Segmentation and reconstruction of vascular structures for 3D real-time simulation
Paper Review

Miguel A. Veganzones

Grupo de Inteligencia Computacional
Universidad del País Vasco

2012-03-02

Outline

Paper

Segmentation and reconstruction of vascular structures for 3D real-time simulation ☆

Xunlei Wu [a,e], Vincent Luboz [b,e,*], Karl Krissian [c], Stephane Cotin [d,e], Steve Dawson [e]

[a] Renaissance Computing Inst., Univ of North Carolina at Chapel Hill, USA
[b] Dept. of Biosurgery and Surgical Technology, Imperial College London, UK
[c] Dept de Informática y Sistemas, Univ. Las Palmas de Gran Canaria, Spain
[d] ALCOVE Team, INRIA Lille – North Europe, France
[e] The SIM Group and Department of Imaging, Mass. General Hospital, Cambridge, USA

Motivation

- Authors propose a technique to obtain accurate and smooth surfaces of patient specific vascular structures, using two steps:

  1. Segmentation: provides accurate and smooth centerlines of the vessels, together with cross section orientations and cross section fitting.
  2. Reconstruction: consists of reconstructing the surface by generating a set of topologically preserved quadrilateral patches of branching tubular structures.

- Experiments, on phantom and real datasets, show that the proposed technique reaches a good balance in terms of smoothness, number of triangles, and distance error.

- This technique can be applied in interventional radiology simulations, virtual endoscopy and in reconstruction of smooth and accurate 3D models for use in simulation.

Outline

1. Introduction

2. Description of the method
   - Segmentation of the vessels
   - Surface reconstruction

3. Experiments
   - Phantom
   - Clinical data

Outline

Common techniques

Different techniques can be used for the segmentation of vascular structures. One of the simplest and sometimes most efficient techniques is thresholding, which usually needs some preprocessing to allow outlining tiny tubular structures. More complex methods, such as centerline enhancement and contour extraction, have more potential for computing the topological representation of a vascular network. Both techniques are very efficient for segmenting the carotids and other relatively large vessels, but have trouble outlining small vessels accurately, especially in the brain when datasets have resolutions coarser than $0.5 \times 0.5 \times 1 \, \text{mm}^3$. In this

- Regardless of the technique chosen, it is usually preceded by a noise reduction process, such as anisotropic diffusion filtering.

Proposed segmentation technique

1. enhancement and cleaning the patient data with anisotropic diffusion and morphological operators;
2. segmentation of the vessels through a level set evolution, initialized from a manually selected threshold;
3. skeletonization to obtain the centerlines of the vessels;
4. estimation of the vessel surface through circle or ellipse fitting;
5. cross section post-processing.

- Steps one to three have been achieved using the open source software amilab available at amilab.sourceforge.net.

1.Enhancement and cleaning
Noise reduction

To reduce the data noise while preserving and enhancing small vascular structures, an anisotropic filter is used. This filter is detailed in Krissian (2002), and makes use of a matrix anisotropic diffusion with different coefficients in the gradient and in the principal curvatures directions, computed on a smoothed version of the image.

1.Enhancement and cleaning
Skin removal

2.1.1.2. Skin removal. We applied the following process to remove the skin from the dataset: first, the image is thresholded to separate the tissue intensities from the background, intensities higher that 665 are set to 1 and intensities lower than 665 are set to 0. Then, a closing operation of size 12 mm is applied to fill cavities within the tissues, and finally, an opening of size 5 mm allows taking out the skin area from our mask. A voxel-wise multiplication of the original dataset by this mask performs this erasure.
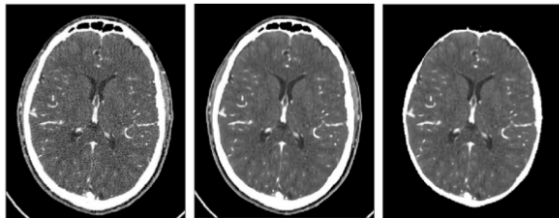
1.Enhancement and cleaning
Results

Fig. 1. Image denoising and skin removal: left, axial slice from the initial dataset, middle: same slice after denoising, right: result after removal of the skin area.

1.Enhancement and cleaning
Bones and skull removal

2.1.1.3. Bones and skull. Then, the bones and their interface with soft tissues are treated. Because bones and skull have much higher intensities than the vessels have, it is easy to remove them from the data set by using classic thresholding. A mask is thus created by thresholding the image at intensity 1530, setting intensities lower than 1530 to one, and intensities higher to zero. Additionally, this mask is eroded by 1.2 mm (approximately 2 voxels) to fully exclude the bones intensities, and the original image is then multiplied by this mask. The result is depicted in Fig. 2. This is the most important part of the cleaning process since the vessels have intensities very similar to this interface and their proximity in some locations might lead to leakage or problems in the level set evolution and in the segmentation process when approximating the vessel's contour or propagating the segmentation.
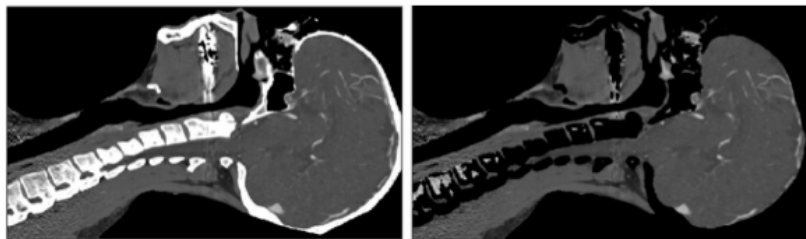
1.Enhancement and cleaning
Results

Fig. 2. Illustration of the removal of the bones on a sagittal slice.

2.Level set segmentation

We segment the vessel contours by means of a level set evolution (Sethian et al., 1999; Osher and Sethian, 1988) applied on the enhanced data set. For better efficiency, the active contour is initialized by a threshold on the image intensity. Manual seed selection and region growing methods would be time consuming and less robust since disconnected parts could be missing. Our implementation for fast evolution of the contour is based on Krissian and Westin (2005). Although the level set technique is computationally more expensive than a straightforward threshold, it allows a better estimation of the contours based on both intensity values and edges. The level set evolves a surface according to three different forces: an advection force that pushes the surface towards the edges of the image; a smoothing term, proportional to the minimal curvature of the surface (Lorigo et al., 2001), that keeps the surface smooth; and a balloon force, defined as

$$F_b(I) = \begin{cases} e^{-\frac{(I-m)^2}{\sigma^2}} - \tau & \text{if } I < m, \\ 1 - \tau & \text{if } I \geqslant m. \end{cases} \tag{1}$$

allows the contours to expand within vascular structures, this force is defined by Eq. (1), and depends on the intensity of each voxel. It relies on the intensity statistics to either expand or shrink the evolving contour. I is the intensity image, m stands for the mean intensity of the vessels, σ is their standard deviation, and τ is a threshold that allows shrinking the contour when the position is unlikely to belong to a vessel.

3.Skeletonization and centerline extraction

From the images segmented by the level set equation, a skeletonization technique is applied to obtain the vessels centerlines. This process allows our simulator to efficiently perform collision detection and blood flow computation by supplying an abstract topological representation of the vascular network. It is based on homotopic thinning where voxels are removed in the order of the Euclidean distance to the segmented surface. Voxels are iteratively removed from the object if they are simple (i.e., if they can be deleted without changing the object's topology (Malandain et al., 1993)) and if they are not end-points, such that they have more than one neighbor in a $3 \times 3 \times 3$ neighborhood. The advantage of the homotopic thinning is to preserve the topology of the vascular network, providing connection at junctions.

3.Skeletonization and centerline extraction

Skeletonization leads to a set of rough centerlines which can still have connectivity discrepancies, especially near small branches. Consequently, automatic pruning is applied to remove small leaves (lines connected at only one extremity, or lines with no connection) from the centerline tree. Given the resolution of the medical data set, some lines remain disconnected when they should be part of the same vessel. We connect them using an automatic process which selects close lines having a similar direction. This direction criterion, computed between the two nodes at the extremity of a line, helps to match lines within a small curvature difference. This step often requires manual adjustment since a few lines might be too long to be deleted by pruning or too distant to be connected automatically. These areas need human judgment to delete or connect appropriate lines and to complete the vascular skeleton. Manual intervention in this streamlined reconstruction is needed only for "breaks" in some vessels. These breaks are discontinuities produced by artifacts, such as CT artifacts caused by dental amalgams, or by low data resolution that makes thin vessels look like dashed lines. This manual interaction is not required in good data sets, e.g. higher resolution datasets with vessels clearly separated from each other and few artifacts without vessels or ambient areas having similar grey scale.

4.Circle and ellipse fitting and centerlines adjustment

Once the vascular structures are finally connected, one possibility is to model the cross sections as circles and to estimate their radii around the centerlines (Luboz et al., 2005). This process is based on an algorithm that grows a circle in the orthogonal plane to the centerline points. It computes, along circles of increasing radii, the intensity gradient, i.e. the derivatives of a Gaussian kernel with a given standard deviation, in the medical data set. It stops when a relevant local maximum of the intensity gradient is found on the cross section therefore estimating the radii along the centerlines. We introduced in Krissian et al. (2006) an algorithm for fitting an ellipse instead of a circle. The ellipse fitting is based on the method described in Fitzgibbon et al. (1999). Starting from the estimated centerlines and a rough segmentation of the vessel boundaries, ellipses are fitted in the planes of the vascular cross sections, defined as the planes orthogonal to the centerlines. The ellipses are fitted at points regularly distributed along each centerline. The problem of fitting ellipses is described in the Appendix A.

5.Cross section post-processing

Once the cross sections are estimated along the centerlines, post-processing is applied to the data (ellipse or circle). This step allows filling of any the possible gaps in the estimation, due to a low resolution data set, artifacts, or errors when two or more vascular structures are touching each other. This post-processing method enhances the cross section estimation and guarantees a smoother surface, which closely approximates the segmented one. It is based on a computation of the average axis of the ellipse at the current sample point, which is compared to other cross section's estimates (circles or ellipses). If the obtained value is outside a 40% variation range, we replace the estimated cross section by a linearly interpolated one (circle or ellipse).

5.Cross section post-processing
Diagram
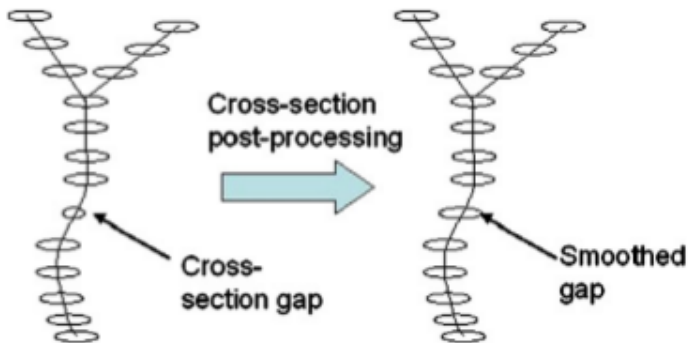
Fig. 4. Post processing of the cross sections. Each gap and each consequent geometrical variation are smoothed out.

Outline

1 Introduction

2 Description of the method
  - Segmentation of the vessels
  - Surface reconstruction

3 Experiments
  - Phantom
  - Clinical data

Improvements

- Requirement: need to supply real-time interventional radiology simulations.

1. General tree structures: existing methods were limited to bifurcations.
2. Improved trunk vessel selection: considers both branching angle and vessel cross section areas in order to reduce under sampling artifacts.
3. Adaptive cross section distribution: considers both vessel's cross section area profiles and centerline curvature profiles.
4. New bifurcation tiling scheme: allows holes in the surface at the reconstruction to be avoided.

1.General tree structures

- Each vessel is discretized into segments.
- A segment is formed with two circular or elliptical end cross sections and a line segment connecting the two.
- Approximating these cross sections by square or diamond polygons creates the base mesh.
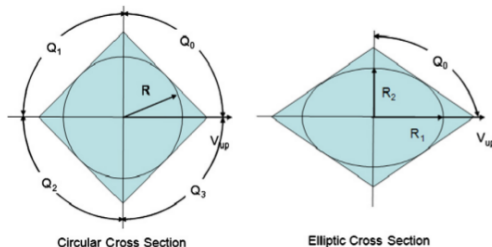
Fig. 5. Illustration of **up** for both circular cross section and elliptical cross section cases. Four quadrants, Q_0 to Q_3, are ordered in both cases, creating the base mesh (either a square or a diamond) for each cross section.

1.General tree structures

The procedure initially patches segments in the bifurcation regions, and then the other segments between two end segments of each vessel.

The patching procedure consists of three tasks:

1. tiling the surface from the second segment of each vessel's centerline to the one preceding the last segment, assuming that the first and last segments have been tiled during the bifurcations patching;
2. tiling the joint by patching the end segment of the trunk vessel and the beginning segments of other branches that share this joint;
3. finishing by recursive calls of itself to the parent and children branches of the reference vessel.

2.Improved trunk vessel selection
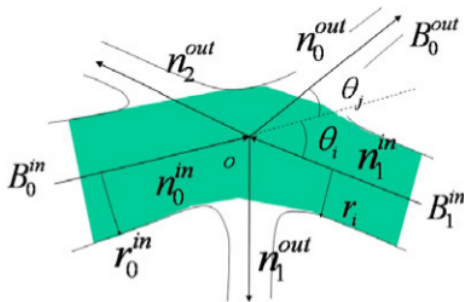
Fig. 6. Trunk vessel selection: using both vascular structures average radii and branching angle to determine the continuation trunk of current vessel, B_0^{in}. Although $\theta_j < \theta_i$, our algorithm chooses B_1^{in} as the trunk vessel of B_0^{in}, due to the similarity of their average radii.
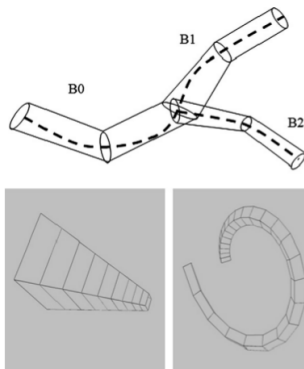
3.Adaptive cross section distribution

Fig. 7. Top: Vessel B_0, B_1, and B_2 are sampled under the same rate. While the digitized B_2 still resembles much of its original geometry, the discretized version of B_0 and B_1 loses critical curvature details. Bottom left: cross section distribution is denser at thinner regions of a vascular structure. Cross sections are farther apart as the vascular structure diameter increases. Bottom right: the distribution density is higher where a vascular structure turns or twists. Few cross sections are placed where the centerline curve is flatter.

2.New bifurcation tiling scheme
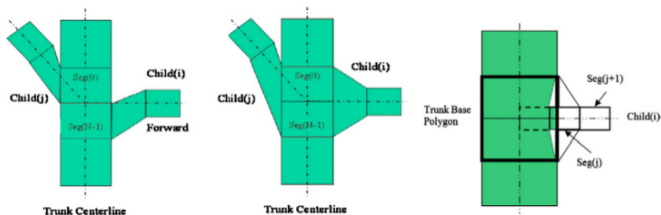
Fig. 8. Left: Depending on the outgoing vascular structure's orientation, Felkel's method, on the left, connects $Child(j)$ to $Seg(0)$ and $Child(i)$ to $Seg(N-1)$. Our method, on the right, connects $Child(i)$ and $Child(j)$ both $Seg(N-1)$ and $Seg(0)$. The bottle-neck effect is reduced and skeleton symmetry is preserved. Right: Additional intersection tests between segment polygons are conducted. Starting from the bifurcation, each segment of $Child(i)$ branch is sequentially examined against the trunk's base polygon. The first segment from the bifurcation node that does not intersect with and is not enclosed by the trunk base polygon will be used to connect with the trunk. In this example, if we use the left cross section of $Seg(j)$ to connect with the trunk as drawn in the dotted line, the resulting surface mesh will have an artifact of eroding into the trunk. To avoid this, we continue such intersection tests and select the left cross section of $Seg(j+1)$ of $Child(i)$ as the candidate for connection drawn as a solid line.
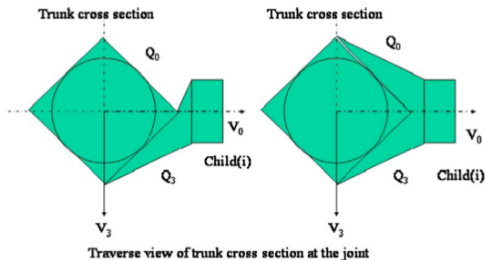
2.New bifurcation tiling scheme

Fig. 9. *Child(i)* centerline lies close to the boundary of Q_0 and Q_3. Using only one quadrant Q_3 induces an unwanted twisting artifact, on the left. Adjacent-quadrant-grouping uses both Q_0 and Q_3 to connect *Child(i)*'s base mesh to the trunk surface, on the right.
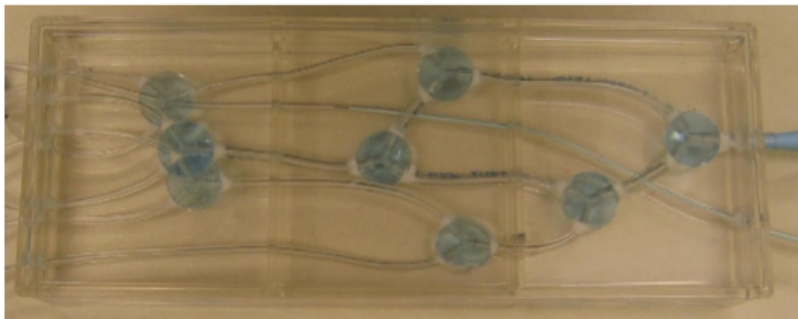
Outline

Outline

Phantom

Fig. 10. The phantom with its nylon tubing embedded in silicon.

Results
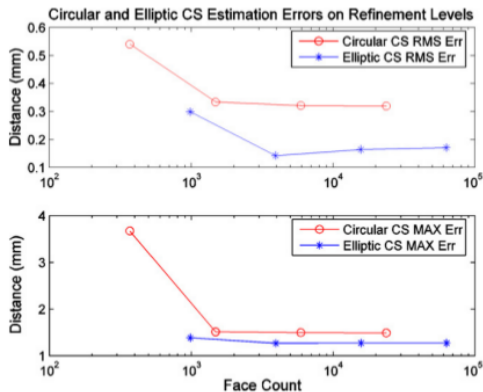
Fig. 12. Elliptic cross section based vascular surface patching consistently improves the surface accuracy over circular cross section based technique in both RMS and maximum error tests.

Outline

Computed tomography angiogram (CTA)

A first evaluation of a clinical case has been performed in Luboz et al. (2005). Our method was applied to the basilar section of the brain vessels. This partial network is quite representative of the brain vasculature since it is composed of vessels with various sizes and shapes. A distance of 0.4 mm or less is noted between most of the vertices of the surface reconstructed by our method and the one obtained with the level set (Luboz et al., 2005). Given the resolution of the CTA dataset ($0.48 \times 0.48 \times 1$ mm for a voxel), it shows little difference between the two surfaces. Furthermore, our method has 11 times fewer polygons than the reference surface and their aspect is visually smoother. The goal of our method
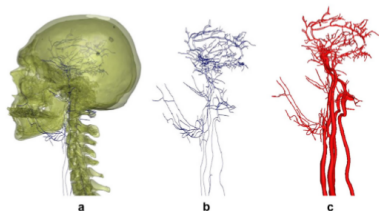
Results

Fig. 13. From left to right: (a) iso-surface of the skull obtained from a threshold of the CTA (b) skeleton generated from our method, and some manual work, is seen in transparency, vessels skeleton from the patient CTA, and (c) 3D surface resulting from our method.
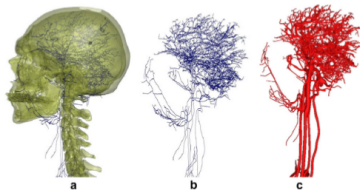

Fig. 14. From left to right: (a) iso-surface of the skull obtained from a threshold of the CTA, (b) skeleton generated from our method, and considerable manual work, is seen in transparency, vessels skeleton from the patient CTA, and (c) 3D surface resulting from our method.
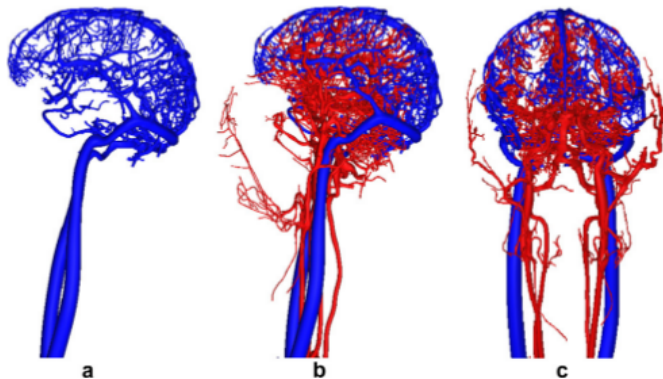
Results

Fig. 15. From left to right: (a) venous network reconstructed from a patient CTA, (b) and (c) arterial and venous sides from the same patient CTA.

Magnetic Resonance Angiography (MRA)

Our method has also been applied to a Magnetic Resonance Angiography (MRA) of the same patient's head and neck. Though less used than CTAs in the neuro-interventional setting, because of a poorer resolution, MRAs have the advantage of presenting a wider grey level repartition for the pixel intensity for the vessels. In the end, this leads to an easier segmentation of the vascular structure and therefore fewer user interactions to connect the skeleton. The drawback of using MRA is that fewer small vessels are visible. This medium is therefore less interesting for our simulator but can be used for the main vessels diagnosis.
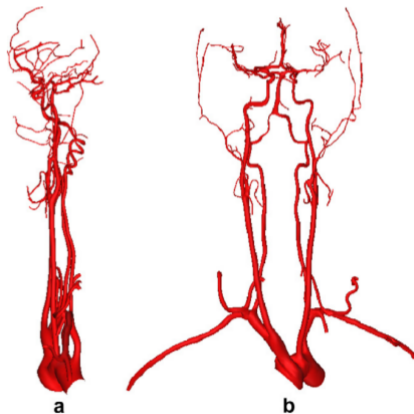
Results

Fig. 16. (a) Sagittal view of the reconstructed surface of the arteries from a patient MRA. (b) Frontal view of the same network.