Acquisition, segmentation and tracking of the cerebral vascular tree on 3D magnetic resonance angiography images

Nicolas Flasque*, Michel Desvignes, Jean-Marc Constans, Marinette Revenu

GREYC-ISMRA, 6 Boulevard Marechal Juin, 14050 Caen Cedex, France

Abstract

This paper presents a method for the detection, representation and visualisation of the cerebral vascular tree and its application to magnetic resonance angiography (MRA) images. The detection method is an iterative tracking of the vessel centreline with subvoxel accuracy and precise orientation estimation. This tracking algorithm deals with forks. Centrelines of the vessels are modelled by second-order B-spline. This method is used to obtain a high-level description of the whole vascular network. Applications to real angiographic data are presented. An MRA sequence has been designed, and a global segmentation of the whole vascular tree is realised in three steps. Applications of this work are accurate 3D representation of the vessel centreline and of the vascular tree, and visualisation. The tracking process is also successfully applied to 3D contrast enhanced MR digital subtracted angiography (3D-CE-MRA) of the inferior member vessels. In addition, detection of artery stenosis for routine clinical use is possible due to the high precision of the tracking algorithm. © 2001 Elsevier Science B.V. All rights reserved.

Keywords: Cerebral vascular tree; Magnetic resonance angiography

1. Introduction

Magnetic resonance angiography (MRA) offers an interesting alternative to computerised tomography angiography (CTA) or other 2D angiograms for clinical use. MRA is a non-invasive technique (as opposed to X-ray angiography) and provides high-quality 3D images. The wide variety of acquisition sequences and techniques [2D versus 3D, time of flight (TOF) versus phase contrast (PC)] allows its use for many clinical examinations. The lack of nocivity and the speed of clinical acquisition make MRA a useful and attractive tool for diagnosis and surgical planning of arterio-venous malformations (AVM), aneurysms and stenosis. AVM are congenital disorders in which arteries and veins have direct and often complex connections, depriving the concerned area of a normal blood supply. Aneurysms are pockets of blood, sometimes very large,

caused by weakness in the arterial wall. The rupture of an aneurysm can cause lethal haemorrhaging and severe neurological deficiencies. Stenoses are occlusions of vessels (from partial to even total obstruction as defined by the NASCET criterion) with significant risk for the vascularisation of the brain. In these cases, real 3D manipulation and visualisation is helpful for cautiously examining the pathological area for diagnosis and for surgical planning. MRA has become a robust, available and cost-effective tool in clinical routine.

Although numerous 3D visualisation tools exist, handling and visualisation of thin 3D structures such as peripheral vessels is not easy. The most common tool is the maximum intensity projection (MIP), which is a very simple ray tracer. It remains a 2D projection of a 3D object along a single direction, and presents many drawbacks such as visual occultation of vessels and artificial crossings that can only be detected by matching on several other projections. Furthermore, low-intensity structures are always masked by high-intensity ones. Due to MIP image formation, the size of objects is not correctly estimated.

Accurate visualisation by surface or volume rendering techniques requires a segmentation step followed by a rendering of blood vessels.

Obtaining quantitative measurements such as the diameter or cross-sectional area of vessels which is necessary to reveal pathologies is a very difficult task in 3D (Wilson and Noble, 1997). The diameter is estimated visually, and is largely underestimated in stenotic areas as the flow that generates the MRA signal becomes irregular (Hoogeveen et al., 1998). A correct estimation of diameter and cross-sectional area also requires that the direction of processing lies in an interval of only a few degrees away from the real axis of the blood vessel. Reliable visualisation and quantification requires a processing of the 3D images to properly extract vessel parameters.

In this paper, we will present a general method for the accurate detection of 3D tubular vessels on 3D images. We will first describe related works. Next, we will present the detection and tracking process. Results for synthetic images are examined. This method has been applied to the extraction of blood vessels on 3D MRA images. The overall application will be presented, including MRA image sequence design, segmentation and tracking results. Synthetic images, real healthy and pathological examples are examined. This application is a first step towards the measurement of parameters such as the diameter and cross-sectional area of a vessel.

2. Related works

Several works address the problem of segmentation and tracking of vessels on 2D or 3D images to visualise vessels or to extract parameters. Most of them share a common approach: in the first stage, the image is filtered to reduce the noise and to highlight the structures of interest. The detection of small structures is thus more robust (Klose et al., 1995; Du and Parker, 1998). Linear and median filters cannot be used as the finer structures to be detected reach the resolution of the image. Diffusion filtering (Gerig et al., 1992; Wilson and Noble, 1999) reduces noise and preserves geometry, topology and localisation of vessels. The diffusion process can be guided with a priori knowledge of the relevant image features (Sato et al., 1998).

A second common characteristic is the description of structures by the centrelines. Junctions and forks are easy to model, and geometrically complex structures are handled by a graph or tree representation with centrelines as nodes. To extract the centreline, tracking algorithms are based upon one of the following strategies.

- Segmentation/detection of edges and extraction of the centreline without a priori knowledge. This method can be used to extract curvilinear structures in many application fields such as medical imaging (vessels, bronchi) or map analysis (roads) (Kutka and Stier, 1996). Skeletonisation of structures is often used to detect centrelines of structures. The accuracy for the

localisation is limited to the resolution of the image, since this processing takes place in the discrete 3D space of voxels.

- Modelling and matching. The model is in this case largely constrained by the application field (Wilson and Noble, 1999). A priori knowledge has to be brought by specialists.
- Extraction of the centreline by snakes (Cohen and Kimmel, 1996). This method requires user interaction to define the start and the end of a single line. Junctions are not modelled.
- Multi-scale analysis (Krissian et al., 2000). Tubular structures are identified by the response of voxels to second derivatives of a set of Gaussian kernels with which the image is convolved.

Some of these methods are purely dedicated to 2D images and come from 2D image processing of classical angiography. Sun et al. (1995) used continuity constraints of vessel centrelines with orientation and density information to detect the centreline. Directional low-pass filter responses are compared to estimate the borders of a vessel and its diameter with subpixel accuracy. Figueiredo and Ao (1995) used morphological operators and dynamic programming to track the centre and the edges of a vessel. In this work, the morphological approach is preferred to smoothing, as the latter process is known to modify the edge location. Directional filters and Gaussian modellisation are used, but the accuracy of the orientation estimation is directly related to the number of filters with which the image is convolved (Klose et al., 1995). A priori knowledge and local constraints enhance the detection and the localisation (Sonka et al., 1993).

Summers et al. (1997) proposed an octree decomposition of a velocity field image in order to find an optimal tessellation. Each block of the octree contains at most one feature, defined by a real grey-level function and local orientation. This work is strictly dedicated to phase contrast MRA, containing information about velocity of the flowing blood.

The main drawbacks of these techniques are the difficulty of handling the artificial junctions that occur in 2D projections.

To take these junctions into account, Chen et al. (1998) used directional Gabor filters. Another interesting approach was proposed by Coatrieux et al. (1995), where the tracking of the edges is processed with the uniform cost algorithm A*. This finds an optimal border path in a graph representation of the image, by searching the path with the minimum cost. The cost is defined in order to match the border of the vessels. The same idea is exploited by Lecornu et al. (1994) to simultaneously detect the two borders of a vessel. Overlapping structures are correctly detected.

Most of these works are not easily transposable to 3D images. Centrelines are always curvilinear, but edges become 3D surfaces, which are difficult to follow and require more sophisticated modelling.

Reuzé et al. (1993) and Hernandez-Hoyos et al. (1999) detected vessels by cylinder matching. The method is based upon minimisation of inertia moments of a cylinder and a priori knowledge of the intensity profiles in and at the edge of a vessel. This work does not take into account the junctions. Calculation of moments leads to strong hypotheses on the local configuration of the image and on the vessel geometry: there must be only one vessel in the processed area, and the cross-section of the vessel must be circular. While the latter hypothesis can be accepted in non-pathological situations, the shape of a blood vessel can be greatly modified near stenoses and aneurysms (Juhan et al., 1997). Verdonck et al. (1995) used generalised cylinders. The cross-section of a vessel is approximated by a polygon. Continuity and orientation between consecutive slices are used to calculate a locally optimal shape for the polygon with good accuracy, without quantitative measurement. Cohen and Kimmel (1996) proposed a multi-scale space approach and a Gaussian modellization of the intensity profile inside the vessels. The convolution of the image with a series of Gaussian kernels with different standard deviations gives a scale-set of images where centrelines are detected as local extrema of a response function designed for a set of models (Sato et al., 1998; Krissian et al., 1998; Lorenz et al., 1997; Koller et al., 1995; Poli and Valli, 1998). The localisation of vessels is quite accurate and junctions can be detected. Wilson and Noble (1997) estimated the centreline after segmenting the image by skeletonisation. These works use convolutions with several Gaussian kernels and are thus time consuming; furthermore, the accuracy of the centreline is limited to one voxel. Frangi et al. (1999) used a model of the centreline coupled with the wall surface of linear segments of vessels. This model provides accurate detection and quantification of stenoses but junctions are not modelled.

Many works give accurate results in 2D, but are not suited for 3D processing. In 3D approaches, accuracy of the detection and/or handling of junctions is not completely achieved. We chose to emphasise the following features:

- minimum user interaction in the segmentation and tracking process;
- the junctions and branches should be efficiently handled;
- accuracy of centreline detection has to be less than 1 voxel for the location and 5° for the cross-section orientation in order to provide reliable quantitative measurements;
- processing should not be time-consuming in order to be usable in a routine clinical environment.

3. Vessel tracking

The goal of the tracking process is to extract points of the centreline of the vessel on a 3D angiographic image. Tracking is an iterative search with subvoxel accuracy. The

two main features of the process are the subvoxel accuracy and the detection of junctions.

The initial 3D angiographic image is filtered to suppress noise. The tracking process can be done on grey-level images, but the junction handling needs a binary one for connectivity purposes.

The result of the tracking process is a set of centreline points which allows a continuous modelling by B-spline and precise measurements.

The global methodology of the tracking is an iterative process: given an initial point P_0 and an initial direction D_0 , an ordered set of points $\{P_i\}$ is computed step by step in a local area around the previous point P_{i-1} .

3.1. Local search area

At a step or iteration i of the tracking process, a local search area is built from the point P_i and the direction D_i . This area is a parallelepiped (or box) of variable dimensions L_i , l_i (plane orthogonal to D_i) and H_i (parallel to D_i). The next centreline point P_{i+1} and the next orientation D_{i+1} are estimated in this box (Fig. 1). We call face F the face of the box including the point P_i , orthogonal to the direction D_i . The grey levels in this box are over-sampled from the initial image. In our experiments, we use trilinear interpolation. The sampling within the local search area has the same resolution as the image.

3.2. Detection of point P_{i+1}

Let P_i be the current point, D_i the current tracking direction and $D_{\perp 1}$, $D_{\perp 2}$ two directions so that $D_i \perp D_{\perp 1} \perp D_{\perp 2} \perp D_i$. D_i , $D_{\perp 1}$ and $D_{\perp 2}$ are unitary vectors.

Let L_i , l_i and H_i be the dimensions of the box in integer number of voxel.

Let $\alpha \in [-l_i; l_i]$, $\beta \in [-L_i; L_i]$, $\gamma \in [0; H_i]$.

Let $X_{\alpha, \beta, \gamma} = P_i + \alpha \cdot D_{\perp 1} + \beta \cdot D_{\perp 2} + \gamma \cdot D_i$, and $A_{\alpha, \beta, \gamma}$ a weight associated with $X_{\alpha, \beta, \gamma}$.

$$V = \{X_{\alpha, \beta, \gamma} / A_{\alpha, \beta, \gamma} > 0\},$$

Fig. 1. Local search area.

$$B_i = \frac{\sum_{X_{\alpha,\beta,\gamma} \in V} A_{\alpha,\beta,\gamma} X_{\alpha,\beta,\gamma}}{\sum_{X_{\alpha,\beta,\gamma} \in V} A_{\alpha,\beta,\gamma}},$$

$$P_{i+1} = \frac{(k_p P_i) + B_i}{k_p + 1},$$

$$\overrightarrow{D_{i+1}} = \frac{\overrightarrow{D_i}}{\|\overrightarrow{D_i}\|} + k_d \left(\frac{\overrightarrow{B_i - P_i}}{\|B_i - P_i\|} \right).$$

Parameters k_p and k_d have a fixed value corresponding to the maximal curvature of vessels that can be tracked by this technique. For a given set $\{k_p, k_d\}$, the maximum curvature that can be estimated by the tracking algorithm is limited by the dimensions of the parallelepiped which are computed at each step of the tracking. $k_p = 1/2$ and $k_d = 30$ are used in the following examples.

B_i is the weighted centre of mass of the pertinent voxels in the local search area ($X_{\alpha,\beta,\gamma} \in V$). A priori knowledge on the image can be exploited through the set $\{A_{\alpha,\beta,\gamma}\}$ of weightings. $A_{\alpha,\beta,\gamma}$ represents the weight at the point $X_{\alpha,\beta,\gamma}$, depending on the imaging modality. If the shape of the structure is more informative than its intensity distribution, like in 3D-CE-MRA or in DSA, binary weights $A_{\alpha,\beta,\gamma} = 1 \Leftrightarrow \text{Im}(\alpha, \beta, \gamma) > 0$ are used. If the vessel presents a Gaussian-like profile, an image weighting $A_{\alpha,\beta,\gamma} = \text{Im}(\alpha, \beta, \gamma)$ takes this information into account. If the vessel profile is known to show a narrow grey-level maximum at its centre, an exponential weighting $A_{\alpha,\beta,\gamma} = e^{\text{Im}(\alpha, \beta, \gamma)}$ still allows us to use the image features.

3.3. Detection of junctions

The detection of junctions is realised by the analysis of connected components in the volume and on the surface of the parallelepiped obtained by clustering voxels in or on the box. For a single vessel, the incoming flow through face F must exit from the parallelepiped by only one connected component.

Let NV be the number of volumic connected components in the parallelepiped P and NS the number of surfacic connected components on its surface. NV is the number of vessels inside P and NS is the number of ingoing/outgoing vessels. A single vessel gives values of 1 for NV and 2 for NS . A vessel with a junction will give values of 1 for NV and 3 for NS . For other values, each surfacic component is also labelled with the label of the volumic component it is attached to (Fig. 2). NS can be restricted to 2 or 3 for each vessel and junctions are detected with no ambiguity. Voxels that do not belong to the main vessel introduce a bias in the computation of B_i , P_{i+1} , D_{i+1} and are removed from the box.

When a junction is detected, the gravity centre of each outgoing connected component gives a new starting point for a tracking step. The initial orientation is given by the

Fig. 2. Possible local configuration in the parallelepiped (from left to right and top to bottom): (a) single vessel, $NS = 2$, $NV = 1$; (b) two vessels, $NS = 4$, $NV = 2$; (c) simple fork, $NS = 3$, $NV = 1$; (d) simple fork and a vessel, $NS = 5$, $NV = 2$.

difference between this new initial point and the last B_i point.

3.4. Computation of adaptive dimensions

The dimensions L_i , l_i and H_i are calculated at each step of the tracking algorithms to suit the local data with maximal accuracy and speed.

We must ensure that only one vessel is tracked at a time by the parallelepiped: dimensions L_i and l_i are computed at each step so that only one connected component is visible through the 'lower' face F . Non-circular sections for vessels imply that L_i and l_i are computed separately. Dimensions are first reduced, possibly several times, if all the border voxels of the face F have a value of 0. They are then augmented once if any border pixel of the face F belongs to the main connected component in this face.

The change of direction between two steps i and $i + 1$ is evaluated by the following value:

$$\Delta \tilde{\epsilon} = \arccos \left(\frac{\langle \overrightarrow{D_i} | \overrightarrow{(P_{i+1} - B_{i+1})} \rangle}{\|\overrightarrow{D_i}\| \cdot \|\overrightarrow{(P_{i+1} - B_{i+1})}\|} \right).$$

Large values of $\Delta \tilde{\epsilon}$ indicate a significant change of direction between steps i and $i + 1$. This situation effectively occurs in carotid artery and venous sinus tracking. In this case, the direction change between D_i and D_{i+1} must be suited to those abrupt changes and the position update must be as small as possible in order to stay in the correct area. A large number of points must be computed in sinuous areas for good accuracy.

H_i is the horizon of the parallelepiped. Straight sections do not require as many sampling points as sinuous ones. H_i evolves inversely proportional to the local curvature,

estimated by the direction change $\Delta \Xi$, with a fixed maximum value k_h for straight sections,

$$H_i = k_h e^{-\Delta \Xi}.$$

The accuracy of the point location depends on k_p and H_i : if r is the current curvature of the vessel (Fig. 3), the expected position error is

$$e_p = \frac{1}{r} \left(1 - \sqrt{1 - \frac{r^2 H_i^2}{4k_p^2}} \right) \approx \frac{H_i^2 r^3}{2k_p^2}.$$

This dynamic handling of the dimensions ensures that the tracking process remains as close as possible to the vessel. It also overcomes small disconnections of vessels that occur in the segmentation step. Typical values of the dimensions are 5 ± 2 for L_i and l_i , and 3 ± 1 for H_i .

3.5. B-spline modelling for the centreline

The list $\{P_i\}$ obtained by the tracking algorithm is interpolated by a second-order B-spline curve. We choose this modelling for three reasons:

- we have an analytic representation of the centreline which allows resampling of the cross-section along the whole course of the vessel, even at points that were not directly computed by the tracking algorithm;
- those curves minimise their torsion, just as any physical system submitted to a set of constraints;
- the underlying model for vessel description must deal with pathological cases, namely stenosis.

We use two sets of descriptors for the vessel: a set $\{C_i, i \in [1, \dots, n]\}$ of centre points and a set $\{S_i, i \in [1, \dots, n]\}$ of data on cross-sections. A section S_k corresponds to the point C_k .

The points C_i are computed along the second-order B-spline interpolating the set $\{P_i\}$ of detected points. This

Fig. 3. Position error, e_p , estimation for arbitrary values of r , H_i and k_p .

analytic representation allows resampling and computation of a continuous first derivative Δ . Samples of Δ at the C_i points (Δ_i) give the tangent direction to the B-spline curve. The S_i cross-sections are built by interpolation in the plane orthogonal to Δ_i and are designed for quantification and virtual endoscopy. C_i and Δ_i are 3D vectorial values.

3.6. Evaluation of the accuracy

Several different synthetic image sets were used to test the various aspects of the tracking algorithm (Fig. 4): straight tubes, 3D-Lissajoux shapes, and Y-fork shapes with the angle varying from 20 to 160° in steps of 20°. All images were generated using a second-order B-spline curve. The intensity of the cross-section has a Gaussian profile. The intensity of each voxel of the image is computed using the distance of the nearest point of the B-spline curve. For each shape, four diameters are generated: 1, 2, 4 and 6 voxels, with a resolution of 1 mm. These synthetic images are used directly as binary masks for the tracking algorithm. The weights $A_{\alpha,\beta,\gamma} = \text{Im}(\alpha,\beta,\gamma)$ are used.

The following parameters are measured between the synthetic object real centreline $\{L_i, i \in [1, \dots, p]\}$ and the detected centreline $\{C_i, i \in [1, \dots, p]\}$: the difference ΔP in position between the nearest points and the difference $\Delta \Theta$ in the orthogonal section orientation between those points. The detected centreline is the bicubic B-spline curve. The sets $\{A_i, i \in [1, \dots, p]\}$ and $\{\Delta_i, i \in [1, \dots, n]\}$ give the derivative for each L_i or C_i point.

The distance between the centrelines and the orientation difference is measured between $\{(C_i, \Delta_i), i \in [1, \dots, n]\}$ and $\{(L_j, A_j), j \in [1, \dots, p]\}$.

$$\Delta P_i = \min_{j \in [1, \dots, p]} (\|C_i L_j\|) = C_i L_{jmi},$$

$$\Delta \Theta_i = \arccos \left(\frac{\langle \Delta_i | A_{jmi} \rangle}{\|\Delta_i\| \cdot \|A_{jmi}\|} \right).$$

3.6.1. Tracking a single structure

Table 1 presents the results of the tracking process on the following images: straight tubes in the X, Y and Z direction and three 3D Lissajoux shapes. For each shape, four diameters are tested: 1, 2, 4 and 6 isotropic voxels of 1 mm³. On each image, the minimum, maximum and

Fig. 4. (from left to right) 3D rendering of a Lissajoux shape, a 20° Y fork shape and a 160° Y fork shape. Diameter is 4 voxels.

Table 1

Global results for the tracking of straight tubes and Lissajoux shapes

	ΔP (voxel)		$\Delta \Theta$ (deg)	
	Avg	S.D.	Avg	S.D.
Min	0.12	0.01	0.01	0
Max	1.49	0.12	26.32	6.58
Avg	0.68	0	1.31	0.10

average value of ΔP and $\Delta \Theta$ are measured. The average and standard deviation of these values on the 24 images are listed in Table 1.

We estimate the expected error on the cross-sectional area as $err = 1 - \cos^2(\Delta \Theta)$. err is always inferior to 1%. Maximum $\Delta \Theta$ values are reached during the first iterations of the tracking process. $\Delta \Theta$ values oscillate around the calculated average value after a few iterations ($i > 10$).

To evaluate the influence of the curvature on the method, a 3D helix with increasing curvature from 4 to 100 mm and 1 mm diameter was processed with a constant mean error $\Delta P = 0.4$ mm and $\Delta \Theta = 1.45^\circ$ (Fig. 5).

3.6.2. Robustness to user interaction

Initial values D_0 and P_0 are interactively provided by the user. The assumptions made on the choice of P_0 and D_0 are that P_0 is inside the vessel, P_0 is not located on a junction, and that the angle between the true orientation of the vessel and D_0 is less than 45° . The limits of the tracking accuracy with respect to the interactive initial values are evaluated for three specific values of P_0 and D_0 on a vessel of diameter r (Fig. 6).

The tracking fails for orientation shifts over 85° .

3.6.3. Robustness to noise

The vessels in real angiographic images do not have an ideal grey-level profile. Images are noisy and vessels sometimes present discontinuities or a varying intensity, especially in MRA. The robustness of the tracking method is evaluated relative to both phenomena.

Fig. 5. Tracking of synthetic structure with increasing curvature.

Initial position outside the structure with good orientation : $\langle D_0$

$$D_{\text{real}} \geq 0^\circ ; d(P_0 ; P_{\text{real}}) = 2r$$

	ΔP (voxel)	$\Delta \Theta$ (°)
Min	0.11	0.01
Max	0.98	44.97
Avg	0.68	1.56

Initial position in the structure - : $\langle D_0$

$$D_{\text{real}} \geq 45^\circ ; d(P_0 ; P_{\text{real}}) = r$$

	ΔP (voxel)	$\Delta \Theta$ (°)
Min	0.12	0.01
Max	2.53	73.47
Avg	0.74	2.20

Initial position in the structure - : $\langle D_0$ $D_{\text{real}} \geq 60^\circ ; d(P_0 ; P_{\text{real}}) = r$

Fig. 6. ΔP and $\Delta \Theta$ values for different initial settings of P_0 and D_0 .

The synthetic image of a straight structure is modified as follows (Table 2, Fig. 7, from left to right): two consecutive slices are removed from the image perpendicularly to the structure, the intensity is progressively reduced to 50% of its original value, and another slice is removed from the image.

An additive uniform noise is added to the image. The noise level is defined by its amplitude relative to the maximum intensity in the image.

The position of the detected centreline method is very robust to noise and the orientation is really robust below a noise amplitude of 10% (Table 3). The tracking method fails when the noise value reaches the intensity of the

Table 2

Values of ΔP and $\Delta \Theta$ for the tracking of a discontinuous tubular structure

	ΔP (voxel)	$\Delta \Theta$ (deg)
Min	0.66	0.01
Max	1.44	17.55
Avg	0.74	2.01

Fig. 7. Tracking of a discontinuous structure.

Table 3

Average values of ΔP and $\Delta \theta$ at different levels of noise

Noise amp	Avg ΔP (voxel)	Avg $\Delta \theta$ (deg)
5%	0.68	1.49
10%	0.69	1.46
20%	0.73	5.02

Fig. 8. Curvature and intensity discontinuity.

structure to be detected. Voxels belonging to the background become connected to the structure.

We also tested the handling of intensity discontinuity and noise in the case of curved vessels (Fig. 8). Several images were generated with different discontinuity gap length and curvature values. This process fails when the gap length reaches the dimension of H_i .

3.6.4. Evaluation of the junction detection algorithm

On synthetic images, the junctions from 20° to 160° are always detected by our algorithm. Failure occurs for angles of 10° or less and 170° or more. Another synthetic image with three different junctions is used to demonstrate that the algorithm is able to track complex networks (Fig. 9).

Fig. 9. Tracking of a complete tree structure from P_0 and D_0 .

4. Acquisition and segmentation of the vascular tree

The tracking method is designed to process images composed of unstructured sets of connected voxels. We apply this method to 3D MRA images. A segmentation of those images is thus necessary to generate the sets of connected voxels. The purpose of this segmentation is not to extract only the cerebral vasculature, but rather to ensure that the cerebral vasculature is included in the segmented sets.

4.1. Image acquisition

MRA images were acquired on a General Electric Medical System scanner with a 1.5 Tesla magnetic field. In all cases, a conventional head coil was used. Whole head images of healthy volunteers were acquired. Pathological images address carotid artery stenosis near the main carotid artery separation and are located in the area of the neck and lower head.

Recovering the cerebral vascular tree with accuracy requires cautious design for an imaging sequence (Hoogeveen et al., 1998). Clinical use of MRA often assumes a previous examination with CTA or X-ray angiography to localise the area to be imaged. MR scan time is then optimised to a few minutes for an average resolution of $4 \text{ mm} \times 1 \text{ mm} \times 1 \text{ mm}$.

We want, on the contrary, a high-resolution image of the whole vascular tree. The presence of neighbouring anatomic structures is required for multi-modality matching. PC sequences are very efficient for imaging arteries and veins, even thinner ones, but the subtraction process that is the base of this acquisition technique suppresses anatomical information. Furthermore, two acquisitions are needed to obtain an image and the signal-to-noise ratio is not very good.

With TOF sequences, particular care has to be taken when choosing the various parameters. Only one acquisition is needed. In 3D mode, only thin volumes can be correctly acquired for our purpose. TOF sequences are also used in clinical routine to acquire images for diagnosis of arterial stenosis with good confidence (Wentz et al., 1994).

We use a 3D-TOF SPGR (gradient-echo) imaging sequence in order to obtain a high signal-to-noise ratio. Furthermore, TOF sequences do not entirely suppress the neighbouring anatomical structures. Multi-modality matching can be useful for further registrations with anatomical (sulci location, gray/white matter segmentation) and/or functional maps (PET/functional MRI). The sequence was designed with MOTSA (multislab overlay thin slice acquisition) options to maximise the TOF effect along the whole image. Each slab is 16 mm thick and is acquired in 16 slices. Thicker slabs lead to unrecoverable signal loss (Hoogeveen et al., 1998). The signal is maximal for vessels orthogonal to the slab and for a flow speed slightly

inferior to TE across a single slice. Faster flow generates a signal void (bolus effect).

Healthy volunteer image dimensions are $256 \times 256 \times 152$ voxels in size, with each voxel having an isotropic size of $1 \text{ mm} \times 1 \text{ mm} \times 1 \text{ mm}$. The scan time is 23 min. An axial (from bottom to top) acquisition direction was chosen because the vessels of greater interest for our work, the carotid arteries and the vertebro-basilar artery, are predominantly orthogonal to the axial plane in their main course. The high speed of the flow in those vessels guarantees a sufficient amount of signal in the other portions.

The segmentation of the images is realised in three distinct steps:

- correction of the intensity artefacts generated by the MOTSA acquisition;
- diffusion filtering to highlight the vessels while preserving their geometry, topology and location;
- region growing to isolate vessel-like structures.

4.2. Intensity correction

The blood signal decreases inside a single slab due to the partial saturation effect and the sequence generates intensity artefacts in areas where the slabs overlay. An intensity correction is needed before the filtering step.

This method is based on histogram homothesis after histogram modelling. The background is removed by thresholding. The value of the threshold is computed as the first local minimum on the slice histogram. It is actually a Gaussian modelling of a blank noise with a low standard deviation and zero average. The cumulative histogram of the remaining voxels is built, and two intensity values I_1 and I_2 are extracted from the histogram. I_1 (I_2) is the intensity reached at x_1 % (x_2 %) of the total population (Fig. 10). An extreme value of x_2 (>90) implies values of I_2 given by marginal populations of voxels (essentially noise); close values of x_1 and x_2 prioritise a particular

Fig. 10. Determination of I_1 and I_2 on the cumulative histogram from x_1 and x_2 .

grey-level distribution in the slice certainly associated with a given structure which may not be present in all the volume.

The best results were obtained empirically for $15 \leq x_1 \leq 20$ and $80 \leq x_2 \leq 85$. This selection is processed on all the image slices. A reference slice is then chosen, giving two reference intensities I_{R1} and I_{R2} . The reference slice is chosen according to two criteria:

- the slice does not belong to an overlay area;
- it must contain as many different structures (i.e. grey-level distributions) as possible.

Each of the other slices is recomputed by linear interpolation so that all voxels with I_1 intensity reach I_{R1} intensity and all voxels with I_2 intensity reach I_{R2} intensity. All slices of the volume have the same dynamic after this processing.

4.3. Diffusion filtering

The filtering step uses a priori knowledge about the intensity features of the vessels. This information is directly available from the parameters of the imaging sequence.

Blood vessels appear as homogeneous regions of high intensity and produce high contrast at their border. To enhance these regions, an adaptive diffusion filtering is used. Diffusion filtering is an adaptive and iterative process based upon the heat diffusion equation (Orkisz et al., 1996; Gerig et al., 1993). For each voxel of intensity u , a quantity δu is computed at each step, depending on the local gradient value ∇u and a function g . The intensity u of vessels is greater than for the other structures. We exploit this information by introducing a term depending on the intensity in the function g , instead of using only the gradient estimation ∇u . The voxel intensity value becomes $u + \delta u$, and δu may be negative,

$$\frac{\partial u}{\partial t} = \text{div}(g(u, |\nabla u|) \nabla u).$$

Different functions g exist (Perona and Malik, 1990), and we choose the second Perona–Malik form,

$$g(u, |\nabla u|) = \frac{1}{1 + (u|\nabla u|/k)^2}$$

This function is numerically stable and easy to compute. The strength of filtering depends on gradient and intensity, and we are thus able to filter areas with high gradient values if the intensity is low.

High intensity and high contrast areas are highlighted, while other structures are strongly filtered. The border of the vessels is well located, which is important for further measurements. This process is computationally expensive and converges slowly. The criterion used to end the diffusion filtering process is reported in the application section.

4.4. Segmentation of the connected sets of voxels

The vessels are curvilinear and connected structures. We use voxel clustering to build regions of interest. Voxels are merged into an existing region if they are likely to belong to a vessel, according to two criteria: the intensity of the voxel and its neighbourhood configuration.

The voxel clustering stage is initiated by the detection of a seed voxel $S_{x,y,z} / \text{Im}(S_{x,y,z}) > I_3$, which is a fixed threshold. S defines the first voxel of the region R_1 . The region R_1 grows according to the criterion: $V_{x',y',z'} \in R_1 \Leftrightarrow \text{Im}(V_{x',y',z'}) > I_3$ and one of the 26 neighbours of $V_{x',y',z'} \in R_1$.

When no new voxel can be added to R_1 , another seed voxel is selected in the image to build a second region R_2 , and this process continues until there are no seed voxels left.

Candidate voxels are then chosen among the voxels that have not already been selected, depending on the number of neighbours belonging to the same region.

Let $N_i(V_{x,y,z})$ be the number of 26 neighbours of a voxel $V_{x,y,z}$ belonging to the region R_i . We define

$$N = \max_i(N_i), \quad j/N_j = N,$$

$V_{x,y,z}$ is a candidate to belong to $R_j \Leftrightarrow N > 13$.

Among these candidates, only those with intensity greater than a threshold I_4 are kept. Regions including less than 25 voxels are discarded. Both I_3 and I_4 threshold values are determined on the filtered image histogram in the following way: a least-squares fit of a Gaussian function $G(x_i, \sigma_i)$ is performed on the whole image histogram which exhibits a strong peak corresponding to the anatomical structures. We segment this peak using the following values: $I_3 = x_i + 5\sigma_i$ and $I_4 = x_i + 3\sigma_i$.

We now have a binary mask constituted of unstructured sets of voxels. The core structure of the vascular network has to be rebuilt from these data by a centreline tracking algorithm able to handle junctions.

5. Results and discussion

5.1. Segmentation

5.1.1. Intensity correction

Results of one correction scheme are presented in Table 3. The homogeneity of the signal has been restored, and clearer strips have disappeared from sagittal slices.

Fig. 11 shows that the average intensity value for the corrected image is globally inferior to the average value for the raw image. This effect is due to background removal. The background grey level is set to zero in corrected images.

Intensity gaps are computed on the average intensity of

Fig. 11. Average intensity value on raw and corrected images.

a given slice. Before correction, gaps reach values of 3, and the maximum gap is 1 after correction.

The correction process takes 2 min for a $256 \times 256 \times 152$ voxel image. All computational times are given for an Ultra2 Sun computer.

5.1.2. Diffusion filtering

The diffusion filtering process mainly enhances peculiar areas in the intensity-corrected image. The enhanced areas are connected to vessel-like structures (Fig. 12). The criterion used to end this process depends on the global intensity change in the image. When this intensity change reaches 10% of its value at the first iteration, the diffusion process ends.

These areas are located near high intensity and high contrast structures. At each step of the process, the border voxels of the structures located along the axis of the area are enhanced, thereby allowing restoration of missing vessel parts.

The diffusion process takes 5 min for a $256 \times 256 \times 152$ voxels image.

5.1.3. Voxel clustering

The resulting image is a grey-level mask of the cerebral vascular network. The grey-level values can be used in the tracking process by selecting an appropriate weighting set $A_{\alpha,\beta,\gamma}$. For connectivity purposes, this image is also considered as binary: a voxel at location $X_{\alpha,\beta,\gamma}$ is part of a

Fig. 12. Non-filtered image (opaque) versus filtered image (transparent) 3D isosurfaces of the cerebral middle artery at intensity $I = 90$.

Fig. 13. Left and right carotid artery and vertebral artery tracking on real MRA images.

structure $\Leftrightarrow A_{\alpha,\beta,\gamma} > 0$. All results in the next section are displayed relative to 3D isosurfaces of this binary image.

The clustering process takes less than 1 min for a $256 \times 256 \times 152$ voxels image.

5.2. Detection of the cerebral vascular network

5.2.1. Healthy volunteer images

No reference centreline is available to evaluate the quality of our tracking algorithm. Results were examined visually. Detected centrelines seem well located for both carotid arteries, the vertebral artery, anterior cerebral arteries, and middle cerebral arteries. The junction between the vertebral artery and both posterior cerebral arteries is detected (Fig. 13). The vessel diameter ranges from 3 to 6 voxels on these images.

5.2.2. Pathologic images

As for healthy volunteer images, no reference centreline is available. Fig. 14 shows the centreline tracking on the

primitive carotid artery, before it splits into the internal and external carotid arteries. Many stenoses are observed near this junction. The tracking process does not stop at the stenosis location, and the detected centreline follows the axis of the vessel.

5.2.3. 3D contrast-enhanced MRA image

Blood vessels are well depicted on 3D-CE-MRA images. The connected sets of voxels are obtained by thresholding the original image which is highly contrasted. We successfully applied the tracking method to the arteries of the inferior members and detected the junction (Fig. 15).

Fig. 15. Tracking on a 3D-CE-MRA image of the inferior members.

Fig. 14. Centreline tracking on a stenotic left carotid artery. The interruption is an artefact of the visualisation algorithm. Left column: standard MIP views of the original dataset. Middle and right columns: surface view of the vessels.