate studies with a considerable x-ray exposure, such as multi-planar conventional angiography or spiral computed tomography (CT) with thin slices, have to be carried through to achieve an accurate assessment of the vasculature. Three-dimensional CT angiography and three-dimensional time-of flight magnetic resonance angiography (TOF-MRA) yield spatial information, but lack more subtle information. Furthermore, the three-dimensional CT needs a significant amount of contrast administration. All these studies cannot provide a spatial representation of small vessels. These vessels, and their branches which exhibit much variability, 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. The surgical interest is to differentiate between the feeding vessel and the transgressing vessel which needs to be preserved. In interventional neuroradiology this knowledge is utilized to selectively close the feeding vessel through the artery itself.

*Corresponding author
(e-mail: liana@ai.mit.edu)

During surgery the vessels serve as landmarks and guidelines to the lesion. The more minute the information is, the more precise the navigation and localization of computer guided procedures. Present representations do not yield this kind of information. A more precise spatial representation of this complex anatomic structure is needed.

For these reasons, we would like an automated analysis tool to interpret these images, with the capability to obtain as much of the fine detail as possible. For this study, we consider the segmentation of volumetric vasculature images, such as the magnetic resonance angiography (MRA) images pictured in section 5, with a focus on segmenting the small vessels. Areas of flowing blood can appear bright in this imaging modality. The MRA 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 very small vessels which may not appear connected in the volumetric image. Instead, we have developed the *CURVES* system which models the vessels as three-dimensional curves with arbitrary branching and uses an active contours approach to segment these curves from the medical image (Lorigo *et al.* 1999).

This paper is organized as follows. Before describing CURVES, we review previous approaches to vessel segmentation. We then discuss CURVES' theoretical and experimental foundations, followed by the system description itself. The paper concludes with experimental results on MRA and CT datasets and a brief summary.

2. Other Approaches

Multiscale filtering has been proposed for the segmentation of curvilinear structures in three-dimensional medical images (Sato *et al.* 1998; Krissian *et al.* 1999; Krissian *et al.* 1998; Frangi *et al.* 1998; Lorenz *et al.* 1997). The primary application addressed is the segmentation of vasculature in MRA images. This method involves 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. For example, if the voxel corresponds to a linear structure such as a bright vessel in an MRA image, the eigenvalues would be different than if the voxel corresponds to a planar structure, speckle noise, or no structure. Some methods use the output of the multiscale filter directly to define a new image in which curvilinear structures are brightened and bright voxels corresponding to speckle noise and planar structures such as skin are darkened (Sato *et al.* 1998; Frangi *et al.* 1998; Lorenz *et al.* 1997). This enhanced image is visualized directly (Frangi *et al.* 1998), thresholded (Sato *et al.* 1998), or segmented using an active contour method (Lorenz *et al.* 1997). Other methods use the eigenvalues so obtained to define a candidate set of voxels which could correspond to the centerlines of vessels (Krissian *et al.* 1999; Krissian *et al.* 1998). Multiscale response functions are evaluated at each of these voxels to determine the likelihood that the voxel is a vessel of various diameters. The maximal response over all choices of diameters (scales) is retained at each voxel, and a surface model of the entire vascular structure is reconstructed from knowledge of centerlines and diameters. A final method which obtains segmentations by thresholding a filtered MRA image uses anisotropic diffusion to remove noise without removing small vessels (Krissian *et al.* 1997; Perona, Malik 1990; Catte *et al.* 1992).

A different multiscale approach based on medial axes uses that assumption that the centerlines of the vessels often appear brightest to detect these centerlines as intensity ridges of the image (Aylward *et al.* 1996). The width of a vessel is then determined by a multiscale response function. This algorithm has been used in conjunction with two-dimensional/three-dimensional registration to incorporate information from a pair of x-ray angiograms (Bullitt *et al.* 1997). Other work has taken a differential geometry approach in which the volumetric MRA image is treated as a hypersurface of 4D space whose extrema of curvature correspond to vessel centerlines (Prinet *et al.* 1996), and a statistical approach in which Gaussian (Wilson, Noble 1997) or Rician (Chung, Noble 1999) intensity distributions are assumed for background and for vessel intensities and the expectation maximization (EM) algorithm is applied to find appropriate thresholds for classification.

Deformable model approaches have been applied to three-dimensional vascular segmentation as well. In such methods, an initial boundary estimate is deformed iteratively to optimize an energy function which depends both on image information and on the smoothness of the surface. One such algorithm is called *minimal surfaces*. Our work follows closely from this approach, so its discussion is deferred until section 3. Moreover, we compare results obtained with our algorithm to results obtained with a variant of the minimal surfaces approach in section 5.4. Another volumetric deformable model approach is *t-surfaces*, or *topologically adaptive surfaces* (McInerney, 1999), which are an extension of classical snakes (Kass *et al.* 1988) that can handle changing topologies and are independent of parameterization of the evolving surface model. An overview of the use of deformable models in medical image analysis is found in (McInerney, Terzopoulos 1996). The CURVES algorithm belongs to this class of approaches; the next section reviews the work from mathematics and computer vision on which it is founded.

3. Background

Curve evolution schemes for segmentation, implemented with level set methods, have become an important approach in computer vision (Caselles *et al.* 1997; Kichenassamy *et al.* 1995; Malladi *et al.* 1994; Sethian 1996). This approach is an extension of classical active contour models (Kass *et al.* 1988), and uses partial differential equations to control the evolution of an initial boundary estimate toward the true object boundary. An overview of the superset of techniques using related partial differential equations can be found in (Caselles *et al.* 1998). The fundamental concepts from mathematics from which these schemes derive were explored several years earlier when smooth closed curves in 2D were proven to shrink to a point under mean curvature motion (Gage, Hamilton 1986; Grayson 1987), which is the evolution of a manifold over time defined so that the temporal derivative of the manifold is equal to the mean curvature vector (the normal vector scaled by the mean curvature). Evans and Spruck and Chen, Giga, and Goto independently framed mean curvature flow of any hypersurface as a level set problem and proved existence, uniqueness, and stability of viscosity solutions (Chen *et al.* 1991; Evans, Spruck 1991). For application to image segmentation, a vector field was induced on the embedding space so that the evolution could be controlled by an image gradient field or other image data. This model is called *geodesic active contours*. The same results of existence, uniqueness, and stability of viscosity solutions were obtained for the modified evolution equations for the case of planar curves, and experiments on real-world images demonstrated the effectiveness of the approach (Caselles *et al.* 1997; Kichenassamy *et al.* 1995; Caselles *et al.* 1993).

When extended to three dimensions, curves evolving in the plane became surfaces evolving in space, called *minimal surfaces* (Caselles *et al.* 1997). Although the theorem on planar curves shrinking to a point could not be extended to the case of surfaces evolving in three dimensions, the existence, uniqueness, and stability results of the level set formalism held analogously to the 2D case. Thus the method was feasible for evolving both curves in two dimensions and surfaces in three dimensions. Beyond elegant mathematics, success on real-world data sets established the method as an important segmentation tool in both domains. Recent work has extended and applied the method to medical image segmentation (Zeng *et al.* 1999; Zeng *et al.* 1998). A related method called *bubbles* has also been applied to medical imagery (Tek, Kimia 1997). One fundamental limitation of these schemes has been that they describe only the flow of hypersurfaces, i.e., surfaces of codimension one: the *codimension* of a manifold is the difference between the dimension of the evolving space and the dimension of the manifold. It is also the number of equations that are necessary to define the manifold.

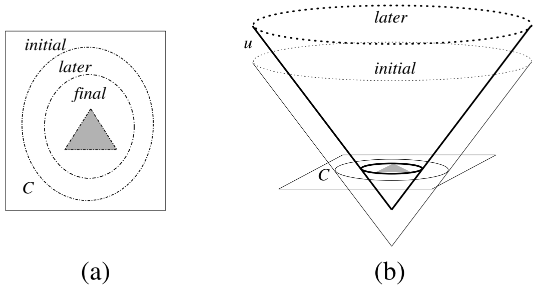

Figure 1. Simple segmentation example: (a) Evolving curve. (b) Level set implementation of curve evolution.

Regarding surfaces of codimension larger than one, Altschuler and Grayson studied the problem of curve-shortening flow for three-dimensional curves (Altschuler, Grayson 1992), and Ambrosio and Soner generalized the level set technique to arbitrary manifolds in arbitrary dimension. They provided the analogous results and extended their level set evolution equation to account for an additional vector field induced on the space (Ambrosio, Soner 1996). Subsequent work developed and analyzed a diffusion-generated motion scheme for codimension-two curves (Ruuth *et al.* 1998). We have developed the first implementation of geodesic active contours in three dimensions (Lorigo *et al.* 1999), based on Ambrosio and Soner's work. Our system, CURVES, uses these techniques for automatic segmentation of blood vessels in MRA images. The dimension of the manifold is one, and its codimension is two. The CURVES algorithm is an extension of geodesic active contours research, also using a level set implementation. We next review the basic geodesic active contour model, the level set technique employed for implementation, and this more recent extension of the level set method to higher codimensional curves.

3.1. Geodesic Active Contours

The task of finding the curve that best fits the object boundary is posed as a minimization problem over all closed planar curves $C(p):[0,1]\to\mathbb{R}^2$ (Caselles *et al.* 1997; Kichenassamy *et al.* 1995; Caselles *et al.* 1993). The objective function is

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

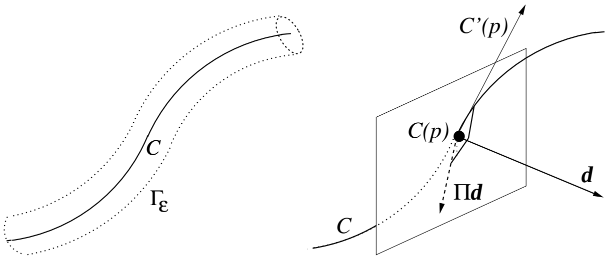

(a) (b)

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.

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 $r \to \infty$, e.g., $g(|\nabla I|) = \frac{1}{1+|\nabla I|^2}$.

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

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

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 (Sethian 1996). 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; it is for example positive outside the region determined by the curve and positive inside. 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 (Chen *et al.* 1991; Evans, Spruck 1991) 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 three dimensions is straightforward and is called *minimal surfaces* (Caselles *et al.* 1997). 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 codimension is one. The examples of a planar curve and a three-dimensional surface have codimension one, but space curves (curves in three-dimensions) have codimension 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 codimension 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 thus its gradient is singular on the curve itself. The discovery of more general level set equations for curvature-based evolution (Ambrosio, Soner 1996), however, motivated the development of CURVES, which uses image information to create the auxiliary vector field used to evolve 1D curves.

Let $C(p) : [0,1] \to \mathbb{R}^3$ be some curve and $v : \mathbb{R}^3 \to [0,\infty)$ be an auxiliary function whose zero level set is identically C, 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

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

as the projector onto the plane normal to $\mathbf{q}$ (I is the identity matrix) . 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 for mean curvature flow $\vec{C}_t = \kappa\vec{N}$ is then (Ambrosio, Soner 1996)

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

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 v be 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

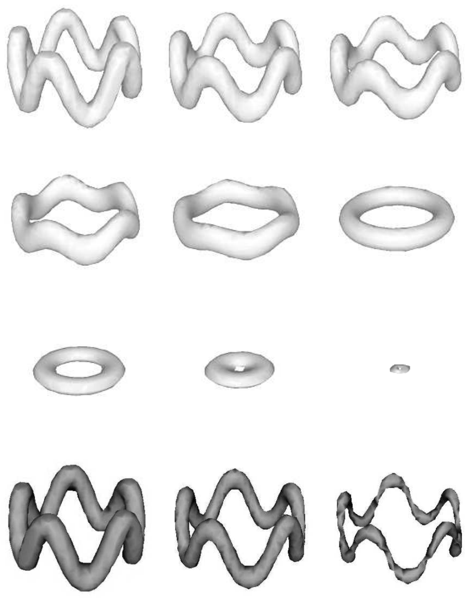

Figure 3. First three rows demonstrate the tubular object evolving to smooth the underlying curve, as in CURVES. Notice the bumps are first smoothed out until the shape approximates a torus, then the torus shrinks to a point. Fourth row demonstrates the shape evolving under codimension-one flow. The high curvatures corresponding 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.

(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.

The first rows of Figure 3 demonstrate the behavior of a shape undergoing this motion, where the smoothing force corresponds 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 (Ambrosio, Soner 1996)

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

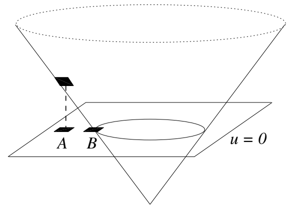

Figure 4. To evolve a point on the distance function, CURVES chooses image information from A instead of B.

4. CURVES

The curve evolution equation we use follows directly from an energy-minimization problem statement. When embedding that curve evolution in the evolution of a volume, we make a non-traditional choice for incorporating the image information. Beyond that equation, several additional features of the program are incorporated for numerical and application-specific reasons, as described below.

4.1. Evolution Equation

For the case of one-dimensional structures in three-dimensional 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 $r \to \infty$. For our current implementation, we use $g(r) = \exp(-r)$ because it works well in practice. By computing the Euler-Lagrange equations, we find that the curve evolution equation is

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

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. Locality of Image Information

For geodesic snakes of any dimensionality and codimensionality one must compute some curvature and some external image-related term at each point on the higher-dimensional manifold (the surface in the case of a planar curve, the volume in the case of a space curve). For each of these terms, one can use the values defined at the particular location or those defined at the closest point on the zero level set (Figure 4). Traditional level set segmentation methods use the image term from the closest point on the level set, but compute the curvature term locally (Caselles *et al.* 1997; Kichenassamy *et al.* 1995; Caselles *et al.* 1993). The reason is that the curvature term is defined locally, and the level-set-equivalence relation says that indeed one should use that local curvature. The image-term, conversely, is not defined locally if one regards the problem as evolving the curve directly. One must, then, "invent" an image term at those points off the zero level set. The best choice is to choose the image term at the nearest point on the zero level set. This choice keeps the evolving "distance function" as close to a true distance function as possible without modifying the curvature term. Alternative formulations keep the evolving manifold a distance function throughout the evolution (Gomes, Faugeras 2000; Zhao *et al.* 1996) using image or other information from the object boundary as well as curvature information from the boundary only; no local information is used at all.

The CURVES method, however, uses the image term at each location on the higher dimensional manifold instead of propagating the image data off the current zero level set. This choice was made to enable the evolving surface to be attracted to edge gradients that are not on the current surface. For example, if there are two neighboring tubes in the image and the curve or surface is initialized near one, CURVES can capture the other tube; the traditional method cannot. However, this also means that the CURVES method is not equivalent to explicit Lagrangian evolution, which would not find the second tube. The reason that neither an explicit evolution nor a tradition level set evolution would find the second tube is that they are stopped by the local minimum found at the outline of the single tube. CURVES is thus less sensitive to initialization than previous level set methods are.

This choice also has implications in the need to reinitialize the evolving higher-dimensional manifold to be a distance function. Even in the absence of an image force, all of the level sets are evolving toward the same local minima in traditional methods; they thus become increasingly close together so the manifold is no longer a distance function (Gomes, Faugeras 2000). In the case of CURVES, the image force is a more severe force that invalidates the distance function constraint. It follows that CURVES requires far more reinitializations than the traditional method. However, if we wish the image information off the zero level set to affect the evolution, we cannot reinitialize too frequently. For example, reinitializing after every step in the evolution is equivalent to using only the image information on the zero level set, since the reinitialization maintains only those values, updating all other values to be their distance to the zero level set.

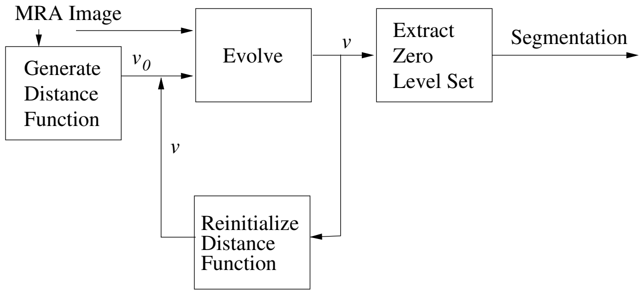

Figure 5. Overview of segmentation algorithm. v is the evolving volume whose zero level set is the current segmentation estimate. An initial volume v_0 is generated and passed to the "Evolve" routine along with the raw image data, where it is evolved according to the partial differential equation derived from the energy minimization formulation. Periodically, v is reinitialized to be a distance function to its zero level set. At convergence or when desired, the zero level set is extracted from v for visualization of the segmentation.

4.3. System Details

A flowchart of the CURVES system is shown in Figure 5. v is the evolving volume whose zero level set is the current segmentation estimate. An initial volume v_0 is generated from the image data and is passed with that data to the main body of the system which evolves v_0 (v) iteratively, according to the partial differential equation given in Equation 2 below. Periodically, v is reinitialized to be a distance function to its zero level set. At convergence or when desired, the zero level set is extracted from v for visualization of the segmentation.

This section discusses issues that have arisen in converting the theory above to practice. 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 gradient of the distance function was singular exactly there. Moreover, the projection operator $P_{\mathbf{q}}$ is defined only for non-zero vectors q, so the method is undefined at $\nabla v = \vec{0}$, which is the curve itself, and is numerically unstable near the curve. 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. We stress that this is an approximation to evolving the underlying curve but is not equivalent. If we were to constrain the width of the tube to remain constant along the tube, it would

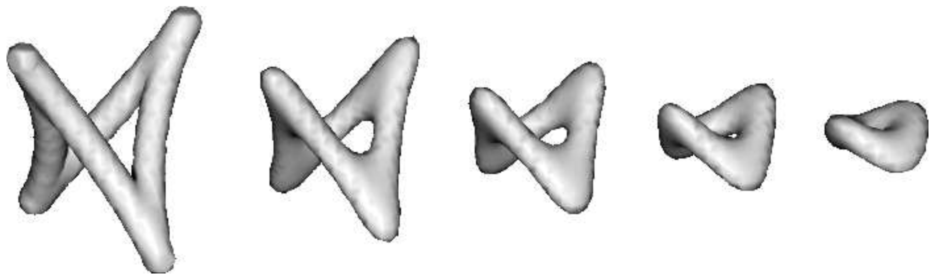

Figure 6. A tubular shape with corners under codimension-two flow.

be equivalent; however, allowing the image to attract local surface areas independently causes the width to vary, so the tube is no longer as isolevel set of the distance function to its centerline. Thus, we are now evolving surfaces similar to minimal surfaces (Caselles *et al.* 1997), 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.

We stress that this technique is an approximation to evolving the underlying curve but is not equivalent. If we were to constrain the width of the tube to remain constant along the tube, it would be equivalent; however, allowing the image to attract local surface areas independently causes the width to vary, so the tube is no longer as isolevel set of the distance function to its centerline.

The next implementation details relate to the volume evolution equation. 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. In practice, since our experimental datasets vary widely in dynamic range and noise content, segmentations are performed separately for a few different values of ρ, and the best segmentation is selected; however, if the system were routinely run on datasets acquired with the same parameters, the correct value would be known *a priori*. 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

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

The last comment on the update equation deals with the λ term. For computational efficiency and because of numerical instability of the gradient computations near $\nabla v = \vec{0}$, we remark that the level sets of the function v flow in the direction of the normal with velocity equal to the sum of their smaller principal curvature and the dot product of ∇v with the image-based vector field $\vec{d}$. Therefore, we compute the smaller principal curvature λ directly from v instead of as an eigenvalue of $P_{\nabla v}\nabla^2 v P_{\nabla v}$.

The initial surface (and thereby the initial volume v_0) is usually generated by thresholding the MRA dataset. However, the method does not require that the initial surface be near the target surface but may use any initial surface, so arbitrary tubular surfaces have been used to explore the behavior of the algorithm. Before running CURVES the image dataset is smoothed by a small isotropic Gaussian since the algorithm inherently requires some smoothness of gradients. The sigma normally used is 0.75mm in each dimension, where a typical dataset has voxel dimensions of $.9375 \times .9375 \times 1.5\text{mm}^3$.

Instead of evolving the entire volume, we evolve only the portion of the volume within a narrow band of the zero level set (the current surface). Normally, we set the band to include voxels that are 4 to 6 voxels away from the surface. This aspect of the implementation does not have the same meaning as "banding" (Adalsteinsson, Sethian 1995; Malladi *et al.* 1994; Chopp 1993) in previous geodesic active contour methods where the image data on the zero level set is propagated throughout the band. We simply mean that only those points are evolved. Note that, unlike these other methods, CURVES is sensitive to the width chosen for the band since image values therein are indeed used in the evolution.

Further, the points in the band are periodically reinitialized to be a distance function: the zero level set S is extracted, then the value at each point is set to be its distance to S. For our implementation, this reinitialization is itself a level set method. To obtain the positive distances, the surface is propagated outward at constant speed of 1, and the distance at each point is determined to be the time at which the surface crossed that point. A second step propagates the surface

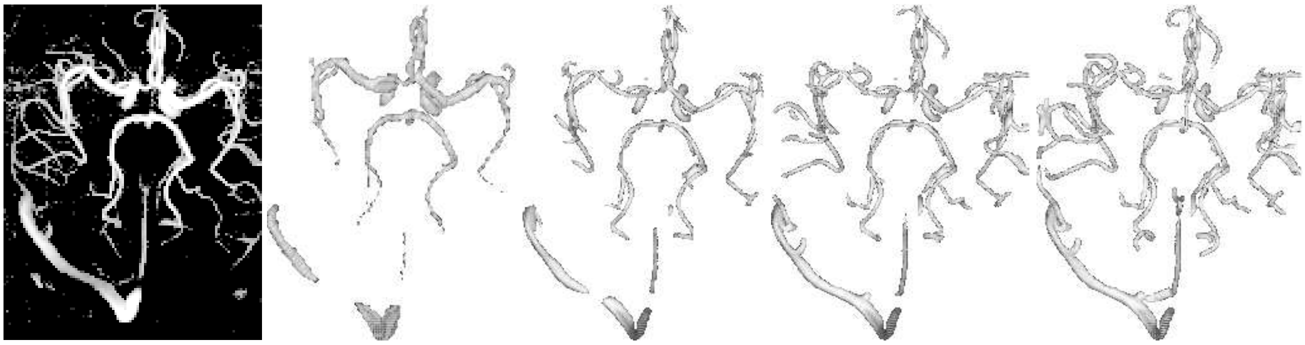

Figure 7. Surface evolution over time: maximum intensity projection of raw data, initialization, then successive boundary estimates.

inward to obtain the negative distances analogously. We reinitialize the distance function every 3 to 5 steps; this is much more frequently than previous level set techniques for reasons discussed above.

Convergence of the algorithm is detected when the volume of the segmented region changes less than some specified percentage of total volume, across a specified number of iterations. The user may then choose to omit from the resultant segmentation all structures whose volume is less than some threshold. This step can remove "noise" that may have been incorrectly segmented or can enable the user to focus only on the largest connected structures.

A final attribute of CURVES is the inherent capability to estimate vessel radii directly from the volumetric distance function v. Of the two surface curvatures, the smaller is used in the segmentation procedure, but the larger curvature can also be useful as it corresponds to the radii of the vessels. We demonstrate this observation by color-coding the segmentation result according to the larger curvature at each point on the surface in section 5.2.

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. After an illustration on synthetic data, we 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. The next example illustrates performance on an aorta dataset and the capability to estimate vessel radii. 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.

Qualitative comparisons are shown only due to the difficulty of obtaining ground truth segmentations for datasets of this level of complexity. Even the manually-obtained segmentations to which we compare CURVES cerebral vasculature segmentations cannot be considered "ground truth" since many vessels are not obtained and bright areas not corresponding to vessel are included in some cases. The images show the thin structures that CURVES obtains beyond those obtained by the manual method, but there is not currently a quantitative measure to evaluate the segmentations in these regions.

5.1. Example Evolutions

To illustrate the codimension-two regularization force applied to tubular shapes, we show the evolution of two synthetic shapes using regularization only, without an image force. One shape was shown in the first rows of Figure 3 and was discussed there. The second is shown in Figure 6. This shape has sharp corners, and we see that although the one-dimensional centerline of the shape is singular at those corners, the evolution simply uses the smaller principal curvature of the surface at these points. This procedure has the advantage of enabling the evolution to proceed in a natural way to smooth out these corners and the disadvantage of causing the tube to become fatter at these corners over the course of the evolution. The fattening is caused by the difference in the curvature estimates obtained on the inside and the outside of the corners. This disadvantage exists for all tubular shapes that have non-zero underlying curvature, but is especially prominent for sharp corners because the interior curvature estimate is very high in those cases. As expected, the tube shrinks according to its underlying centerline, modulo this non-constant change in width.

The next example shows the use of the image force derived from MRA data. In particular, Figure 7 illustrates the behavior of our system over time on a PC-MRA image of cerebral vessels. The initial surface is obtained by thresholding the raw dataset, then CURVES evolution produces the subsequent images. The results are discussed below; these images indicate the temporal change in the tubular structure.

5.2. Aorta Segmentations

Figure 8 shows the segmentation of a contrast-enhanced MRA image of an aorta. This image was acquired on a Siemens scanner at New York University, with voxel resolution of $1.75 \times 1.75 \times 2.39\text{mm}^3$ and a size of $256 \times 256 \times 45$ voxels. The segmentation is shown from two orthogonal viewpoints. For each viewpoint, the maximum intensity projection of the raw data is shown first, followed by the original segmentation and that segmentation color-coded by local radii estimates. The colormap shows the widest vessels in red, intermediate vessels in green, and the thinnest vessels in blue. Recall that cylinders are not fit globally, but only local curvature properties are used, so the color can vary between adjacent regions where the structure is not perfectly cylindrical.

5.3. Cerebral Vasculature, Comparison to Manual

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 (Grimson *et al.* 1996). 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 9 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 without contrast agent, 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 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 illustrating the differences between the segmentations. Notice that CURVES is able to capture much more of the thin vessels than is the manual procedure which is based on simple thresholding. One negative aspect of CURVES performance on this example is that some large vessels such as the middle cerebral arteries and the superior sagittal sinus appear too thin in the CURVES segmentation. This artifact occurs because CURVES places the vessel boundary at the location of sharpest intensity gradient. If the vessel intensity profile in MRA is assumed to be Gaussian, then the true vessel wall should be placed farther out along the tails of the Gaussian than at the points of sharpest gradients. This problem does not occur in CT data below which does not have a Gaussian profile, and future work will explore the modification of the objective function dependent on the imaging modality used.

Figure 10 shows CURVES results on three more datasets for which we also have manual segmentations. Here, each row shows a different dataset. The first dataset was acquired using the same protocol as that in Figure 9, 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. Each row shows, left to right, the maximum intensity projection of raw data, the CURVES segmentation, the manual segmentation, and a combination image showing the differences between the segmentations. The first row shows an MRA image containing considerable pulsatile flow artifacts which appears as a bright horizontal area surrounding the middle cerebral arteries. For this example, the thresholding-based manual method must include much of this "noise" in order to also obtain the thin vessels; since CURVES depends on intensity gradients it is better able to distinguish those arteries from the surrounding region, without losing the small vessels.

5.4. Bronchi, Comparison to Codimension-One

For comparison purposes, we have created a version of the CURVES program which uses the codimension-one regularization force, the mean curvature of the surface, as in previous level set segmentation schemes (Caselles *et al.* 1997); otherwise, all parameter settings were identical to those used in the CURVES experiment. Figure 11 shows the CURVES segmentation of bronchi in a lung CT dataset compared to the codimension-one segmentation, for the same parameter settings. The dataset had voxel resolution of $.64 \times .64 \times 1\text{mm}^3$ and a size of $512 \times 512 \times 224$ voxels. In the figure, the first image is the maximum intensity projection of a sub-block of the CT data, the second and third are the CURVES and codimension-one segmentations respectively, and the fourth is a combination image of the two segmentations. Notice that the codimension-two regularization force in CURVES does indeed allow the segmentation of more thin structures than does the codimension-one force. This affect is intuitive because the codimension-one algorithm incorporates a smoothness constraint which acts to prevent high curvatures anywhere on the resultant surface, which is inappropriate for the segmentation of thin tubular structures which must have high curvatures corresponding to their small radii. Conversely, the regularization force in CURVES allows this high curvature, regularizing against only the curvature of the underlying one-dimensional curve.

We display this comparison for pulmonary CT images instead of for the cerebral MRA images because the difference between the performance of the two algorithms is more dramatic for the pulmonary images. Since the difference in the algorithms is only in the regularization term, and the weight of the regularization term versus the weight of the image-related term is a parameter of the algorithm, the results are more different when the regularization term is more heavily weighted. The cerebral images shown appeared to contain less imaging artifacts than did the pulmonary image, so the image-force was set higher for the cerebral images, thus lessening the effect of changing the regularization force to be codimension-two. Regarding the sensitivity of the CURVES algorithm to this parameter, in general, similar settings are appropriate for images obtained with common settings. That is, appropriate defaults are possible for a particular application, but the user must re-determine the settings if the image modality or acquisition parameters are changed.

6. Summary

This paper has presented a novel method for the segmentation of curvilinear structures in volumetric medical imagery. Its primary application has been the segmentation of blood vessels in MRA data, and it has also been applied to the segmentation of bronchi in lung CT data. The method itself is an extension of geodesic active contours and minimal surfaces, with the distinction that its regularizing force derives from an underlying one-dimensional curve in three dimensions, which can be considered intuitively as the centerline of the tubular structures. Experimental results have been shown for an aorta MRA dataset, for several cerebral MRA datasets, and for a lung CT dataset. Manual segmentations were used for validation for the cerebral MRA datasets and a codimension-one segmentation algorithm was run on the lung CT dataset for comparison.

In summary, our goal is to replace tedious manual outlining of small vessels by an automated algorithm. The comparisons shown herein show feasibility for the brain vessels in these cases. The automatically obtained vessel models are at least as detailed as the manually editted ones. While MIPs are usually sufficient for purely diagnostic purposes, they do not usually provide the spatial differentiation required for navigation in the vicinity of tumors. In such scenarios, we routinely use 3D surface models in our research. One of the reasons that these models are not routinely used in clinical work is the amount of time required for their preparation. This issue is at the heart of the research presented in this paper.

Acknowledgments

Medical datasets pictured provided by Surgical Planning Laboratory, Brigham & Women's Hospital. L. Lorigo was funded by NSF Contract IIS-9610249, NSF Contract DMS-9872228, and NSF ERC (Johns Hopkins University agreement) 8810-274. A. Nabavi was funded by a grant from the DFG (NA356/1-1). R. Kikinis and C.-F. Westin were partially funded by NIH grants P41-RR13218, P01-CA67165, R01-RR11747 and a grant from CIMIT. R. Kikinis is also funded by NSF ERC (Johns Hopkins University agreement) 8810-274. We thank Dr. Yoshinobu Sato of Osaka University Medical School for discussions on the algorithm and validation, and we thank Dan Kacher of Brigham & Women's Hospital for acquisition of the cerebral MRA scans used in this study.

References

Adalsteinsson, D., Sethian, J.A., 1995. A fast level set method for propagating interfaces. Journal of Computational Physics 118, pp. 269–277

Altschuler, S., Grayson, M., 1992. Shortening space curves and flow through singularities. Journal of Differential Geometry 35, pp. 283–298

Ambrosio, L., Soner, H.M., 1996. Level set approach to mean curvature flow in arbitrary codimension. Journal of Differential Geometry 43, pp. 693–737

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

Bullitt, E., Liu, A., Aylward, S., Pizer, S.M., 1997. Reconstruction of the intracerebral vasculature from

MRA and a pair of projection views. Proc. Int'l Conf. Information Procession Medical Imaging. Lect. Notes Comp. Sci. 1230, pp. 537–542

Caselles, V., Catte, F., Coll, T., Dibos, F., 1993. A geometric model for active contours. Numerische Mathematik 66 pp. 1–31

Caselles, V., Kimmel, R., Sapiro, G., 1997. Geodesic active contours. Int'l Journal of Computer Vision 22(1), pp. 61–79

Caselles, V., Morel, J.M., Sapiro, G., Tannenbaum, A., 1998. 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), pp. 269–273

Catte, F., Coll, T., Lions, P.L., Morel, J.M., 1992. Image selective smoothing and edge detection by nonlinear diffusion. SIAM Journal Numerical Analysis 29(1), pp. 182–193

Chen, Y.G., Giga, Y., Goto, S., 1991. Uniqueness and existence of viscosity solutions of generalized mean curvature flow equations. Journal of Differential Geometry 33, pp. 749–786

Chopp, D.L., 1993. Computing minimal surfaces via level set flow. Journal of Computational Physics 106, pp. 77–91

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

Evans, L.C., Spruck, J., 1991. Motion of level sets by mean curvature: I. Journal of Differential Geometry 33, pp. 635–681

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

Gage, M., Hamilton, R.S., 1986. The heat equation shrinking convex plane curves. Journal of Differential Geometry 23, pp. 69–96

Gomes, J., Faugeras, O.D., 2000. Reconciling Distance Functions and Level Sets. Journal of Visual Communication and Image Representation 11, pp. 209–223

Grayson, M., 1987. The heat equation shrinks embedded plane curves to round points. Journal of Differential Geometry 26, pp. 285–314

Grimson, W.E.L., Ettinger, G.J., Kapur, T., Leventon, M.E., Wells III, W.M., Kikinis, R., 1996. Utilizing Segmented MRI Data in Image-Guided Surgery. Int'l Journal Pattern Recognition and Artificial Intelligence

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

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

Krissian, K., Malandain, G., Ayache, N., 1999. Model based detection of tubular structures in 3D images. INRIA Technical Report 373

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

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

Lorenz, C., Carlsen, I.-C., Buzug, T.M., Fassnacht, C., Weese, J., 1997. A multi-scale line filter with automatic scale selection based on the Hessian matrix for medical image segmentation. Proc. Scale-Space Theories in Computer Vision. Lect. Notes Comp. Sci. 1252, pp. 152–163

Lorigo, L. M., Faugeras, O., Grimson, W.E.L., Keriven, R., Kikinis, R., Westin, C.-F., 1999. Co-dimension 2 geodesic active contours for MRA segmentation. Proc. Int'l Conf. Information Procession Medical Imaging. Lect. Notes Comp. Sci. 1613, pp. 126–139

Malladi, R., Sethian, J.A., Vemuri, B.C., 1994. Evolutionary fronts for topology-independent shape modeling and recovery. Proc. European Conf. Computer Vision. Lect. Notes Comp. Sci. 800, pp. 3–13

McInerney, T., Terzopoulos, D., 1996. Deformable models in medical image analysis: a survey, Medical Image Analysis 1(2), pp. 91–108

McInerney, T., Terzopoulos, D., 1999. Topology adaptive snakes. Medical Image Analysis, *In Press*

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

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

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

Sapiro, G., 1996. Vector-valued active contours. In Proc. IEEE Conf. Computer Vision and Pattern Recognition, pp. 680–685

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

Sethian, J.A., 1996. Level Set Methods. Cambridge University Press

Tek, H., Kimia, B., 1997. Volumetric segmentation of medical images by three-dimensional bubbles. Computer Vision Image Understanding 65(2), pp. 246–258

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

Zeng, X., Staib, L.H., Schultz, R.T., Duncan, J.S., 1998. 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, pp. 519–530

Zeng, X., Staib, L.H., Schultz, R.T., Tagare, H., Win, L., Duncan, J.S., 1999. A new approach to 3D sulcal ribbon finding from MR images. Proc. Int'l Conf. Medical Image Computing Computer-Assisted Intervention. Lect. Notes Comp. Sci. 1679, pp. 148–157

Zhao, H-K., Chan, T., Merriman, B., Osher, S., 1996. A variational level set approach to multiphase motion. Journal of Computational Physics 127, pp. 179–195

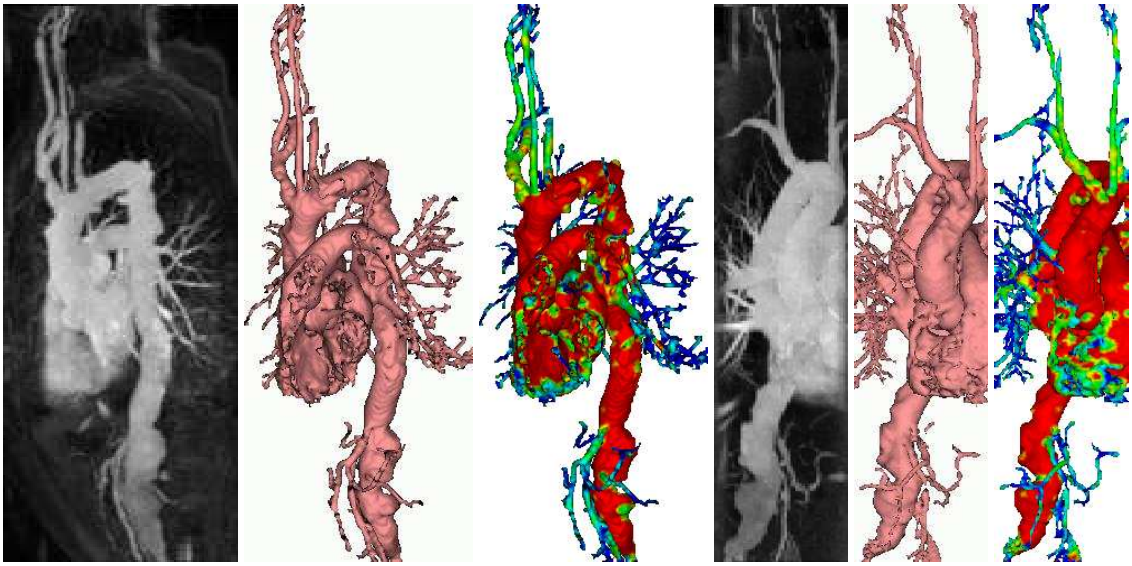

Figure 8. Segmentation of a contrast-enhanced aorta MRA image, courtesy of Siemens, acquired on a Siemens scanner at New York University. From each viewpoint is shown the maximum intensity projection of the raw data, the CURVES segmentation, and the CURVES segmentation color-coded by local radii, where the colorscale ranges from blue to red in order of increasing radius.

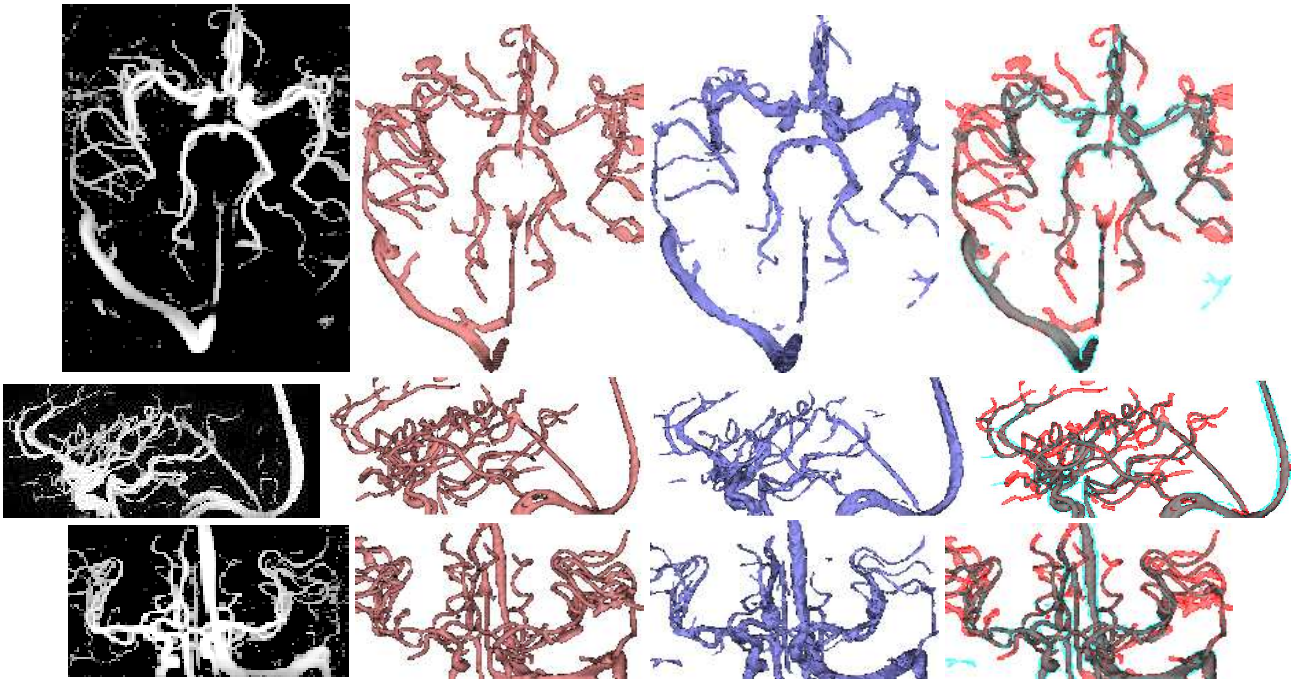

Figure 9. The same cerebral MRA dataset is shown from three orthogonal viewpoints. For each viewpoint, the maximum intensity projection of the raw data is shown, followed by the CURVES segmentation (red), the manual segmentation (blue), and a combination image showing the differences between the segmentations.

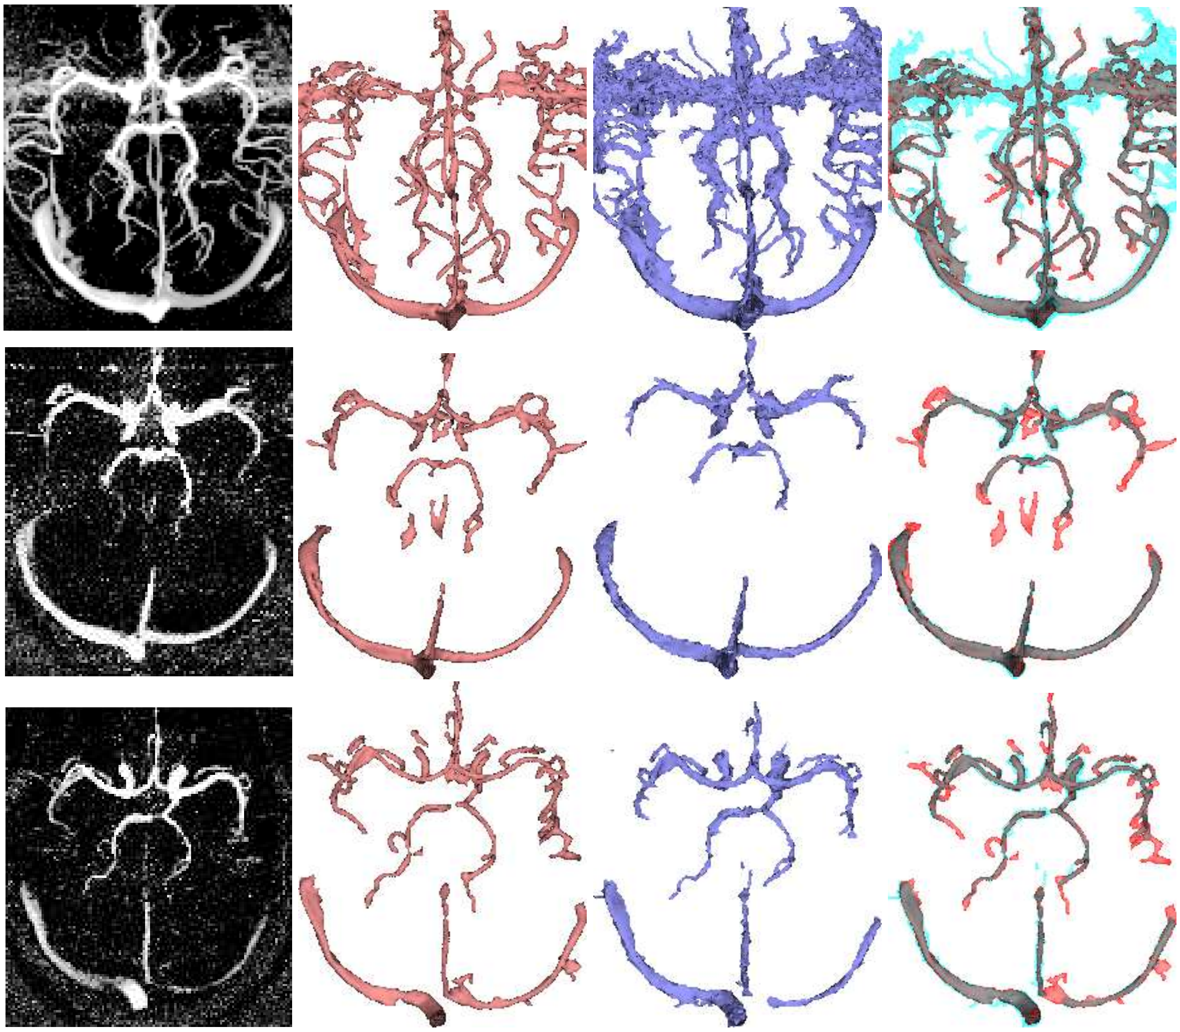

Figure 10. Each row shows a different dataset. Left to right: maximum intensity projection of raw data, CURVES segmentation (red), manual segmentation (blue), combination image showing the differences between the segmentations.

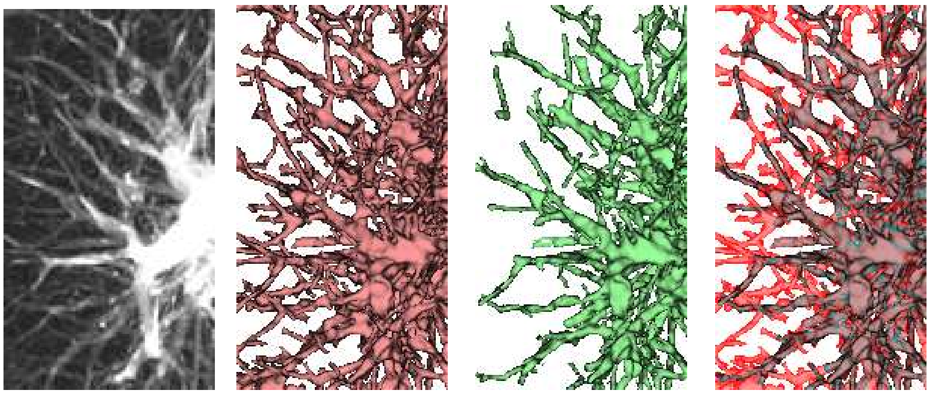

Figure 11. Segmentation of a volumetric subregion of a CT lung scan by CURVES (red) compared to one obtained by a codimension-one (green) method.