

3D Flux Maximizing Flows

Kaleem Siddiqi and Alexander Vasilevskiy

McGill University
School of Computer Science &
Center for Intelligent Machines
3480 University Street
Montréal, QC H3A 2A7, Canada
{siddiqi,sasha}@cim.mcgill.ca

Abstract. A number of geometric active contour and surface models have been proposed for shape segmentation in the literature. The essential idea is to evolve a curve (in 2D) or a surface (in 3D) so that it clings to the features of interest in an intensity image. Several of these models have been derived, using a variational formulation, as gradient flows which minimize or maximize a particular energy functional. However, in practice these models often fail on images of low contrast or narrow structures. To address this problem we have recently proposed the idea of maximizing the rate of increase of flux of an auxiliary vector field through a curve. This has led to an interpretation as a 2D gradient flow, which is essentially parameter free. In this paper we extend the analysis to 3D and prove that the form of the gradient flow does not change. We illustrate its potential with level-set based segmentations of blood vessels in a large 3D computed rotational angiography (CRA) data set.

1 Introduction

Level-set based numerical methods for hyperbolic conservation laws developed by Osher and Sethian [14] for curvature-dependent flame propagation were introduced to the computer vision community for shape analysis by Kimia *et al.* [8]. Such models were later adapted to the problem of shape segmentation independently by Caselles *et al.* [3] and Malladi *et al.* [13]. Here the essential idea was to halt an evolving curve in the presence of intensity edges by multiplying the evolution equation with an image-gradient based stopping potential. This led to new active contour models which, when implemented using level set methods, handled changes in topology due to the splitting and merging of multiple contours in a natural way. These geometric flows for shape segmentation were later given formal motivation as well as unified with the classical energy minimization formulations through several independent investigations [4, 7, 16, 17]. The main idea was to modify the Euclidean arc-length or the Euclidean area by a scalar function and to then derive the resulting gradient evolution equations. Mathematically this amounted to defining a new metric on the plane, tailored to the given image, and then deriving the corresponding gradient flows. The results generalized to the case of evolving surfaces in 3D by adding one more dimension to the variational formulation.

Recently there have been other advances in the use of geometric flows in computer vision, which have both theoretical and practical value. First, it has been recognized that a practical weakness of most geometric flows with stopping terms based purely on local image gradients is that they may “leak” in the presence of weak or low contrast boundaries, are not suitable for segmenting textures and typically require the initial curve or curves to lie entirely inside or outside the regions to be segmented. Thus, a number of researchers have sought to derive flows which take into account the statistics of the regions enclosed by the evolving curves [15, 21]. Further developments include multi-phase motions, which allow triple points to be captured [5], as well as the incorporation of an external force field based on a diffused gradient of an edge map [20]. Second, most geometric flows are not able to capture elongated low contrast structures well, such as blood vessels viewed in 2D and 3D angiography images. At places where such structures are narrow, edge gradients may be weak due to partial volume effects and it is also unclear how to robustly measure region statistics. Approaches to regularizing the flow in 3D by introducing a term proportional to mean curvature have the unfortunate effect of annihilating such structures. To address this issue, Lorigo et al. have proposed the use of active contours with co-dimension 2 (curves in 3D) [12]. The idea is to regularize the flow by a term proportional to the curvature of a 3D curve. The approach is grounded in the level set theory for mean curvature evolution of surfaces of arbitrary co-dimension [1] and has a variational formulation along with an energy minimizing interpretation. However, the derived flow is later modified with a (heuristic) multiplicative term to tailor it to blood vessel segmentation [12].

We have recently suggested an alternate approach to segmenting blood vessels in angiography images, which is motivated by the observation that blood flows in the direction of vessels. Brightness in angiography images is proportional to the magnitude of the blood flow velocity. This leads to the constraint that in the vicinity of blood vessel boundaries, the gradient vector field of the image should be locally orthogonal to them. Thus, a natural principle to use towards the recovery of these boundaries is to maximize the inward flux of the gradient vector field through an evolving curve (in 2D) or surface (in 3D). The derivation of the 2D flux maximizing flow was presented in [18] and lead to an elegant interpretation which is essentially parameter free. In the current paper we prove that the extension to 3D has the same form, a calculation which is more subtle. We also illustrate the potential of the 3D flux maximizing flow with several new simulations of blood vessels segmented from a large 3D computed rotational angiography (CRA) data set.

2 3D Flux Maximizing Flows

Let $\mathcal{S} : [0, 1] \times [0, 1] \rightarrow \mathcal{R}^3$ denote a compact embedded surface with (local) coordinates (u, v) . Let \mathcal{N} be the inward unit normal. We set

$$\mathcal{S}_u := \frac{\partial \mathcal{S}}{\partial u}, \quad \mathcal{S}_v := \frac{\partial \mathcal{S}}{\partial v}.$$

Then the infinitesimal area on \mathcal{S} is given by

$$dS = (\|\mathcal{S}_u\|^2 \|\mathcal{S}_v\|^2 - \langle \mathcal{S}_u, \mathcal{S}_v \rangle^2)^{1/2} du dv$$

$$= \|\mathcal{S}_u \wedge \mathcal{S}_v\| dudv.$$

Let $\mathcal{V} = (V_1(x, y, z), V_2(x, y, z), V_3(x, y, z))$ be a vector field defined for each point (x, y, z) in \mathcal{R}^3 . The total inward flux of the vector field through the surface is defined by the surface integral

$$Flux(t) = \int_0^{A(t)} \langle \mathcal{V}, \mathcal{N} \rangle dS, \quad (1)$$

where $A(t)$ is the surface area of the evolving surface. The main contribution of the current paper is the proof of the following theorem

Theorem 1. *The direction in which the inward flux of the vector field \mathcal{V} through the surface \mathcal{S} is increasing most rapidly is given by $\frac{\partial \mathcal{S}}{\partial t} = \text{div}(\mathcal{V})\mathcal{N}$.*

It turns out that the flux maximizing flow has the same form in 2D, a calculation which we presented in [18]. However, the proof is more subtle for the 3D case.

Proof: The essential idea is to calculate the first variation of the flux functional with respect to t :

$$Flux'(t) = \underbrace{\int_0^{A(t)} \langle \mathcal{V}_t, \mathcal{N} \rangle dS}_{I_1} + \underbrace{\int_0^{A(t)} \langle \mathcal{V}, \mathcal{N}_t \rangle dS}_{I_2}.$$

With $\mathcal{S} = (x(u, v, t), y(u, v, t), z(u, v, t))$, the unit normal vector is given by the normalized cross product of two vectors in the tangent plane:

$$\begin{aligned} \mathcal{N} &= \frac{\mathcal{S}_u \wedge \mathcal{S}_v}{\|\mathcal{S}_u \wedge \mathcal{S}_v\|} = \frac{(N_1, N_2, N_3)}{\|\mathcal{S}_u \wedge \mathcal{S}_v\|} \\ &= \frac{(y_u z_v - y_v z_u), (x_v z_u - x_u z_v), (x_u y_v - x_v y_u)}{\|(y_u z_v - y_v z_u), (x_v z_u - x_u z_v), (x_u y_v - x_v y_u)\|}. \end{aligned} \quad (2)$$

I_1 is then given by

$$\int_0^1 \int_0^1 \langle \mathcal{S}_t, (N_1 \nabla V_1 + N_2 \nabla V_2 + N_3 \nabla V_3) \rangle dudv,$$

where the integrand is the inner product of \mathcal{S}_t with another vector. We shall now simplify I_2 so that it takes on a similar form. It turns out to be advantageous to express the unit normal vector in Eq. (2) as $\frac{\mathcal{S}_u \wedge \mathcal{S}_v}{\|\mathcal{S}_u \wedge \mathcal{S}_v\|}$ and expand it in terms of the partial derivatives $x_u, x_v, y_u, y_v, z_u, z_v$ only later. With $dS = \|\mathcal{S}_u \wedge \mathcal{S}_v\| dudv$, I_2 can be rewritten as

$$\int_0^1 \int_0^1 \langle \mathcal{V}, (\mathcal{S}_u \wedge \mathcal{S}_v)_t \rangle dudv.$$

The trick now is to exploit the fact that for any vectors A, B and C , the following properties of inner products and cross products hold:

$$\begin{aligned} A \wedge B &= -B \wedge A \\ \langle A, (B \wedge C) \rangle &= \langle (A \wedge B), C \rangle \\ (A \wedge B)_t &= (A_t \wedge B) + (A \wedge B_t). \end{aligned}$$

Hence, I_2 can be written as

$$\begin{aligned}
I_2 &= \int_0^1 \int_0^1 \langle \mathcal{V}, (\mathcal{S}_{ut} \wedge \mathcal{S}_v + \mathcal{S}_u \wedge \mathcal{S}_{vt}) \rangle dudv \\
&= \int_0^1 \int_0^1 \langle \mathcal{V}, (\mathcal{S}_{ut} \wedge \mathcal{S}_v) \rangle dudv + \int_0^1 \int_0^1 \langle \mathcal{V}, (\mathcal{S}_u \wedge \mathcal{S}_{vt}) \rangle dudv \\
&= \int_0^1 \int_0^1 -\langle \mathcal{V}, (\mathcal{S}_v \wedge \mathcal{S}_{ut}) \rangle dudv + \int_0^1 \int_0^1 \langle \mathcal{V}, (\mathcal{S}_u \wedge \mathcal{S}_{vt}) \rangle dudv \\
&= \int_0^1 \underbrace{\left[\int_0^1 -\langle (\mathcal{V} \wedge \mathcal{S}_v), \mathcal{S}_{ut} \rangle du \right]}_{I_3} dv + \int_0^1 \underbrace{\left[\int_0^1 \langle (\mathcal{V} \wedge \mathcal{S}_u), \mathcal{S}_{vt} \rangle dv \right]}_{I_4} du
\end{aligned}$$

Using integration by parts, I_3 works out to be

$$-\underbrace{\langle (\mathcal{V} \wedge \mathcal{S}_v), \mathcal{S}_t \rangle_0^1}_{\text{equals 0}} + \int_0^1 \langle \mathcal{S}_t, (\mathcal{V} \wedge \mathcal{S}_v)_u \rangle du$$

Similarly, using integration by parts, I_4 works out to be

$$\underbrace{\langle (\mathcal{V} \wedge \mathcal{S}_u), \mathcal{S}_t \rangle_0^1}_{\text{equals 0}} - \int_0^1 \langle \mathcal{S}_t, (\mathcal{V} \wedge \mathcal{S}_u)_v \rangle dv$$

Combining I_3 and I_4 , I_2 works out to be

$$\int_0^1 \int_0^1 \langle \mathcal{S}_t, (\mathcal{V} \wedge \mathcal{S}_v)_u - (\mathcal{V} \wedge \mathcal{S}_u)_v \rangle dudv.$$

It can now be seen that the integrand in I_2 has the desired form of the inner product of \mathcal{S}_t with another vector. Hence, combining I_1 and I_2 , the first variation of the flux is

$$\int_0^1 \int_0^1 \langle \mathcal{S}_t, N_1 \nabla V_1 + N_2 \nabla V_2 + N_3 \nabla V_3 + (\mathcal{V} \wedge \mathcal{S}_v)_u - (\mathcal{V} \wedge \mathcal{S}_u)_v \rangle dudv.$$

Note that

$$\begin{aligned}
(\mathcal{V} \wedge \mathcal{S}_v)_u - (\mathcal{V} \wedge \mathcal{S}_u)_v &= (\mathcal{V}_u \wedge \mathcal{S}_v) + (\mathcal{V} \wedge \mathcal{S}_{vu}) - (\mathcal{V} \wedge \mathcal{S}_{uv}) - (\mathcal{V}_v \wedge \mathcal{S}_u) \\
&= (\mathcal{V}_u \wedge \mathcal{S}_v) - (\mathcal{V}_v \wedge \mathcal{S}_u).
\end{aligned}$$

Hence, the first variation of the flux can be written as the surface integral

$$\int_0^{A(t)} \left\langle \mathcal{S}_t, \frac{N_1 \nabla V_1 + N_2 \nabla V_2 + N_3 \nabla V_3 + (\mathcal{V}_u \wedge \mathcal{S}_v) - (\mathcal{V}_v \wedge \mathcal{S}_u)}{\|\mathcal{S}_u \wedge \mathcal{S}_v\|} \right\rangle dS.$$

Thus, for the inward flux to increase as fast as possible, the two vectors should be made parallel:

$$\mathcal{S}_t = \frac{N_1 \nabla V_1 + N_2 \nabla V_2 + N_3 \nabla V_3 + (\mathcal{V}_u \wedge \mathcal{S}_v) - (\mathcal{V}_v \wedge \mathcal{S}_u)}{\|\mathcal{S}_u \wedge \mathcal{S}_v\|}. \quad (3)$$

The above expression for the 3D flux maximizing gradient flow can be further simplified by noting that the components of the flow in the tangential plane to the surface \mathcal{S} affect only the parametrization of the surface, but not its evolved shape. Hence, they can be dropped. The normal component of the flow can be calculated by taking the inner product of the right hand side of Eq. (3) with the unit normal vector in Eq. (2) to give

$$\mathcal{S}_t = \left\langle \frac{N_1 \nabla V_1 + N_2 \nabla V_2 + N_3 \nabla V_3 + (\mathcal{V}_u \wedge \mathcal{S}_v) - (\mathcal{V}_v \wedge \mathcal{S}_u)}{\|\mathcal{S}_u \wedge \mathcal{S}_v\|}, \frac{\mathcal{S}_u \wedge \mathcal{S}_v}{\|\mathcal{S}_u \wedge \mathcal{S}_v\|} \right\rangle \mathcal{N}$$

It is now a straightforward task to expand the terms in the expression by using Eq. (2):

$$\begin{aligned} \mathcal{S}_t &= \frac{(y_u z_v - y_v z_u)}{\|\mathcal{S}_u \wedge \mathcal{S}_v\|^2} (V_{1x}(y_u z_v - y_v z_u) + V_{1y}(x_v z_u - x_u z_v) + V_{1z}(x_u y_v - x_v y_u)) \\ &+ \frac{(x_v z_u - x_u z_v)}{\|\mathcal{S}_u \wedge \mathcal{S}_v\|^2} (V_{2x}(y_u z_v - y_v z_u) + V_{2y}(x_v z_u - x_u z_v) + V_{2z}(x_u y_v - x_v y_u)) \\ &+ \frac{(x_u y_v - x_v y_u)}{\|\mathcal{S}_u \wedge \mathcal{S}_v\|^2} (V_{3x}(y_u z_v - y_v z_u) + V_{3y}(x_v z_u - x_u z_v) + V_{3z}(x_u y_v - x_v y_u)) \\ &+ \frac{1}{\|\mathcal{S}_u \wedge \mathcal{S}_v\|^2} \langle (\mathcal{V}_u \wedge \mathcal{S}_v), (y_u z_v - y_v z_u, x_v z_u - x_u z_v, x_u y_v - x_v y_u) \rangle \\ &- \frac{1}{\|\mathcal{S}_u \wedge \mathcal{S}_v\|^2} \langle (\mathcal{V}_v \wedge \mathcal{S}_u), (y_u z_v - y_v z_u, x_v z_u - x_u z_v, x_u y_v - x_v y_u) \rangle. \end{aligned}$$

With

$$\begin{aligned} \mathcal{V}_u \wedge \mathcal{S}_v &= (z_v(V_{2x}x_u + V_{2y}y_u + V_{2z}z_u) - y_v(V_{3x}x_u + V_{3y}y_u + V_{3z}z_u), \\ &x_v(V_{3x}x_u + V_{3y}y_u + V_{3z}z_u) - z_v(V_{1x}x_u + V_{1y}y_u + V_{1z}z_u), \\ &y_v(V_{1x}x_u + V_{1y}y_u + V_{1z}z_u) - x_v(V_{2x}x_u + V_{2y}y_u + V_{2z}z_u)). \end{aligned}$$

and

$$\begin{aligned} -\mathcal{V}_v \wedge \mathcal{S}_u &= (-z_u(V_{2x}x_v + V_{2y}y_v + V_{2z}z_v) + y_u(V_{3x}x_v + V_{3y}y_v + V_{3z}z_v), \\ &-x_u(V_{3x}x_v + V_{3y}y_v + V_{3z}z_v) + z_u(V_{1x}x_v + V_{1y}y_v + V_{1z}z_v), \\ &-y_u(V_{1x}x_v + V_{1y}y_v + V_{1z}z_v) + x_u(V_{2x}x_v + V_{2y}y_v + V_{2z}z_v)) \end{aligned}$$

the terms can be grouped and simplified. The curious result is that most cancel, leaving the following simple and elegant form for the 3D flux maximizing flow:

$$\mathcal{S}_t = (V_{1x} + V_{2y} + V_{3z})\mathcal{N} = \text{div}(\mathcal{V})\mathcal{N} \quad \square \quad (4)$$

Remark: The flux maximizing flow is a hyperbolic equation since it depends solely on the external vector field \mathcal{V} and not on properties of the evolving surface. It is easy to see that the flow will drive towards and then converge to a zero level set of the divergence of \mathcal{V} . Thus, the existence and uniqueness of a solution to Eq. (4) is guaranteed, unless the vector field is everywhere non-conservative.

3 Blood Vessel Segmentation

3.1 Background

We shall now show how the 3D flux maximizing flow can be tailored to the problem of segmenting blood vessels in angiography images. We begin by reviewing some of the

recent approaches which have been proposed in the literature. Wilson and Noble have introduced a Gaussian mixture model to characterize the physical properties of blood flow [19]. The parameters are estimated using the EM algorithm and structural criteria are then used to refine the initial segmentation. Krissian et al. propose a method which incorporates a Gaussian model for the intensity distribution as a function of distance from vessel centerlines, and exploits properties of the Hessian to obtain geometric estimates [11]. Koller et al. have also introduced a multi-scale method for the detection of curvilinear structures in 2D and 3D data [9] which combines the responses of steerable linear filters and also exploits the Hessian matrix to obtain geometric estimates. Bullitt et al. have introduced a method for obtaining 3D vascular trees which calculates vessel centerlines as intensity ridges in the data and estimates vessel width via medialness calculations [2]. It should be noted that several of the above approaches require second derivative computations, e.g., to compute the Hessian. Numerically accurate estimates of principal curvature magnitudes and directions are obtained only when the intensity images have been suitably smoothed. Approaches to smoothing the data while preserving and enhancing vessel-like structures include [10, 6].

Whereas the potential of several of the above approaches has been empirically demonstrated, their ability to recover low contrast thin vessels remains unclear. A recent framework which has been developed with this as one of its goals is the work of Lorigo et al. [12]. The main idea is to regularize a geometric flow in 3D using the curvature of a 3D curve, rather than the classical mean curvature based regularizations which tend to annihilate thin structures. The work is grounded in the recent level set theory developed for mean curvature flows in arbitrary co-dimension [1]. This flow is given by [12]:

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

Here ψ is an embedding surface whose zero level set is the evolving 3D curve, λ is the smaller nonzero eigen value of a particular matrix [1] g is an image-dependent weighting factor, \mathbf{I} is the intensity image and \mathbf{H} is its Hessian. For numerical simulations the evolution of the curve is depicted by the evolution of an ϵ -level set. It should be noted that without the multiplicative factor $\rho(\nabla\psi \cdot \nabla\mathbf{I})$ the evolution equation is a gradient flow which minimizes a weighted curvature functional. The multiplicative factor is a heuristic which modifies the flow so that normals to the ϵ -level set align themselves (locally) to the direction of image intensity gradients (the inner product of $\nabla\psi$ and $\nabla\mathbf{I}$ is then maximized). However, with the introduction of this term the flow loses its pure energy minimizing interpretation.

3.2 The 3D Flux Maximizing Flow

The intuition behind using the 3D flux maximizing flow for blood vessel segmentation is illustrated in Figure 1. Here a cross section through an idealized blood vessel (a bright region in a uniform darker background) is depicted. It is clear that if one considers the gradient $\nabla\mathbf{I}$ of the original intensity image \mathbf{I} to be the vector field \mathcal{V} whose inward flux through the evolving surface is to be maximized, then the optimal configuration is for the evolving surface to align itself locally to the blood vessel boundaries. However, an

important consideration in the implementation of Eq. (4) is that since the divergence of the vector field needs to be calculated, implicitly second derivatives of \mathbf{I} are being used. The numerical computation can be made much more robust by exploiting a consequence of the divergence theorem. The divergence at a point is defined as the net *outward* flux per unit volume, as the volume about the point shrinks to zero. Via the divergence theorem,

$$\int_{\Delta v} \text{div}(\mathcal{V})dv \equiv \int_s \langle \mathcal{V}, \mathcal{N} \rangle dS. \quad (5)$$

Here Δv is the volume, s is its bounding surface, \mathcal{N} is the outward normal at each point on the surface, and dv and dS are volume and surface area elements.

For our numerical implementations we use this outward flux formulation (which gives a measure proportional to the divergence) along the boundaries of spheres of varying radii, corresponding to a range of 3D blood vessel widths. The chosen flux value at a particular location is the maximum (magnitude) flux over the range of radii. In contrast to other multi-scale approaches where combining information across scales is non-trivial [11] normalization across scales is straightforward in our case. One simply has to divide by the number of entries in the discrete sum that approximates Eq. (5). Locations where the total outward flux is negative (or equivalently the total inward flux is positive) correspond to sinks; locations where the total outward flux is positive correspond to sources, as illustrated in Figure 1. Hence, the inward flux maximizing flow in Eq (4) has the desirable effect that, when seeds are placed within blood vessels, the sinks drive the seeds towards the vessel boundaries and the sources outside prevent the flow from leaking.

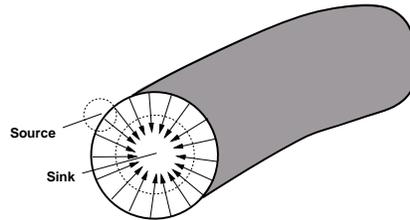


Fig. 1. An illustration of the gradient vector field of an angiography image along a cross-section of a 3D blood vessel. Assuming a uniform background intensity, at the scale of the vessel's width the total inward flux is positive (a sink). Outside the vessel, at a smaller scale, the total inward flux is negative (a source).

3.3 Level Set Implementation

In order to implement the flow, we use the level set representation for compact embedded surfaces, due to Osher and Sethian [14]. Let $\mathcal{S}(u, v, t) : [0, 1] \times [0, 1] \times [0, \tau) \rightarrow \mathbf{R}^3$ be a family of compact embedded surfaces evolving according to the surface evolution equation

$$\mathcal{S}_t = F\mathcal{N},$$

where F is an arbitrary (local) scalar speed function and \mathcal{N} is inward unit normal to the surface. Then it can be shown that if $\mathcal{S}(u, v, t)$ is represented by the zero level set of a smooth and Lipschitz continuous function $\Psi : \mathbf{R}^3 \times [0, \tau) \rightarrow \mathbf{R}$, the evolving hypersurface Ψ satisfies

$$\Psi_t = F \|\nabla\Psi\|.$$

This last equation is solved using a combination of straightforward discretization and numerical techniques derived from hyperbolic conservation laws [14]. For hyperbolic terms, care must be taken to implement derivatives with upwinding in the proper direction. The evolving surface \mathcal{S} is then obtained as the zero level set of Ψ .

4 Examples

In earlier work we have presented simulations of the flux maximizing flow on both 2D (retinal) and 3D (head) MRA data [18]. In this section we present new experiments on a $360 \times 330 \times 420$ computed rotational angiography (CRA) data set of the head, from which we have selected four distinct regions containing vascular networks of varying complexity. All examples were implemented in a level-set framework.

The evolution results are presented in Figures 2, 3, 4 and 5. For each region a maximal intensity projection of the data is shown on the top left, followed by the evolution of a few 3D spheres. These spheres were placed somewhat sparsely in regions of high flux. Notice how the spheres elongate in the direction of blood vessels. The main blood vessels, which have the higher inward flux, are the first to be captured. This is the expected evolution since it maximizes the rate of increase of inward flux through the evolving surface. Put another way, the evolution has the intuitive behaviour that it follows the direction of blood flow to reconstruct the blood vessel boundaries. Our own experience with developing and implementing some of the related geometric flows in the literature is that many would fail in low contrast regions or would not be able to capture the thinner vessels.

We should also point out that although CRA data is of higher resolution than MRA, the vessel structures exhibit a wider range of intensities and there are also a number of other structures whose intensities overlap with those of the thin vessels. Thus, simple thresholding of the intensity data generally gives poor results, although this is a commonly used initialization step in many algorithms including the approach of [12]. This point is illustrated in Figure 6. The first row shows the results of a high threshold on the four regions, where as one would expect, many thin low contrast vessels are not captured. As the threshold is decreased, more thin vessels are captured, but also many voxels are incorrectly labeled as vessels (Figure 6, second row). The segmentation results obtained by the flux maximizing flow are repeated in the last row. The arrows point

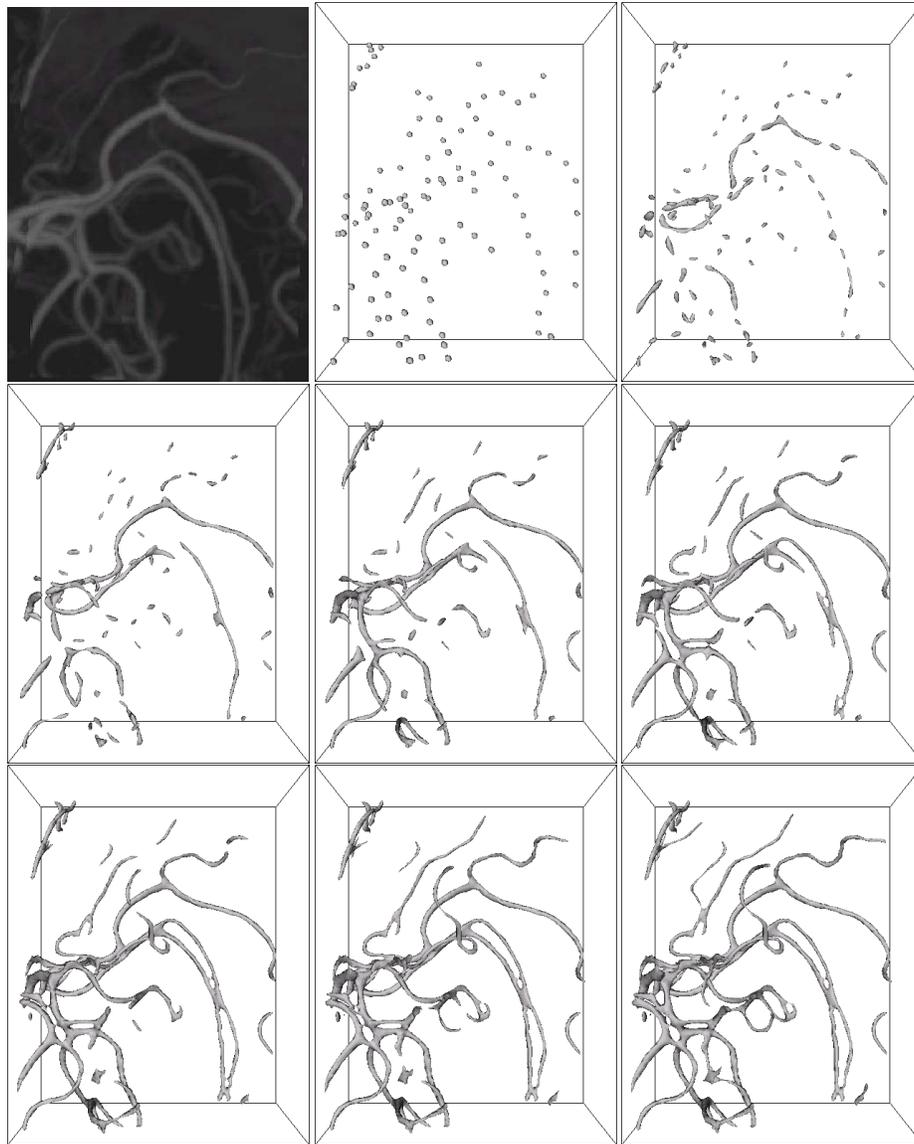


Fig. 2. An illustration of the flux maximizing flow for a portion of a $360 \times 330 \times 420$ 3D CRA image of blood vessels in the head. A maximum-intensity projection of the region being viewed is shown on the top left. The other images depict the evolution of a few isolated spheres. Notice how the evolution follows the direction of blood flow to reconstruct the blood vessel boundaries.

to some of the thin low contrast vessels that are successively captured, but are not seen even in the low threshold case.

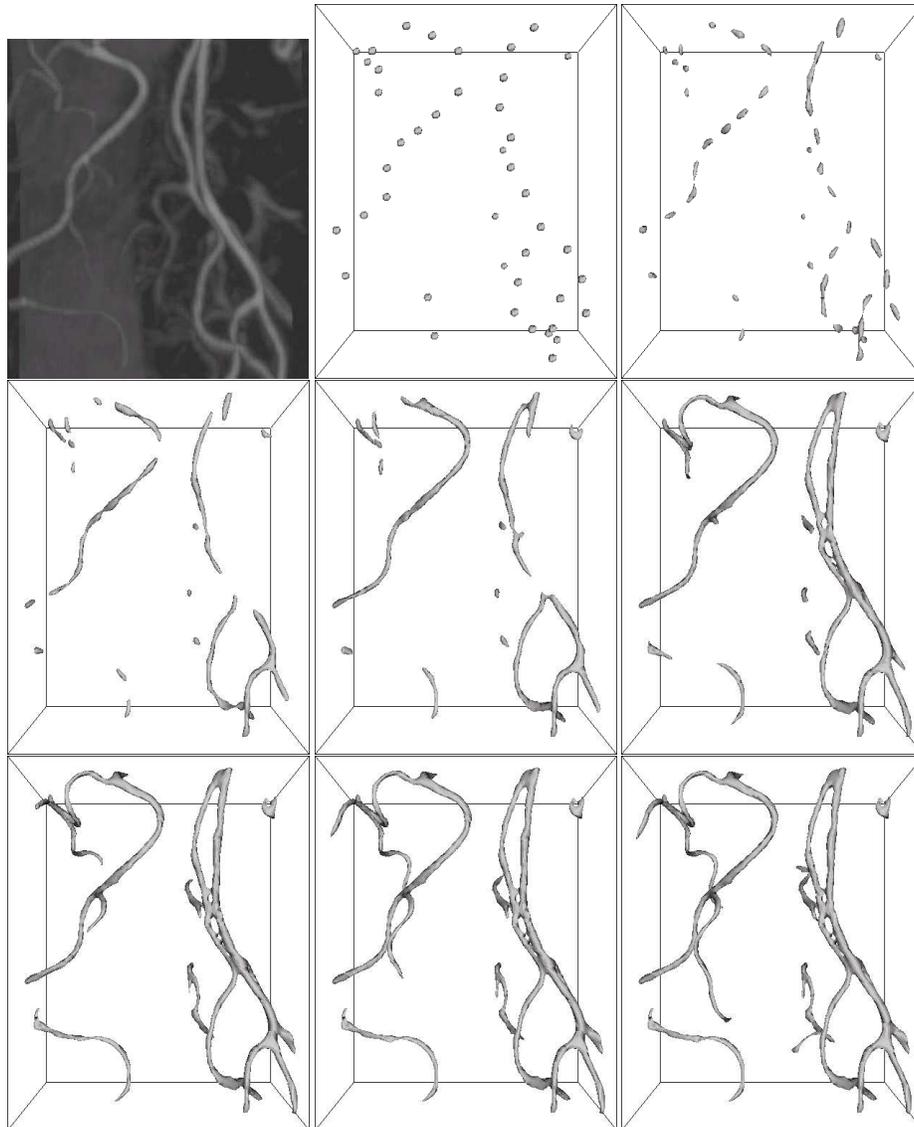


Fig. 3. An illustration of the flux maximizing flow for a portion of a $360 \times 330 \times 420$ 3D CRA image of blood vessels in the head. A maximum-intensity projection of the region being viewed is shown on the top left. The other images depict the evolution of a few isolated spheres. Notice how the evolution follows the direction of blood flow to reconstruct the blood vessel boundaries.

5 Conclusion

In recent work we proposed the flux maximizing flow and derived its form for the case of closed curves evolving in the plane [18]. We also suggested that its form remains the

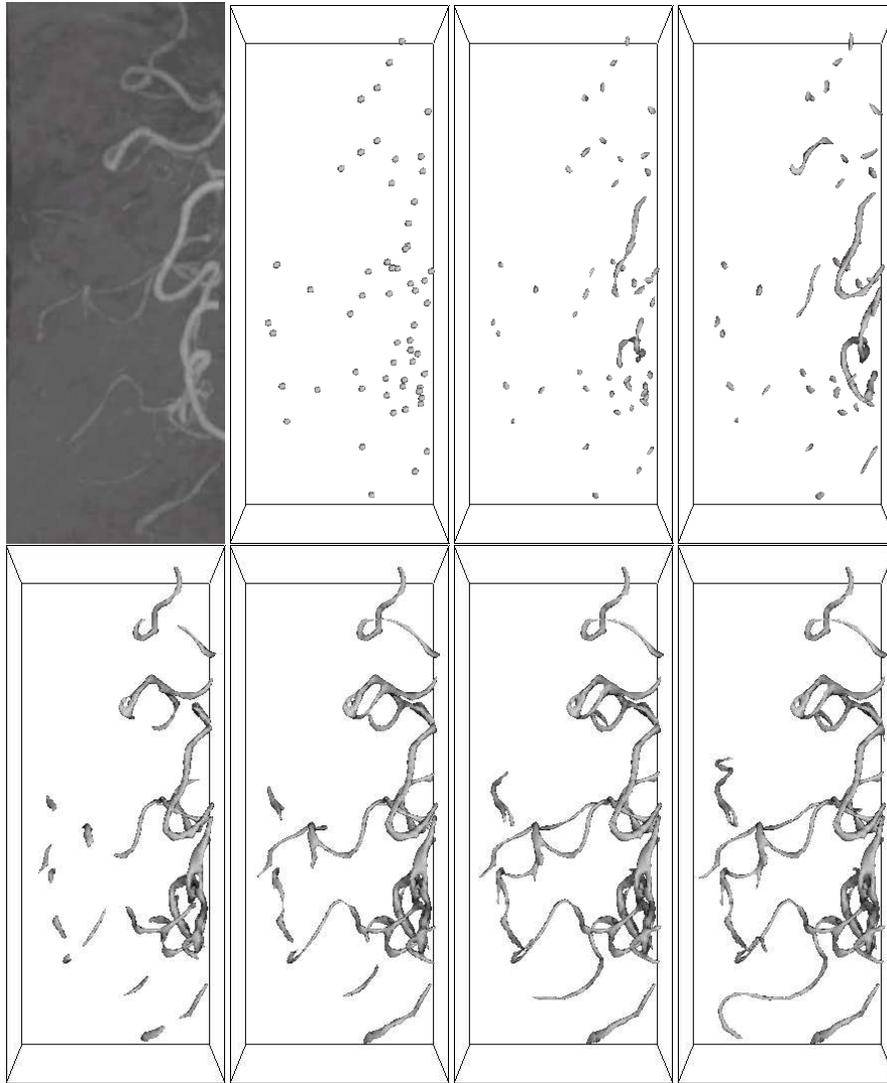


Fig. 4. An illustration of the flux maximizing flow for a portion of a $360 \times 330 \times 420$ 3D CRA image of blood vessels in the head. A maximum-intensity projection of the region being viewed is shown on the top left. The other images depict the evolution of a few isolated spheres. Notice how the evolution follows the direction of blood flow to reconstruct the blood vessel boundaries.

same in higher dimensions. The main contribution of the current paper is the formal derivation of the flux maximizing flow in 3D. We have also carried out a number of new simulations on a large CRA data set. These have the intuitive behaviour that the evolution follows the direction of blood flow to reconstruct blood vessel boundaries. The results suggest the potential of the approach to capture low contrast thin vessels,

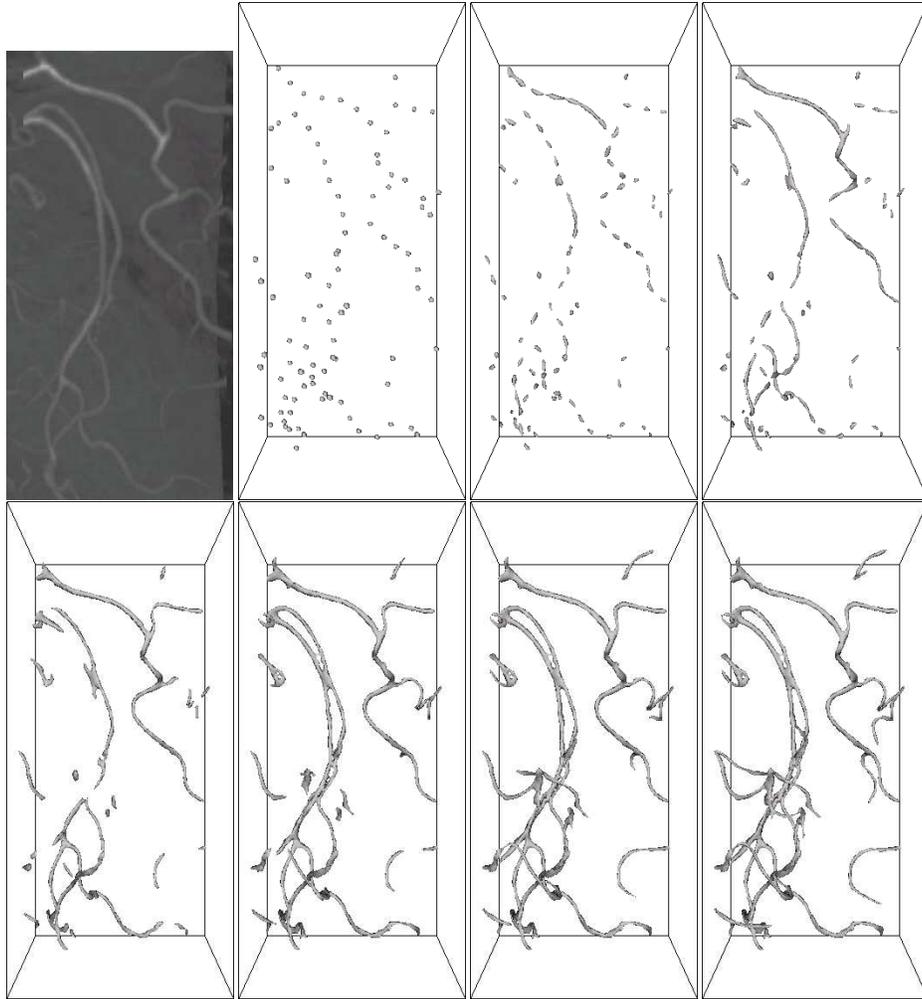


Fig. 5. An illustration of the flux maximizing flow for a portion of a $360 \times 330 \times 420$ 3D CRA image of blood vessels in the head. A maximum-intensity projection of the region being viewed is shown on the top left. The other images depict the evolution of a few isolated spheres. Notice how the evolution follows the direction of blood flow to reconstruct the blood vessel boundaries.

and also illustrate the advantages of the method over thresholding the original intensity image, which is a common initialization step in many vessel reconstruction algorithms.

More work remains to be done to validate this technique against ground truth or expert segmentations, and we are beginning to do this in collaboration with our colleagues in medical imaging. It would also be interesting to see whether a regularization term such as that used in [12] could be incorporated in the derivation of the flux maximizing flow from first principles.

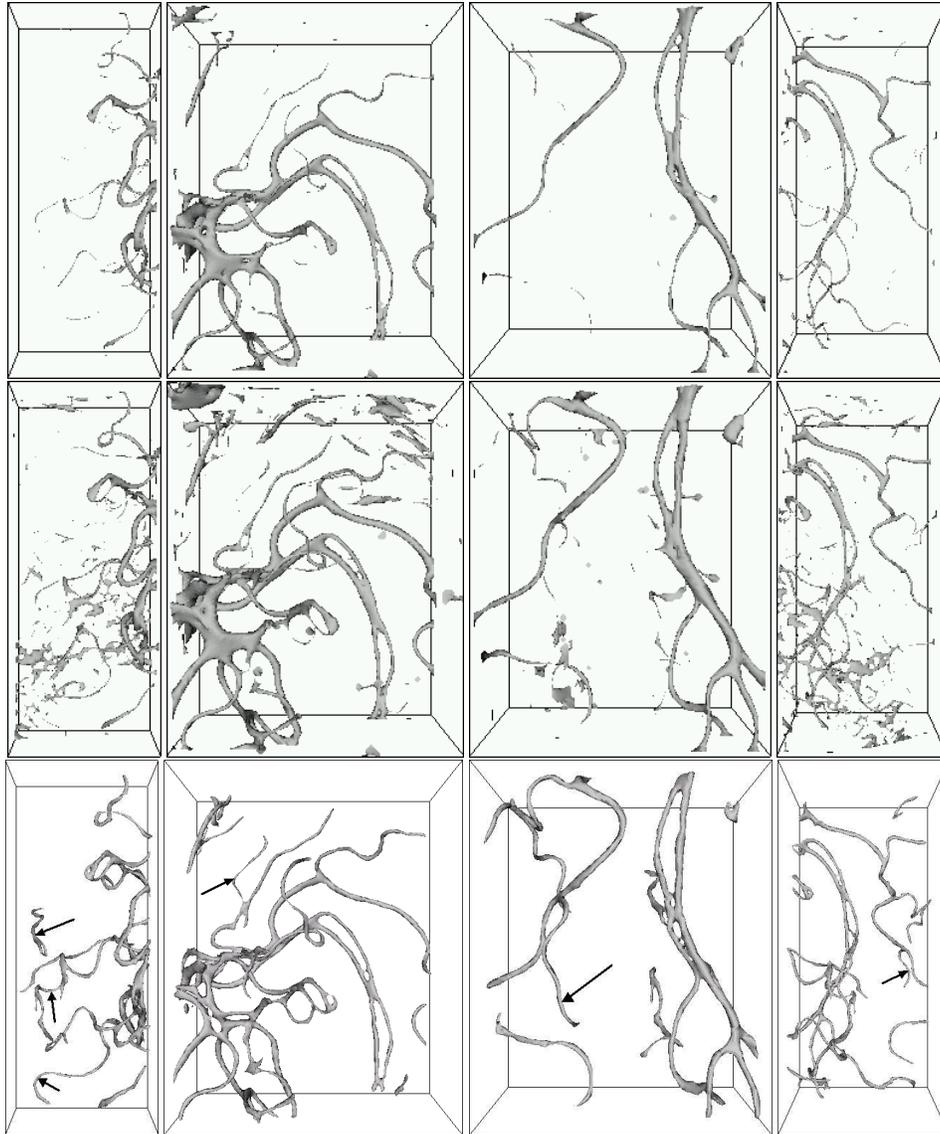


Fig. 6. A comparison of the segmentation results obtained by the flux maximizing flow with simple thresholding on the four different regions of the CRA image. **FIRST ROW:** A conservative high threshold fails to capture many thin low contrast vessels. **SECOND ROW:** A lower threshold captures some of the thinner vessels but also incorrectly labels many voxels. **THIRD ROW:** The segmentation results obtained by the flux maximizing flow. The arrows point to some of the thin low contrast vessels that are successively captured, but are not seen even in the low threshold case.

Acknowledgements This work was supported by grants from the Canadian Foundation for Innovation, FCAR Quebec and the Natural Sciences and Engineering Research Council of Canada. We are grateful to Vincent Hayward, Terry Peters and David Holdsworth for giving us access to the vessel CRA data.

References

1. L. Ambrosio and H. M. Soner. Level set approach to mean curvature flow in arbitrary codimension. *Journal of Differential Geometry*, 43:693–737, 1996.
2. E. Bullitt, S. Aylward, A. Liu, J. Stone, S. K. Mukherjee, C. Coffey, G. Gerig, and S. M. Pizer. 3d graph description of the intracerebral vasculature from segmented mra and tests of accuracy by comparison with x-ray angiograms. In *IPMI'99*, pages 308–321, 1999.
3. V. Caselles, F. Catte, T. Coll, and F. Dibos. A geometric model for active contours in image processing. *Numerische Mathematik*, 66:1–31, 1993.
4. V. Caselles, R. Kimmel, and G. Sapiro. Geodesic active contours. In *ICCV'95*, pages 694–699, 1995.
5. T. Chan and L. Vese. An efficient variational multiphase motion for the mumford-shah segmentation model. In *Asilomar Conference on Signals and Systems*, October 2000.
6. A. Frangi, W. Niessen, K. L. Vincken, and M. A. Viergever. Multiscale vessel enhancement filtering. In *MICCAI'98*, pages 130–137, 1998.
7. S. Kichenassamy, A. Kumar, P. Olver, A. Tannenbaum, and A. Yezzi. Gradient flows and geometric active contour models. In *ICCV'95*, pages 810–815, 1995.
8. B. B. Kimia, A. Tannenbaum, and S. W. Zucker. Toward a computational theory of shape: An overview. *Lecture Notes in Computer Science*, 427:402–407, 1990.
9. T. M. Koller, G. Gerig, G. Székely, and D. Dettwiler. Multiscale detection of curvilinear structures in 2-d and 3-d image data. In *ICCV'95*, pages 864–869, 1995.
10. K. Krissian, G. Malandain, and N. Ayache. Directional anisotropic diffusion applied to segmentation of vessels in 3d images. In *International Conference On Scale Space Theories in Computer Vision*, pages 345–348, 1997.
11. K. Krissian, G. Malandain, N. Ayache, R. Vaillant, and Y. Troussel. Model-based multiscale detection of 3d vessels. In *CVPR'98*, pages 722–727, 1998.
12. L. M. Lorigo, O. Faugeras, E. L. Grimson, R. Keriven, R. Kikinis, A. Nabavi, and C.-F. Westin. Codimension-two geodesic active contours for the segmentation of tubular structures. In *CVPR'2000*, volume 1, pages 444–451, 2000.
13. R. Malladi, J. A. Sethian, and B. C. Vemuri. Shape modeling with front propagation: A level set approach. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 17(2):158–175, February 1995.
14. S. J. Osher and J. A. Sethian. Fronts propagating with curvature dependent speed: Algorithms based on hamilton-jacobi formulations. *Journal of Computational Physics*, 79:12–49, 1988.
15. N. Paragios and R. Deriche. Geodesic active regions for supervised texture segmentation. In *ICCV'99*, pages 926–932, September 1999.
16. J. Shah. Recovery of shapes by evolution of zero-crossings. Technical report, Dept. of Mathematics, Northeastern University, Boston, MA, 1995.
17. K. Siddiqi, Y. B. Lauzière, A. Tannenbaum, and S. W. Zucker. Area and length minimizing flows for shape segmentation. *IEEE Transactions on Image Processing*, 7(3):433–443, 1998.
18. A. Vasilevskiy and K. Siddiqi. Flux maximizing geometric flows. In *ICCV'2001*, July 2001.
19. D. L. Wilson and A. Noble. Segmentation of cerebral vessels and aneurysms from mr angiography data. In *IPMI'97*, pages 423–428, 1997.

20. C. Xu and J. Prince. Snakes, shapes and gradient vector flow. *IEEE Transactions on Image Processing*, 7(3):359–369, 1998.
21. A. Yezzi, A. Tsai, and A. Willsky. A statistical approach to snakes for bimodal and trimodal imagery. In *ICCV'99*, pages 898–903, September 1999.