

Symbolic description of intracerebral vessels segmented from magnetic resonance angiograms and evaluation by comparison with X-ray angiograms

Elizabeth Bullitt*, Stephen Aylward, Keith Smith, Suresh Mukherji, Michael Jiroutek, Keith Muller

Medical Image Display and Analysis Group, University of North Carolina, Chapel Hill, NC, USA

Received 1 June 2000; received in revised form 18 October 2000; accepted 15 December 2000

Abstract

We describe and evaluate methods that create detailed vessel trees by linking vessels that have been segmented from magnetic resonance angiograms (MRA). The tree-definition process can automatically exclude erroneous vessel segmentations. The parent–child connectivity information provided by our vessel trees is important to both surgical planning and to guidance of endovascular procedures. We evaluated the branch connection accuracy of our 3D vessel trees by asking two neuroradiologists to evaluate 140 parent–child connections comprising seven vascular trees against 17 digital subtraction angiography (DSA) views. Each reviewer rated each connection as (1) Correct, (2) Incorrect, (3) Partially correct (a minor error without clinical significance), or (4) Indeterminate. Analysis was summarized for each evaluator by calculating 95% confidence intervals for both the proportion completely correct and the proportion clinically acceptable (completely or partially correct). In order to protect the overall Type I error rate, alpha-splitting was done using a top down strategy. We additionally evaluated segmentation completeness by examining each slice in 11 MRA datasets in order to determine unlabeled vessels identifiable in cross-section following segmentation. Results indicate that only one vascular parent–child connection was judged incorrect by both reviewers. MRA segmentations appeared complete within MRA resolution limits. We conclude that our methods permit creation of detailed vascular trees from segmented 3D image data. We review the literature and compare other approaches to our own. We provide examples of clinically useful visualizations enabled by our methodology and taken from a visualization program now in clinical use. © 2001 Elsevier Science B.V. All rights reserved.

Keywords: Magnetic resonance angiography; Digital subtraction angiography; Registration; Segmentation

1. Introduction

It would be helpful to both neurosurgeons and neuro-radiologists to visualize intracerebral vessels in 3D with information about vascular parent–child relationships. The treatment of skull base tumors is one of several areas in which such descriptions would be useful. Open resection of these lesions is often lengthy, difficult, bloody and

carries a high risk of patient injury. Endovascular embolization is similarly risky, with major complications occurring in 5–20% of patients (Purdy et al., 1990; Higashida et al., 1991; Kupersmith and Berenstein, 1993; Vinuela, 1992).

A large part of the problem is that treatment of such tumors requires interruption of their blood supplies. However, interruption of blood flow to normal brain may produce stroke. As occlusion of a parent vessel will deprive its children of blood, it is critical for the surgeon to visualize vascular parent–child relationships in order to determine which vessels can be safely occluded. Fig. 1 illustrates the clinical problem.

*Corresponding author. Division of Neurosurgery, CB #7060, 148 Burnett-Womack, University of North Carolina-CH, Chapel Hill, NC 27599, USA. Tel.: +1-919-966-1374; fax: +1-919-966-6627.

E-mail address: bullitt@med.unc.edu (E. Bullitt).

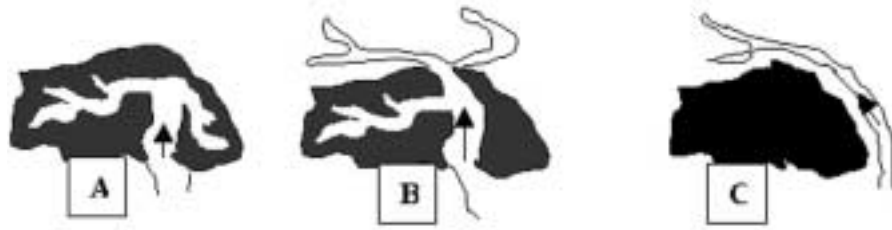


Fig. 1. Vessel types. The black area in A–C represents a solid tumor. The cross indicates the proposed site of vessel occlusion. The arrow shows direction of blood flow. (A) An artery supplying a tumor. (B) An artery that supplies a tumor but that also passes beyond it to supply normal brain. (C) An artery that passes over a tumor but that, from some projection angles, may appear to pass through it.

No current imaging modality provides the required 3D connectivity information. Digital subtraction angiography (DSA), magnetic resonance angiography (MRA), and 3D-DSA are all greyscale images. Furthermore, projection images of the complex, plethoric intracerebral circulation contain so much projection overlap that it is difficult to determine child–parent associations. Without 3D vascular connectivity information, it is impossible to simulate motion of a catheter through a vascular tree, to estimate collateral flow to an anatomical region, or to provide an automated description of vessels by the kind of categorization shown in Fig. 1. The creation of detailed, patient-specific, 3D vascular trees would thus be of great value for surgical planning, as it is often difficult to determine intraoperatively whether a vessel entering a tumor possesses child branches that pass through the tumor to supply normal brain. Similarly, such trees would be of value for guidance of endovascular procedures, which require manipulation of a catheter through a complex maze of vessels whose connections are often obscured on projection views. The purpose of this paper is to describe and evaluate a method of linking disconnected vessels segmented from MRA into vessel trees that possess vascular parent–child connectivity information. Evaluation of the accuracy of the created trees is done through comparison with sequences of DSA images obtained from the same patient. No special image acquisition techniques are required. All programs are cross-platform and run well on home PCs under Windows (Bullitt et al., 1997a). Our methods are now in active clinical use, and a clinical paper illustrating clinically useful visualizations is in press (Bullitt et al., 2000). An overview of our approach and a preliminary evaluation of a single vessel tree were presented at IPMI '99 (Bullitt et al., 1999a).

2. Methods

The approach involves (1) segmentation of vessels from MRA data, (2) definition of vascular trees and (3) evaluation of results against DSA data obtained from the same patient.

2.1. Segmentation

The MRA segmentation method was developed by Dr. Aylward (Aylward et al., 1996). Extraction of a vessel involves three steps: definition of a seed point, automatic extraction of an image intensity ridge representing the vessel's central skeleton, and automatic determination of width at each skeleton point. The approach is applicable to extraction of tubular objects from any type of 3D image data. For this report, we employed three-dimensional, time of flight MRA, which was performed in a Siemens 1.5 T Vision unit with a quadrature head coil. Off resonance (1500 Hz) magnetization transfer suppression (75 Hz) was used to further suppress background stationary tissue. Images were acquired in the axial projection using 1 mm contiguous partitions. MRA acquisition parameters included repetition time: 38 ms; echo time: 7 ms; field of view: $200 \times 150 \text{ mm}^2$; flip angle: 20° ; signal averages: 1; using a 192×256 matrix with superior venous presaturation pulse. The voxel size was $0.78 \times 0.78 \times 1 \text{ mm}^3$.

2.1.1. Seed points

Extraction of each vessel begins with definition of a seed point on the image data. One method of supplying seed points is through user interaction. This approach is useful in extraction of unusually wide vessels such as aneurysms, or when a clinician wants to visualize only a subset of vessels. The extraction program presents the user a set of MRA slices, and the user can click on a point within a vessel, simultaneously supplying a rough estimate of that vessel's width.

The number of seed points required depends on the underlying anatomy, on the quality and volume of the MRA, and on the clinical problem of interest. For normal patients, full extraction of a single carotid circulation requires between 15 and 60 seed points. For patients with huge vascular malformations supplied by both carotids and the vertebrobasilar circulation, hundreds of seed points may be required. For aneurysm patients in which only local anatomy is of interest, five to six seed points may be sufficient.

More recently, we have begun to develop an alternative, fully automated method of supplying seed points that

employs intensity windowing to provide such seeds. Of the seven trees analyzed in this report, the first four were created from vessels extracted with manually supplied seed points and two of the last three were created from a set of vessels whose seed points were automatically determined. The automated method is under evaluation, will be the topic of a future paper, and is not discussed further here.

Regardless of the method of seedpoint definition, the program then queries if the identified point lies on or close to an image intensity ridge. If the identified point does not lie close to an image intensity ridge or if the ridge is less than a voxel in length, the seed point is rejected.

2.1.2. Ridge tracking

Once a seed point is accepted, the method automatically extracts the central skeleton of the indicated vessel. Vessels can be viewed as 3D tubular objects delineated from background by contrast differences. This combination of geometry and intensity means that blurring the image creates a central intensity ridge along each vessel. This intensity ridge is extracted via the height ridge definition:

Define:

$I \in \mathfrak{N}$ as the intensity at $x \in \mathfrak{N}^3$;

$H \in \mathfrak{N}^{3 \times 3}$ as the Hessian of I at x ;

$v_i \in \mathfrak{N}^3$ and $a_i \in \mathfrak{N}$ the eigenvectors and associated eigenvalues of H where $a_1 < a_2 < a_3$.

Then for the program to classify x as being on a ridge, it must be true that

$$a_2 / (a_1^2 + a_2^2)^{1/2} \sim < -0.5, \quad (1)$$

$$v_1 \cdot \nabla I \sim = 0 \quad \text{and} \quad v_2 \cdot \nabla I \sim = 0. \quad (2)$$

Eq. (1) states that the object's cross-section as captured by the intensity curvature about x (i.e. the Hessian H) should be nearly symmetric and negative, corresponding to a ridge rather than to a valley.

From a single ridge starting point, the tracking process proceeds in both directions along a ridge to extract the tube's full extent. Extracted point ordering is maintained to produce a sampled representation of a single, continuous tube. Diagrams illustrating ridge definition are available in (Aylward et al., 1996).

One common concern about this approach is its ability to pass branch points, i.e. singularities at which the eigenvectors of the Hessian become ill defined/may rotate, thus producing ambiguity of the tangent and normal directions. However, rotation of the eigenvectors will only occur at branch points if the ridges of two or more of the branches have the same height. This situation is rare in medical images. As a result, a single seed extracts a vessel that usually tracks past multiple branch points and courses through many dozen (and, for lengthy vessels that dip and curve, sometimes over 100) voxels.

2.1.3. Width estimation

The width of the vessel is automatically estimated at points along its central skeleton. The method takes advantage of the fact that vessels have nearly circular cross-sections. The width of a tube about a central skeleton point is proportional to the scale that produces a maximal response from a cylindrical medialness measure. Earlier extraction methods employed a Laplacian of a Gaussian kernel (LoG) (Aylward et al., 1996). However, Bryan Morse developed a medialness measure that is an improvement over the LoG measure in that it only makes measure at concentric rings just inside of and just outside of the radius being tested (Morse, 1994). As a result, Morse medialness, compared to LoG, is faster and less sensitive to intensity variations within the object whose width is being tested. All extractions described in this report were performed using a version of the Morse medialness measure.

Once a vessel is extracted, the program marks the extracted region so that later segmentation of another vessel will not include any territory already defined by the first. If segmentation of a later vessel attempts to extend into already claimed territory, the segmentation stops.

2.2. MRA/DSA registration

Our methods of tree editing and evaluation compare the 3D tree to DSA data obtained from the same patient. We therefore briefly describe our 3D/2D registration method.

Registration of segmented vessels with DSA is done as described by Liu et al. (1998). The 3D/2D registration process uses as primitives four to eight 2D curves extracted from the DSA and an equivalent number of 3D curves extracted from the MRA. The program then optimizes a viewplane based disparity measure based on the iterative closest point paradigm between the DSA skeletons and the projections of the MRA skeletons. When compared to 3D/2D registration done by manual methods, the approach is both faster and more accurate (Bullitt et al., 1999b).

It should be noted that image distortion errors may preclude perfect 3D/2D registration. Although we use phantoms to correct for DSA pincushion errors and for the machine-specific errors of our MR machine, we cannot correct for MR geometrical distortions in the region of an air-tissue interface. Vessels extracted from such regions of the MR are deformed, and cannot be perfectly registered with DSA. However, the current study does not require perfect registration, but only that the registration be sufficiently good to allow a user to associate the projection of each 3D vessel with its DSA counterpart. This paper contains multiple figures of 3D vessel skeletons projected upon DSA data. As shown by these images, there is little difficulty in associating appropriately paired vessel projections despite imperfect registration.

2.3. Definition of vascular trees and identification of branchpoints

The output of the segmentation program is a set of disconnected, directed, 3D skeleton curves with an associated radius at each point. In order to create meaningful vascular trees it is necessary to link these vessels together, to define branchpoint connections, and to eliminate extraneous segmented ‘vessels’ actually representing noise. We have developed an automated post-processing method that performs these three tasks.

Before discussing the tree definition algorithm, it is necessary to make several points about the individual vessels produced by the segmentation program.

First, although each segmented vessel has a direction, this direction is determined by the extraction process and may not correctly model the direction of blood flow. It is therefore necessary for the tree-definition process to reverse the direction of some vessels and, in the case of some ‘Y’ connections, to break a child vessel into two parts (Bullitt et al., 1999a).

Second, MRA datasets are noisy. For our images, typically acquired with a voxel size of $0.8 \times 0.8 \times 1.0$ mm, vessels with diameters of 1 mm or less are very faint in the image data and are difficult to discriminate from noise. As endovascular procedures may involve catheterization of 1 mm vessels, it is necessary to include even these small vessels in the final tree. However, any intensity-based extraction method attempting to include small, faint vessels will also extract noise. In addition, our automated extraction method may pick up noise in the skull, orbits and skull base, as well as unwanted vessels belonging to the external carotid circulation and pathological lesions of high image intensity.

We use a post-processing program not only to define branchpoints and to identify parent–child connections, but also to automatically exclude noise. The approach is based upon a variant of the minimum spanning tree algorithm. The user first selects one or more roots by clicking on a vessel in a projection image. The program will then automatically define a vessel tree by progressively attaching children to the connected base, using the cost function outlined below. End-to-end ‘I’ connections are permitted, as are ‘Y’ connections (bifurcations) that involve at least one vessel endpoint. As each child is added, its direction of flow is corrected as necessary. ‘X’ connections (trifurcations) are not permitted as the intracerebral vasculature does not contain trifurcations. This exclusion results in automatic elimination of almost all noise as such connections must not only meet the minimum requirements of the cost function but must also involve a piece of noise with an endpoint very close to the proposed parent. Such situations are rare.

The cost function for adding a proposed child is based both upon Euclidean distance and the image intensity data. As previously noted, trifurcations are not permitted. The

distance between a pair of vessels is thus calculated by defining the minimum distance between each orphan vessel’s two endpoints and the closest point along the proposed parent’s skeleton curve, and the minimum distance between the unattached end of the proposed parent and the closest point along the proposed child. Each of the three possible connections thus involves at least one vessel endpoint. This ‘distance’ between the vessel pair is then corrected to allow for the radii of both parent and child. As discussed later, our segmentations are generally complete. When the radii of the parent and child vessels are taken into account, the distance between each true parent and child is usually about one voxel.

The cost function also examines the image intensity data for evidence of connection. As the gap between parent and child is generally in the order of a voxel, the connection can be represented as a straight cylinder. A line is drawn on the image between the proposed connection points, and the average image intensity of this line is calculated, as is the average image intensity of a set of concentric cylinders constructed around this line, with the inner cylinder possessing a radius just larger than that of the proposed child. The two image intensity values are then expressed as a ratio, with a low value indicating a high likelihood of connection as higher image intensity is expected along the connecting line than in the background.

The final cost of any proposed connection is a composite of the two numerical values ‘distance’ (measured in centimeters) and ‘image intensity ratio’ (a dimensionless number). How best to combine and weight these two parameters is unclear. We have experimented with a variety of weighting functions and have concluded only that, while both parameters are important, the ‘image intensity ratio’ should be weighted at least as strongly as ‘distance’. For all vessel trees discussed in this paper we therefore, and admittedly heuristically, added the two values together. If the final composite number is low, there is a high probability of vessel connection.

Both the maximum allowable distance and maximum allowable intensity ratio can be set interactively. Connections that require either larger distances or larger intensity ratios are rejected. All vessel trees discussed in this paper were created using a maximum allowable connection distance of 0.3 cm (approximately three voxels), and a maximum allowable intensity ratio of 0.7.

2.4. Editing vascular trees

The intracranial arterial supply normally consists of three circulations: the posterior and the two carotid circulations. These circulations are normally interconnected at the base of the skull through the Circle of Willis. Flow direction within the Circle of Willis is variable from patient to patient, however.

Our methods of segmentation are based upon geometry alone. As a result, we do not know flow direction in the

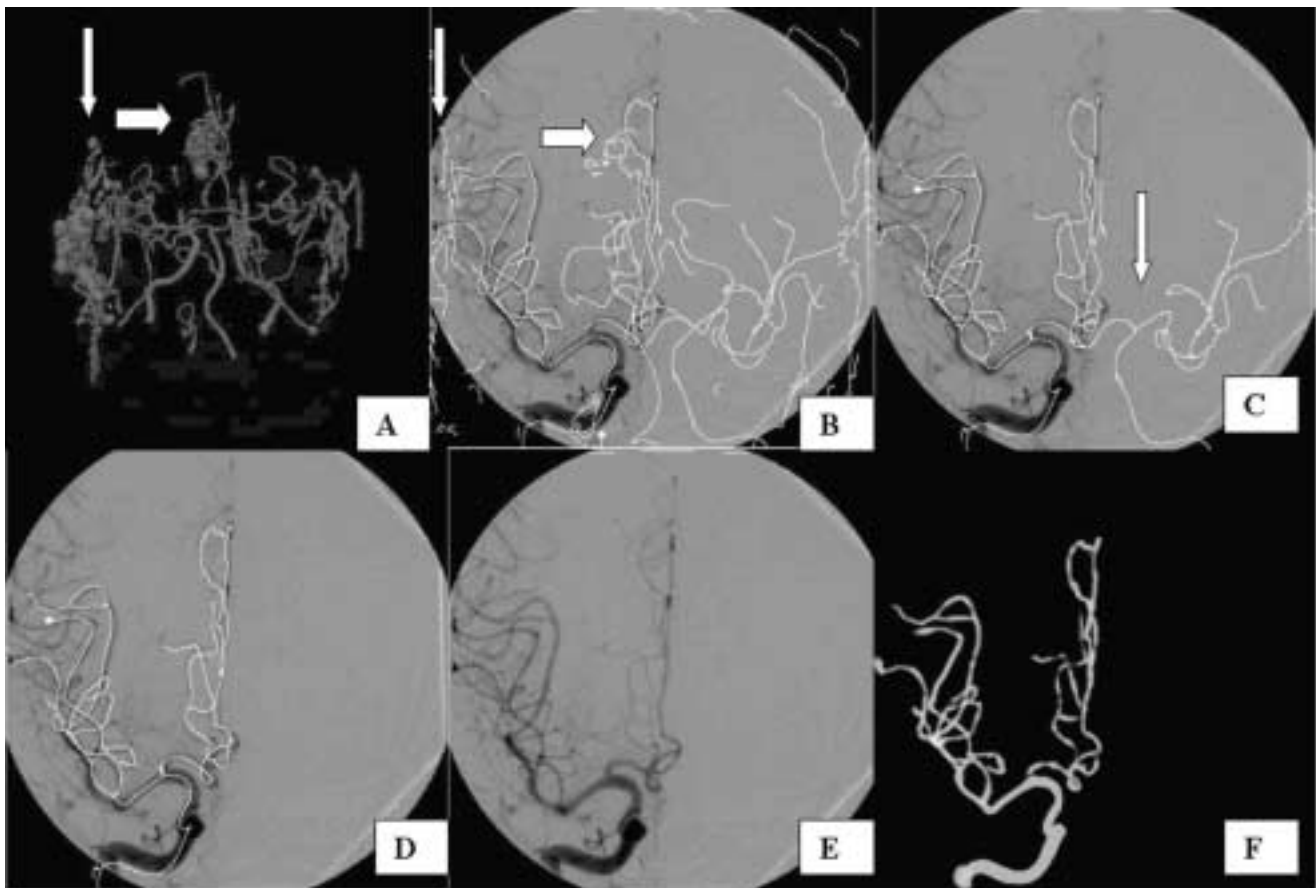


Fig. 2. Steps in tree definition following automatic vessel extraction. (A) Surface rendered AP projection of vessels automatically extracted from an MRA of a patient with a large blood clot. The segmentation includes noise and vessels in the scalp and dura (vertical arrow). The clot has also been segmented (horizontal arrow). (B) Projection of the skeletons of the extracted objects upon a registered DSA. Arrows point to the same 'sheet' of noise and to the clot shown in A. (C) The vessels that remain following tree definition with the right carotid as root. The noise and blood clot are gone but connections remain to the left carotid circulation (arrow). (D) Results following manual editing. The left carotid subtree has been removed by a single 'point and click', as has one anterior cerebral artery and its descendants. (E) DSA without superimposed vessel skeleton projections. (F) 'Pseudoangiographic' projection of the vessels whose skeletons are shown in D. The tree contains 15 branches, each extracted by one seed point.

Circle of Willis. The tree program therefore contains a set of 'point and click' editing tools that can be used, in combination with information obtained from DSA, to exclude subtrees or to redefine the parent of any tree or subtree. Fig. 2 illustrates the steps involved in creating a vascular tree using additional DSA data. In brief, these steps include:

- (a) Segmentation of vessels from MRA. The segmentation will contain noise (Fig. 2(A)).
- (b) Registration of the segmented vessels with a DSA (Fig. 2(B)).
- (c) Identification of one or more roots followed by automated tree definition. This step will also automatically exclude the large majority of extraneous 'vessels' actually representing noise (Fig. 2(C)).
- (d) 'Point and click' removal of undesired trees, such as the contralateral carotid circulation. As vascular trees have now been defined, a tree or subtree that may contain dozens of vessels can be removed by a single mouseclick (Fig. 2(D)).

The editing feature and the use of DSA data are optional. If desired, the user can simply identify arterial roots and request that trees be created, resulting in a user interaction time of a few seconds. If editing is desired to redefine connections in the Circle of Willis, the processing time will be increased by 5 min or so.

2.5. Evaluation

We have evaluated our methods in three ways. The first two approaches involve comparison of vessel trees with DSA data. The third provides an estimate of the completeness of the segmentation method.

2.5.1. Qualitative comparison of vessel trees with DSA images

One can obtain a qualitative measure of the accuracy of our vessel trees by registering each vessel tree with a DSA obtained from the same patient, and then projecting

the tree using a ‘pseudoangiographic’ projection. A quickly rendered approximation of a DSA view can be created by projecting each 3D vessel skeleton point, calculating the width of the projected point, and drawing a filled circle of image intensity ‘ x ’, such that the final intensity of any pixel within that circle is the sum of ‘ x ’ and a proportion of that pixel’s intensity before the circle was added. As is true for DSA images during the arterial fill sequence, the final image will then define larger vessels as brighter and will exhibit increased image intensity at sites of crossing. Fig. 2(F) provides an example of a pseudoangiographic projection.

A side-by-side comparison of the original DSA with the pseudoangiographic projection of a vessel tree provides a subjective comparison of the accuracy and completeness of our vessel trees. Such comparison also provides a sensitive means of evaluating the width determinations made by the segmentation program, since small errors in width estimates are magnified on projection.

2.5.2. Quantitative evaluation of parent–child connections

We quantitatively evaluated the accuracy of 140 parent–child connections in seven vascular trees by asking two neuroradiologists (K.S. and S.M.) to independently evaluate each 3D parent–child connection against the DSA information available for that patient. A third reviewer additionally evaluated the first vessel tree, but then left our institution and so was not available for later evaluations. For reasons of statistical completeness, we consider only the seven pairs of evaluations done by the same two reviewers. The reviewer who provided a third evaluation of the first tree judged no connection as incorrect.

At our hospital, each DSA series consists of a timed sequence of images obtained from the same point of view. From each DSA series, one to two images were selected that best showed the arterial fill phase. The selected images were then grouped by patient and by the circulation shown. We use the term ‘view’ to indicate a timed sequence of images taken from the same angle.

Each neuroradiologist was asked to evaluate the correctness of each parent–child connection by using a program that projected the skeleton of a vessel tree upon each DSA view of that circulation. Each vessel tree was compared to DSAs taken from two to four angles of view, with one to two images in the fill sequence available for each view. The user could step through each parent–child connection one at a time, and could step through each available DSA image. The program would display only that child, its parent, and the connection between the two each in a different color and as superimposed upon each DSA. Each reviewer was asked to rate each connection by a four point scale:

1. ‘Correct’: The proposed connection is clearly correct on

at least one view and does not appear incorrect or partially correct on any view,

2. ‘?’: Projection overlap makes it impossible to determine if the connection is correct or incorrect on any view,
3. ‘Incorrect’: The connection is clearly incorrect on at least one view, resulting in a clinically significant error, and
4. ‘Partially correct’: The connection is not incorrect but at least one view shows a minor error not of clinical significance (e.g. the connection point is slightly displaced on the correct parent).

The results include evaluation of 140 parent–child connections comprising seven different, non-overlapping vessel trees and evaluated against a total of 17 different DSA views obtained from three patients. Some of these 17 views (viewing angles) allowed selection of more than one image in the associated DSA fill sequence. Three trees had the left carotid as root, two trees had the right carotid as root, and two trees had a vertebral or basilar artery as root.

Results were analyzed statistically by calculating the 95% confidence interval for each reviewer both for the proportion of parent–child connections entirely correct and for the proportion of vessels meeting the demands of clinical utility (correct or partially correct). Each tree was viewed as independent of all other trees. For each reviewer’s analysis of each tree, any connection judged as indeterminate from the available DSA data was eliminated from that analysis. In order to protect the overall Type I error rate, alpha-splitting was done using a top down strategy. With a total of seven trees each reviewed twice, there were a total of 14 tree-specific confidence intervals, which we computed at the nominal level of $0.05/14 = 0.0036$.

2.5.3. Completeness of segmentation

As previously noted, it is desirable to define vessel trees that contain even the smallest vessels discernible from MRA. It is therefore important to assess segmentation completeness.

We have evaluated the completeness of segmentation through manual inspection of MRA slice data. During vessel extraction, the segmentation program colors each extracted tube and thus distinguishes segmented objects from the greyscale MRA slices. Bright curvilinear lines that remain uncolored in the MRA slice data are presumably unextracted vessels. One of us (E.B.) manually segmented 11 MRA datasets, providing seed points for every bright (or even slightly bright) pixel not yet included in the segmentation and that the user thought might possibly belong to an as yet unextracted vessel. The segmentation program’s refusal to accept such seed points can be viewed as indicative of program failure. An analysis of such failures can define where the segmentation program appears to succeed and to define conditions under which it may fail.

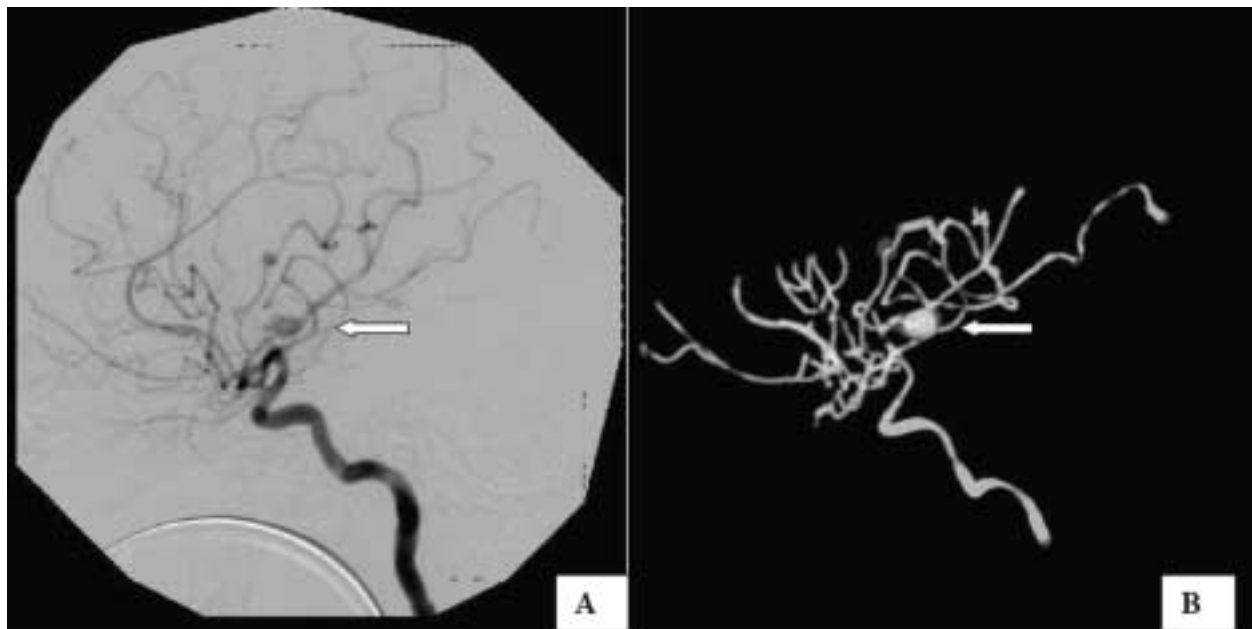


Fig. 3. Qualitative comparison of a vessel tree with an associated DSA image. The patient has a fusiform middle cerebral artery aneurysm. The arrow in each frame points to the aneurysm. The most distal branches are not included since the MRA did not cover the entire head. (A) DSA. (B) Pseudoangiographic projection of the vessel tree. This tree was built from 54 seed points.

3. Results

3.1. Qualitative comparison of vessel trees with DSA images

Fig. 2 shows a pseudoangiographic projection of a vessel tree adjacent to a DSA image of the same circulation. Note that the vessel widths are largely correct in the vessel tree projection even though this patient has a blood clot immediately adjacent to the anterior cerebral group (Fig. 2(A)), making segmentation potentially difficult.

Fig. 3 shows a different vessel tree, defined from a patient with a fusiform middle cerebral artery aneurysm. The aneurysm was manually extracted from the MRA. All other vessels were automatically extracted. Note the level of vascular detail and the range of widths represented.

3.2. Evaluation of parent–child connections

It is difficult for even expert neuroradiologists to define parent–child connections from projection views. This difficulty is mirrored by the different responses of the two neuroradiologists. Although each radiologist faulted or questioned several of the 140 connections, only one parent–child connection was mutually agreed by both radiologists to be incorrect. Table 1 shows the connection ID numbers by tree that each radiologist felt to be incorrect, partially correct or indeterminate.

For statistical analysis, any connection a reviewer viewed as indeterminate was removed from that reviewer’s database. In the analysis in which partially correct connections (connections exhibiting an error of no clinical significance) were viewed as fully incorrect, reviewer A

Table 1
Connection ID numbers by tree and radiologist for each connection judged to be incorrect (IC), partially correct (PC), or indeterminate (?)^a

	Tree 1		Tree 2		Tree 3		Tree 4		Tree 5		Tree 6		Tree 7	
No. vessels	25		23		22		24		15		23		15	
Type	L carotid		L carotid		Basilar		R carotid		L carotid		R carotid		Basilar	
IC	–	–	9	–	9, 10	9, 12	22	–	15	–	–	–	9	–
PC	–	–	–	7	–	13	–	–	9	–	–	–	–	–
?	12, 21	6	–	14	3	–	18	–	–	15	–	21	–	–

^a The IDs given by reviewer 1 are shown on the left of each cell, those given by reviewer 2 on the right. Only one connection (ID 9 within tree 3) was agreed by both reviewers as incorrect.

Table 2
Agreement between radiologists and agreement between each radiologist and our program^a

		Reviewer 1	
		Correct	IC/PC
Reviewer 2	Correct	124	4
	IC/PC	4	1

^a Any connection made by our program and viewed as indeterminate by at least one radiologist is excluded from analysis. Any connection made by our program and viewed as partially correct by any radiologist was classified as fully incorrect. The results show that 124 connections were agreed as correct by both radiologists when compared to our program. One connection made by our program was viewed as incorrect by both radiologists. Each radiologist additionally faulted four connections, each of which was viewed as fully correct by the other radiologist.

judged 94.9% to be correct with a 95% CI of (89.7, 97.9) and reviewer B judged 97.1% to be correct with a 95% CI of (92.6, 99.2). In the ‘clinically acceptable’ analysis in which partially correct connections were judged as correct, reviewer A judged 95.6% to be clinically acceptable to produce a 95% confidence interval of (90.6, 98.4), and reviewer B judged 98.5% to be acceptable to produce a 95% confidence interval of (94.8, 99.8). These findings suggest that our results, while sometimes faulted by either radiologist, can achieve a high level of accuracy.

It should be noted, however, that although we have taken the radiologists’ responses as a ‘goldish standard’ of truth, the two radiologists often did not agree. There is, in fact, no absolute standard of truth. We know of no sufficiently complex phantom upon which to test our methods, and the elegant MNI Brain-web approach, which generates realistic MR images for purposes of validating

image analysis methods (Kwan et al., 1999), does not cover validation of vascular segmentation.

In the absence of an absolute standard of ‘truth’, we asked two radiologists to independently review our findings against the available DSA data for each patient. Under these conditions it would be desirable to show that our program agreed as well with either radiologist as the two radiologists agreed with each other. Formal statistical comparison is not possible, as such analysis would require each radiologist to create trees of his own – a formidably difficult and time-consuming task from complex intracerebral vascular data. However, Table 2 presents the agreement/disagreement between radiologists and between each radiologist and our program. In this table, any connection viewed by either one of the radiologists as indeterminate was eliminated from analysis. For the remaining connections, both radiologists agreed that 124 connections shown by our program were correct, each radiologist judged a different four connections to be either incorrect or partially correct (with the other radiologist judging the same four connections to be fully correct), and only one connection made by the program was judged by both radiologists to be mistaken.

3.3. Completeness of segmentation

Several hundred ‘vessels’, most of which represented noise, were manually extracted from each of 11 MRAs during a test in which the user attempted to click on any even slightly bright pixel that might represent an as yet unsegmented vessel. Fig. 4 shows a randomly chosen, typical MR section in which the user had attempted to extract all possible vessels from one side of the head.

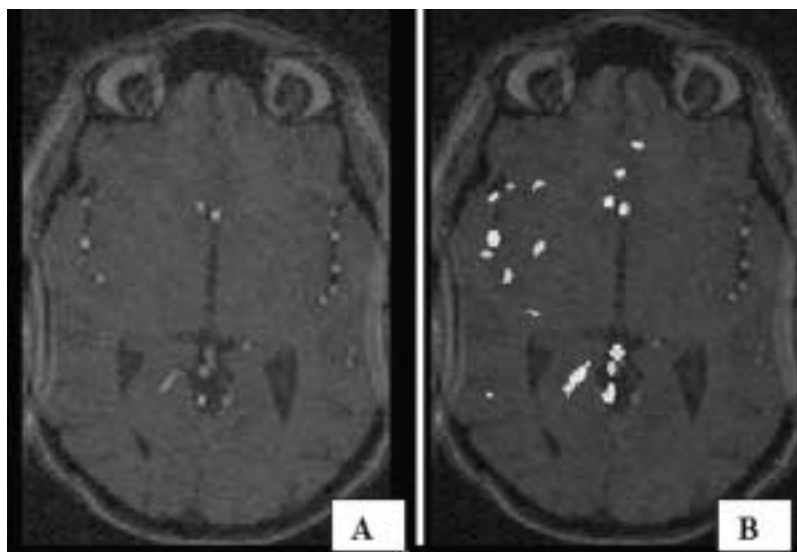


Fig. 4. Manual MRA segmentation. A user has attempted to extract anything that might possibly be a vessel from one side of the head. (A) Typical MRA section. (B) The same section with extracted regions shown in white. Note that all objects that are clearly vessels have been appropriately defined. Additional objects that might be small vessels but are probably noise have also been defined.

There were two situations in which attempted segmentation of a tubular structure failed. The first is that of segmentation in the top or bottom MRA slice. The segmentation protocol seeks a 3D, image intensity ridge and there is insufficient bounding data to support a ridge at the data boundary points. As a result of this analysis, we have therefore implemented a new data boundary handling method that extends the recorded data using dim replications of the boundary slices (dimness increases as a function of the distance from the boundary slice). In this manner, enough data is provided to extract vessels that run along the boundary slices.

The second situation occurs when two vessels course immediately adjacent to each other. MRA resolution is limited by voxel size. If the MRA representation of two adjacent vessels merges the two into a single, contiguous group of pixels, the two vessels cannot be segmented independently. This problem occurred for the paired anterior cerebral arteries in at least one focal location in several patients, and was fixed by manual editing prior to review by the radiologists.

With these two exceptions, the segmentation protocol appeared to be exquisitely sensitive to the detection of tubular objects. No reasonable object identified by the user failed to be accepted by the segmentation program. Additional true vessels, difficult for the user to discern, were reliably extracted, probably by extension of a 3D vessel initially identified on a different slice. A specific example is the anterior communicating artery, which is adjacent to the noisy skull base and contains turbulent flow. Segmentation of even this difficult vessel was performed without difficulty in all 11 cases. With the exception of ‘merged’ vessels produced by limited MRA resolution, we therefore believe that the segmentation protocol can reliably extract even the very small, faint vessels that approach the limit of MR resolution and for which a user can manually define at least one point in MRA slice data.

4. Discussion

The rapid rise of interventional neuroradiology has underlined the need for individualized 3D maps of the intracranial vasculature. As interventionalists may catheterize vessels as small as 1 mm in diameter, clinical utility demands a graph description that includes even the smallest vessels discernible from MRA. This paper describes creation of vascular trees from segmented MRA data.

A number of different groups have described MRA segmentation. Although the described methods may be applicable to other clinical problems, none of the proposed methods appears to provide the kind of detailed vascular trees required by our applications. Most previously described segmentation methods have produced only a few vessels, and graph description has not been attempted

(Alperin et al., 1994; Masutani et al., 1995; Tek and Kimia, 1995; Chapman et al., 1997; Feldmar et al., 1997; Wilson and Noble, 1997). Manual methods are too labor-intensive for routine use. Gerig et al. provide more complete MRA segmentations and also suggest graph description of the vasculature (Gerig et al., 1993; Szekely et al., 1993, 1994; Sato et al., 1998), but the results have not been validated and the number of vessels included in their trees appears sparse. A recent collaborative effort between MIT, Harvard, INRIA, and Cermics extracts vessels from phase contrast MRA using geodesic active contours (Lorigo et al., 1999). The extractions shown illustrate only a few arteries, mix veins and arteries indiscriminately, and contain a great deal of noise. No approach is suggested to segregate veins from arteries or either from noise. Another recent attempt by another excellent group is that of Wiro Niessen and Max Viergever (Frangi et al., 1999; Niessen et al., 1999). This approach requires an initial specification of both ends of a vessel and delineation of vessel walls via a marching-cubes technique. The process requires much more time, user-interaction, and specialized expertise than our own. It also will not work for voxel-width vessels whose walls cannot be defined via marching cubes.

Alternative approaches have also not yet met the needs we intend to fill. Model-based reconstruction of 3D vessels from a small number of fluoroscopic images has been successful for the simple coronary circulation, but is not possible in the highly complex and variable intracerebral circulation in which projection overlap in the absence of a prior 3D model makes curve association difficult. Various methods to associate 2D curves under these conditions have met with only partial success (Rubin and Sayre, 1978; Henri et al., 1990; Delaere et al., 1991; Bahner et al., 1993a,b; Henri and Peters, 1996a,b; Bullitt et al., 1997b,c,d; Smets, 1990). Commercially available programs to reconstruct 3D vessels from a sequence of narrowly apposed DSA views do not provide parent–child information. Stereoscopic DSA imaging is used by a few institutions and works well for some applications (Peters et al., 1996). However, viewing angles are limited to those obtained during angiography, individual subtrees cannot be removed to reduce obscuration, vessels cannot be imaged along with a 3D object such as a tumor, and special glasses make the approach impractical during microscopic surgery. Our approach has none of these limitations.

4.1. Limitations of our approach

We believe that our approach can provide critical information for clinical planning and surgical guidance in ways not heretofore available. Nevertheless, our methodology contains limitations. A potential disadvantage of our segmentation protocol is that it provides geometrical information alone. We have not attempted to employ phase-contrast MRA, as such image acquisition requires a

minimum of four scans, thus making acquisition time too lengthy for routine clinical use and patient motion problematic. Instead, we have applied our methods to time-of-flight MRA, a fast vascular imaging protocol in widespread clinical use, but one that does not provide flow information directly. As a result, we do not know flow direction in the Circle of Willis without incorporating additional DSA information. One automated solution may be to use the width information of segmented data, since arterial flow is normally directed from wider arteries into narrower ones.

A limitation inherent to any method of determining nodal connections is that missing data may produce connectivity errors. There are at least three reasons why our segmentations may contain gaps. First, if the MRA

does not cover a sufficient volume, the tops of vascular curves will be truncated. In such cases the graph description will either fail to include the distal, disconnected curve piece or, worse, will falsely connect the distal curve piece to a neighboring, unrelated vessel (Bullitt et al., 1999a).

A second cause of missing data is that, as previously noted, the resolution of MRA is limited. If two vessels course closely together their profiles may merge on MRA slice data, making it impossible to define two separate vessels. We are currently dealing with this problem by manual editing, which is sometimes required to accurately define the paired anterior cerebral arteries. A long-term, automated solution might involve either a knowledge-based segmentation approach in regions of known difficulty, or the use of higher resolution 3D image data.

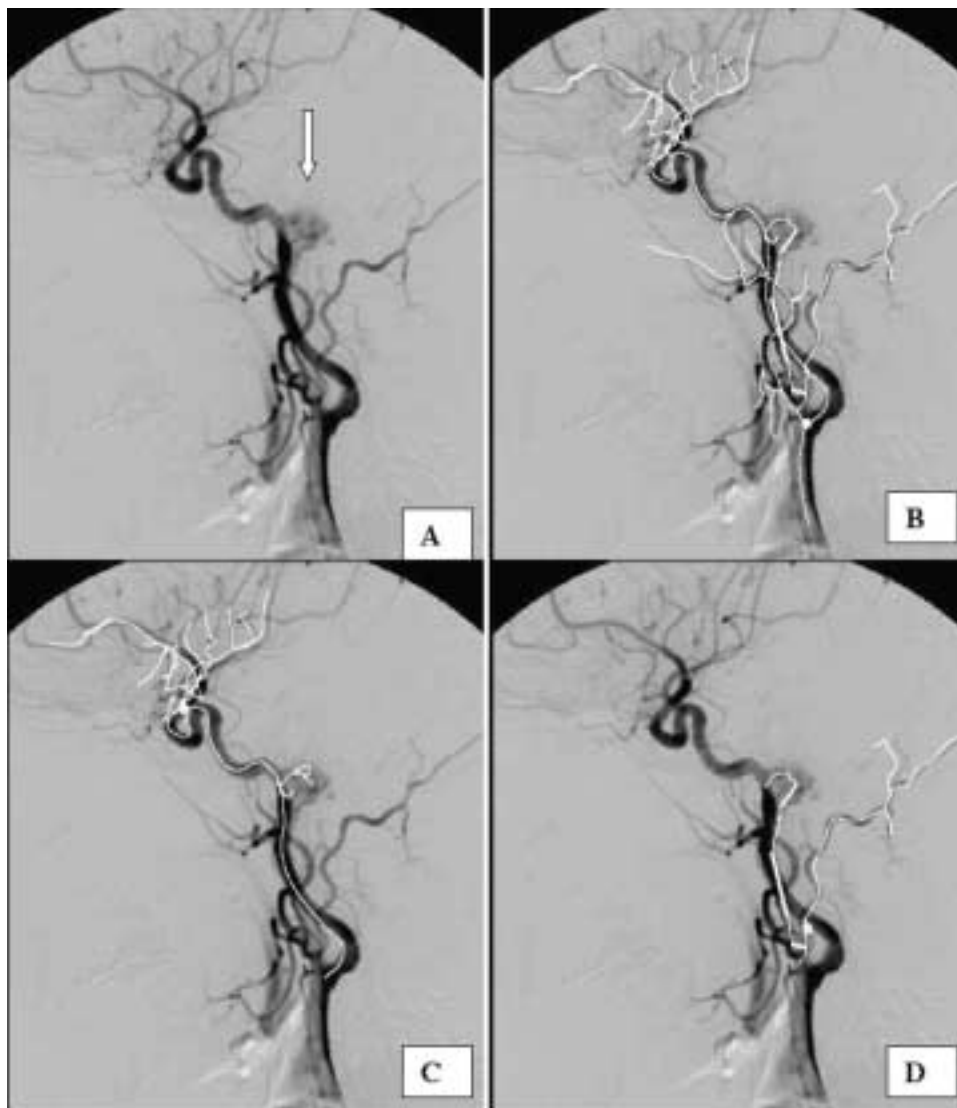


Fig. 5. Vessel trees and endovascular surgery. (A) DSA of a patient with a glomus tumor (arrow). (B) Skeletons of a vessel tree projected on the DSA. This tree contains 32 branches, each derived from a seed point. Note the difficulty of determining parent–child relationships on projection. (C) Subtree arising from the internal carotid. The internal carotid is partially encased by tumor and sends branches to it, but passes through the tumor to supply the intracranial vasculature. (D) Subtree comprised of the ascending pharyngeal and occipital artery and child branches. The ascending pharyngeal artery supplies the tumor directly.

A third potential source of missing data is that MRA may not visualize vessels with slow or turbulent flow. We have not, to date, encountered this problem as our segmentation method appears to recognize tubular objects with great sensitivity.

Another limitation of our approach is that, for some types of surgical planning, the amount of detail provided by MRA may be insufficient. The current study analyzes only the accuracy of nodal connections and the presence of extraneous data. It does not address the issue of missing vessels except as such vessels influence the accuracy of nodal connection. Our methods are applicable to any type of 3D data, however. As the new technique of 3D-DSA improves and/or as MR resolution increases, more 3D vascular detail should become available and our methods will still apply. Images of our methods as applied to high resolution datasets are available elsewhere (Bullitt et al., 2001).

Finally, it would be preferable if our methods were fully automatic. At present, minimal user interaction is required except during the manual definition of seed points. Supplying these seed points may take as little as 1–5 min (to define the Circle of Willis and additional relevant branches in an aneurysm patient), but may also require as much as 90 min (to define a plethoric AVM involving most of the head). However, as stated under Methods, we are actively developing methods of fully automated vessel extraction that we hope should reduce the total user time to a few seconds if no manual editing is planned, and to a few minutes if optional editing is desired.

4.2. Advantages of our approach

Despite these limitations, our approach has several advantages, many of which are inherent to the segmenta-

tion method itself. First, our segmentation method is capable of tracking an individual vessel for long distances. Second, the extractions do not appear to jump from one vessel into another or to extend for noticeable distances into patches of noise. Third, the segmentation of each vessel is close to complete, so that in the majority of cases the parental and child connection points are within a voxel of each other. This feature makes it possible to enforce tight connectivity requirements during tree definition and to deal reasonably effectively with the problem of spurious vessels produced by noise.

Another advantage of our approach is that there are no special image acquisition requirements. The 3D images employed in this report were all routine, time-of-flight MRA sequences previously obtained for clinical use. Our approach is equally applicable to computed tomographic angiography (CTA) and to 3D DSA.

A third advantage of our methodology is that it is not computationally expensive. As a result, no specialized hardware is required. All programs are cross-platform, and run well under Windows on a 433-MHz laptop PC. As a consequence, surgical planning can be done in the office or at home. Figs. 5 and 6 illustrate two different visualizations of a glomus tumor and its associated vasculature. These figures provide examples of the kind of utility we can provide for both surgical planning and endovascular embolization. In Fig. 5, the skeletons of 3D vascular trees are shown superimposed upon a DSA obtained during embolization of a glomus tumor. A 3D graph description of segmented MRA vessels could be used to define endovascular routes through the complex maze of overlapping shadows seen on projection view, without the need for acquiring multiple DSA image sequences. Fig. 6 is an image taken from our surgical planning program, and illustrates the same tumor displayed with vascular subtrees

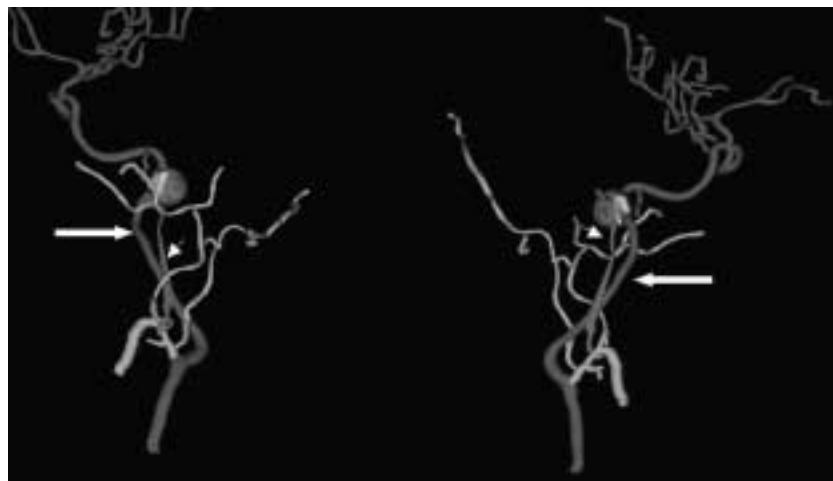


Fig. 6. Vessel trees and surgical planning. The vessels are the same as shown in Fig. 5. The tumor margins have been approximated by defining a sphere in the MR data. A volume rendering of this region is shown in combination with the color-coded vessel tree. Arrows point to the supplying arteries: the large arrow points to the internal carotid and the smaller arrow to the ascending pharyngeal artery. The tumor and vessels can be interactively visualized from any angle of view.

in different colors. The 3D structures can be rotated and viewed from any angle, with color coding used to identify different branches or subtrees seen from various angles of view. Clinicians in our institution are already using this kind of visualization for preoperative planning.

5. Conclusion

This report describes and tests a method of creating intracerebral vascular trees from vessels segmented from MRA. An advantage of the approach is that it can provide 3D graph descriptions that include even the small, faint vessels in MRA data while also excluding noise. No special imaging acquisition requirements are required, and all programs are cross-platform. We evaluate the accuracy of seven vascular trees by asking two radiologists to evaluate each of the 140 parent–child connections comprising these trees. Results indicate the proportion of correct connections to be in the range of 97%.

Acknowledgements

This work was supported, in part, by NIH-NCI R01CA67812, NIH-NCI P01CA47982, a Microsoft Equipment Grant, and an Intel Equipment Award. We have licensed the vascular software to Surgical Navigation Technologies (Medtronic, Minnesota, MN). Siemens will donate hardware/software so that the methods can be applied to 3D DSA.

References

- Alperin, N., Levin, D.N., Pelizzari, C.A., 1994. Retrospective registration of X-ray angiograms with MR images by using vessels as intrinsic landmarks. *J. Magn. Reson. Imaging* 4, 139–144.
- Aylward, S.R., Pizer, S.M., Bullitt, E., Eberly, D., 1996. Intensity ridge and widths for 3D object segmentation and description. In: *IEEE Workshop on Mathematical Methods in Biomedical Image Analysis (WMMBIA)*. IEEE Press, Los Alamitos, CA. pp. 131–138.
- Bahner M., Dick J., Kardatzki B., Ruder H., Matthias S., Steitz A., Bertram C., Hentschel D., Hidebrand T., Hundt E., Kutka R., Stier S., Gerig G., Koller T., Kubler O., Szekely G. Combining two imaging modalities for neuroradiological diagnosis: 3D representation of cerebral blood vessels. In: Pflieger, S., Goncalves, J., Vernon, D. (Eds.), *Data Fusion Applications, Research Reports ESPRIT, Project 5345, DIMUS, Vol. 1*. Springer, Berlin, pp. 1–16
- Bahner, M., Dick, J., Kardatzki, B., Ruder, H., Schmidt, M., Steitz, A., 1993b. Reconstruction of the 3D structure of cerebral vessels from few DSA projections. In: Lemke, H.U., Inamura, K., Jaffe, C.C., Felix, R. (Eds.), *Computer Assisted Radiology: Proceedings of the International Symposium CAR*. Springer, Heidelberg, pp. 553–558.
- Bullitt, E., Liu, A., Aylward, S.R., Soltys, M., Boxwala, A., Rosenman, J., Pizer, S., 1997a. Methods for displaying intracerebral vascular anatomy. *Am. J. Neuroradiol.* 18, 417–420.
- Bullitt, E., Liu, A., Pizer, S.M., 1997b. Three dimensional reconstruction of curves from pairs of projection views in the presence of error. I. Algorithms. *Med. Phys.* 24, 1671–1678.
- Bullitt, E., Liu, A., Pizer, S.M., 1997c. Three dimensional reconstruction of curves from projection views. II. Analysis of error. *Med. Phys.* 24, 1679–1687.
- Bullitt, E., Liu, A., Aylward, S., Pizer, S.M., 1997d. Reconstruction of the intracerebral vasculature from MRA and a pair of projection views. *Lect. Notes Comp. Sci.* 1230, 537–542.
- Bullitt, E., Aylward, S., Liu, A., Stone, J., Mukherji, S., Coffey, C., Gerig, G., Pizer, S.M., 1999a. 3D graph description of the intracerebral vasculature from segmented MRA and tests of accuracy by comparison with X-ray angiograms. *IPMI 99. Lect. Notes Comp. Sci.* 1613, 308–321.
- Bullitt, E., Liu, A., Aylward, S., Coffey, C., Stone, J., Mukherji, S.K., Muller, K., Pizer, S.M., 1999b. Registration of 3D vessels with 2D digital angiograms. Clinical evaluation. *Acad. Radiol.* 6, 539–546.
- Bullitt, E., Aylward, S., Bernard, E., Gerig, G., 2001. Technical report: Computer-assisted visualization of arteriovenous malformations on the home PC. *Neurosurgery* (in press).
- Chapman, B.E., Sanderson, A.R., Goodrich, K.C., Alexander, A.L., Blatter, D.D., Parker, D.L., 1997. Observer performance methodologies for evaluating blood vessel visibility in MR angiograms using accurate geometric registration to high resolution X-ray angiograms. *Magn. Reson. Med.* 37, 519–529.
- Delaere, D., Smets, C., Suetens, P., Marchal, G., Van de Werf, F., 1991. Knowledge-based system for the three-dimensional reconstruction of blood vessels from two angiographic projections. *Med. Biol. Eng. Comput.* 29, NS27–NS36.
- Feldmar, J., Malandain, G., Ayache, N., Fernandez-Vidal, S., Maurincomme, E., Troussel, Y., 1997. Matching 3D MR angiography data and 2D X-ray angiograms. *Lect. Notes Comp. Sci.* 1205, 129–138.
- Frangi, A.F., Niessen, W.J., Hoogeveen, R.M., Walsum, T.V., Viergever, M.A., 1999. Quantification of vessel morphology from 3D MRA. *MICCAI '99. Lect. Notes Comp. Sci.* 1679, 358–367.
- Gerig, G., Koller, T., Szekely, G., Brechbuhler, C., Kubler, O., 1993. Symbolic description of 3-D structures applied to cerebral vessel tree obtained from MR angiography volume data. *IPMI '93. Lect. Notes Comp. Sci.* 687, 94–111.
- Henri, C.J., Collins, D.L., Peters, T.M., 1990. Multimodality image integration for stereotactic surgical planning. *Med. Phys.* 18, 167–177.
- Henri, C.J., Peters, T.M., 1996a. Three dimensional reconstruction of vascular trees. Theory and methodology. *Med. Phys.* 23, 197–204.
- Henri, C.J., Peters, T.M., 1996b. Three dimensional reconstruction of vascular trees: experimental evaluation. *Med. Phys.* 23, 617–627.
- Higashida, R.T., Hieshima, G.B., Halbach, V.V., 1991. Advances in the treatment of complex cerebrovascular disorders by interventional neurovascular techniques. *Circulation* 83 (Suppl. 1), 196–206.
- Kupersmith, M.J., Berenstein, A., 1993. Vascular malformations of the brain. In: Kupersmith, M.J., Berenstein, A. (Eds.), *Neurovascular Neuro-ophthalmology*. Springer, New York, pp. 107–188.
- Kwan, R.K., Evans, A.C., Pike, G.B., 1999. MRI simulation-based evaluation of image-processing and classification methods. *IEEE Trans. Med. Imaging* 18, 1085–1097.
- Liu, A., Bullitt, E., Pizer, S.M., 1998. 3D/2D registration using tubular anatomical structures as a basis. *MICCAI 1998. Lect. Notes Comp. Sci.* 1496, 952–963.
- Lorigo, L.M., Faugeras, O., Grimson, W.E.L., Keriven, R., Kikinis, R., Westin, C.F., 1999. Co-dimension two geodesic active contours for MRA segmentation. *IPMI 99. Lect. Notes Comp. Sci.* 1613, 126–139.
- Masutani, Y., Kurihara, T., Suzuki, M., Dohi, T., 1995. Quantitative vascular shape analysis for 3D MR-angiography using mathematical morphology. In: Ayache, N. (Ed.), *Computer Vision, Virtual Reality and Robotics in Medicine*. Springer, New York, pp. 449–454.
- Morse, B.S., 1994. Computation of Object Cores From Greyscale Images. Department of Computer Science, University of North Carolina, Chapel Hill, NC, Dissertation.
- Niessen, W., van Swindregt, A.M., Elsmann, B., Wink, O., Viergever, M., Mali, W., 1999. Enhanced artery visualization in blood Pool MRA: results in the peripheral vasculature. *IPMI '99. Lect. Notes Comp. Sci.* 1613, 340–345.

- Peters, T., Davey, B., Munger, P., Comeau, R., Evans, A., Olivier, A., 1996. Three-dimensional multimodal image-guidance for neurosurgery. *IEEE Trans. Med. Imaging* 15 (2), 121–128.
- Purdy, P.D., Samson, D., Batjer, H.H., 1990. Preoperative embolization of cerebral arteriovenous malformations with polyvinyl alcohol particles: experience in 51 cases. *Am. J. Neuroradiol.* 11, 501–510.
- Rubin, J.M., Sayre, R.E., 1978. A computer-aided technique for overlaying cerebral angiograms onto computed tomograms. *Invest. Radiol.* 13, 362–367.
- 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, 143–168.
- Smets, C., 1990. A Knowledge-based System For the Automatic Interpretation of Blood Vessels on Angiograms. Ph.D Thesis. Leuven University Press, Belgium.
- Szekely, G., Gerig, G., Koller, T., Brechbuehler, C., Kuebler, O., 1993. Symbolic description of 3-D structures applied to cerebral vessel tree obtained from MR angiography volume data. *CAIP '93. Lect. Notes Comp. Sci.* 719, 687–692.
- Szekely G., Koller T., Kikinis R., Gerig G., 1994. Structural description and combined 3D display for superior analysis of cerebral vascularity from MRA. In: Robb, R. (Ed.), *Proceedings of Visualization in Biomedical Computing (VBC)*. SPIE P, Vol. 2359. pp. 272–281.
- Tek H., Kimia B.B., Volumetric segmentation of medical images by three-dimensional bubbles. In: *Proceedings of the IEEE Workshop on Physics-based Modeling in Computer Vision (PBMCV)*. IEEE Press, Los Alamitos, CA, pp. 9–16.
- Vinuela, F., 1992. Functional evaluation and embolization of intracranial arteriovenous malformations. In: Vinuela, F., Halbach, VV., Dion, J.E. (Eds.), *Interventional Neuroradiology: Endovascular Therapy of the Central Nervous System*. Raven, New York, pp. 77–86.
- Wilson, D., Noble, J.A., 1997. Segmentation of cerebral vessels and aneurysms from MR angiography data. *Lect. Notes Comp. Sci.* 1230, 428–433.