

MODELS, ALGORITHMS AND APPLICATIONS IN VASCULAR IMAGE SEGMENTATION

Maciej Orkisz¹, Leonardo Flórez Valencia^{1,3}, Marcela Hernández Hoyos²

¹ *CREATIS-LRMN, Université Lyon 1, INSA-Lyon, CNRS (UMR 5220) and INSERM (U630) Research Unit, France*

² *Grupo de Ingeniería Biomédica, Grupo Imagine, Universidad de los Andes, Bogota, D.C., Colombia*

³ *Departamento de Ingeniería de Sistemas Pontificia Universidad Javeriana, Bogota, D.C., Colombia*

June 11, 2007

Abstract. A synthesis of the authors' projects in the last decade, in the field of 3D vascular image processing, is provided. This work was motivated by the following applications: display improvement, extraction of geometrical measurements, acquisition optimization, stent-pose planning, phantom generation, blood-flow simulations. The methods are often dependent on the imaging modality and/or on the anatomic region. They involve both: low-level models of intensity patterns and profiles, and higher-level models of cylindrical shapes. Amongst various algorithms used, recursive tracking and fast-marching level-sets are emphasized. Critical analysis of each model and algorithm is done. Problems that remain open, and perspectives associated with the progress of the image acquisition techniques, are listed.

Key words: inertia moments, Hessian, eigen-analysis, centerlines, active contours, deformable models, generalized cylinders, simplex meshes, tracking, fast-marching level sets

1. Introduction

3D magnetic resonance (MR) imaging and computed tomography (CT) develop quickly and play an increasing role in diagnosis, treatment planning and follow up of vascular diseases. Improved acquisition protocols give access to fine resolution images of all vascular territories, recently including even the coronary arteries. These developments also lead to new problems [31], namely huge size of datasets, which lead to tedious visual analysis and to time-consuming processing.

Image processing techniques are used to assist the clinicians: improve the display of the datasets (noise attenuation, vessel enhancement, segmentation) and provide quantitative information. The segmentation is the key to various visualization techniques and to the quantification. Literature in this field is quickly growing. Several surveys report numerous methods [14, 18, 23, 28], but many problems still remain open, *e.g.*: delineation of vessel-wall components, automatic extraction of the vascular tree, etc.

The segmentation of tubular structures such as blood vessels requires the extraction of their axial shape, defined by the centerlines, and of their boundaries. The axial shape, under some conditions, may be extracted independently of the boundary extraction. However, in most cases the centerline is defined relatively to the boundaries and their extraction is interleaved. Depending on the image acquisition type, the boundaries may represent either only the borders of the vascular lumen or both the internal and external border of the vessel wall. When extracting the boundaries, their spatial continuity is to be exploited to cope with noise and with locally weak contrasts. The continuity can be reinforced via the use of an appropriate model.

This article summarizes our experience in 3D vascular image processing during the last ten years, with emphasis on the most recent models and algorithms. In section 2 we quickly describe each of the methods developed or used in our various projects, together with the corresponding applications. Then, we provide a synthesis of the models (section 3) and algorithms (section 4) used, together with a criticism of their strengths and weaknesses. Subsection titles include references to our corresponding previous publications. Work not yet published is highlighted and more details are given.

2. Applications and methods

Our work was motivated by the following applications: display improvement, extraction of geometrical measurements, acquisition optimization in flow imaging, stent-pose planning, phantom generation, blood-flow simulations. The methods used were often dependent on the imaging modality (MR or CT) and/or on the anatomic region. First, MR angiography (MRA) attracted much attention, owing to its non-invasive character. Gadolinium-based contrast agents or blood motion are exploited to increase the contrast between the vascular lumen and the surrounding stationary tissues. In such anatomic regions as head, neck, abdomen and limbs, the stationary tissues can be approximately eliminated by performing a double acquisition before and after the injection of the contrast agent, followed by a digital subtraction of the pre-contrast dataset. More recently, the CT angiography (CTA) gained popularity owing to the reduction of the radiation dose necessary to obtain high-quality images. When processing these images, known correspondence between densities of tissues and intensities expressed in Hounsfield units (HU) is useful. However, some stationary tissues have intensities similar to the contrast-enhanced blood, while calcifications and bones appear brighter than the vessels. Therefore, both visualization and segmentation become more difficult.

2.1. Image improvement

Despite the continuous improvement of the image acquisition devices and protocols, the signal-to-noise ratio often remains locally low. Hence, image denoising and vessel enhancement are useful both to improve the display and to help the segmentation.

2.1.1. Morphology-based denoising and enhancement [6, 12]

In MRA the vascular lumen usually appears brighter than the background, and can be visualized by means of the maximum intensity projection (MIP) technique. Ten years ago, when we initiated our first project in vascular image processing, MIP was the only widely available display mode capable of giving an overview of the vascularization. However, after projection the small low-intensity vessels disappear in noise, while their visualization is important, *e.g.* for the assessment of collateral vascularization of lower limbs in case of arterial obstruction. In order to denoise the dataset before the projection, while preserving the small vessels, we proposed a non-linear anisotropic filter composed of two stages. The first one is used to estimate the local orientation of the structures contained in the image and the second one performs the actual filtering along this orientation [6]. In order to estimate the orientation, a structuring element composed of a set of parallel line segments (sticks, see section 3.2.1) sweeps a number of discrete angles and the following criterion is evaluated for each angle ϕ : $C(\phi) = D - \alpha H$, where D is the average intensity difference between the central stick and its parallel neighbors, while H represents the directional homogeneity and is calculated as the average of intensity standard deviations within each stick. The filtering is performed according to the orientation ϕ that maximizes $C(\phi)$. We proposed a generic directional filter that can perform both denoising and enhancement weighted by appropriate parameters [12]. It computes a truncated mean values in each stick, then it combines them using the Laplacian-based scheme. The filter was clinically evaluated in its purely denoising version. It improved the contrast-to-noise ratio by a factor 2 on average.

2.1.2. Hessian-based enhancement (unpublished)

The visualization of the coronary arteries in both MRA and CTA images is challenging, due to the vicinity of the heart cavities, as well as to the heart motion that hampers the elimination of non-vessel structures. Several filters based on directional second derivatives (see section 3.2.2) have been proposed in the past, in order to highlight the tubular structures (vessels) while attenuating the background [8, 9]. These filters use the eigenanalysis of the Hessian matrix. In a vessel the second derivatives are expected to be close to zero in the direction of the axis, and strongly negative in the plane perpendicular to this axis. Accordingly, the eigen-values ($|\lambda_1| \leq |\lambda_2| \leq |\lambda_3|$) of the Hessian \mathbf{H} should meet the following conditions: $\lambda_3 \approx \lambda_2 \ll \lambda_1 \approx 0$. The output of the filter is set to zero for the voxels having $0 \leq \lambda_3$ or $0 \leq \lambda_2$, and to an appropriate combination of $\lambda_1, \lambda_2, \lambda_3$ otherwise. The actual combination depends on the authors. In order to cope with variable diameters of the vessels, it is necessary to perform the same process at various scales and to choose the scale with the strongest response of the filter. The whole processing is time-consuming and the reported computational times (up to several hours per image volume) were not compatible with physicians' routine.

We attempted to optimize the implementation of this family of filters in two ways.




Fig. 1. Hessian-based enhancement of vascular structures: axial slice of a thorax MR dataset (left), vessels (see arrows) enhanced within the heart sub-volume defined by the rectangle (right).

First, a careful design of computational pipe-lines considerably accelerated the calculus of the Hessian. Second, it can be noted that the vessels typically represent a small percentage of the image volume, and the computational time might be reduced if the voxels having $0 \leq \lambda_3$ or $0 \leq \lambda_2$ were set to zero by some simpler process, while performing the eigen-analysis only for the voxels where the combination of the eigen-values is to be computed. To do this, we used the matrix invariants: trace, determinant and definiteness (since the Hessian is a Hermitian matrix). For example, if $\det(\mathbf{H}) < 0$ and $\text{tr}(\mathbf{H}) > 0$, the considered voxel can be discarded. Indeed, $\text{tr}(\mathbf{H}) = \lambda_1 + \lambda_2 + \lambda_3$, $\det(\mathbf{H}) = \lambda_1 \lambda_2 \lambda_3$ and the above-mentioned condition occurs only when one eigen-value is negative, while the remaining two are positive. When $\det(\mathbf{H}) > 0$ two other cases can be discarded by additionally computing the principal minors of \mathbf{H} , in order to check whether or not \mathbf{H} is definite positive with $\text{tr}(\mathbf{H}) > 0$ (respectively definite negative with $\text{tr}(\mathbf{H}) < 0$). As result, the computational time was further reduced by 25% on average. Now, a $200 \times 200 \times 80$ voxels sub-volume containing the heart within a thoracic scan can be processed in one minute per scale on a Pentium IV, 1.7 GHz laptop. Figure 1 shows coronary arteries enhanced using our implementation of the Frangi's filter [9].

2.2. Pathway extraction for acquisition optimization, display and stenosis quantification

The extraction of the axes of vascular segments has several applications. Knowing the local orientation of the vessels, optimal image acquisition quality can be achieved by an appropriate feedback onto the imaging device settings. The axis can be used to generate a virtual flight-through display or to compute the cross-sectional images essential both for visual evaluation of pathology severity and for more or less automated extraction of

measurements such as lumen diameters and areas, wall thickness, etc. These applications require that the axis orientation be consistent with the local orientation of the vessel. The axis extraction process is usually semi-automatic, since it starts from one or more seed-points interactively selected by the user.

2.2.1. Subtracted magnetic resonance angiography [19, 28, 29, 32]

In the contrast-enhanced subtracted MRA the background is significantly darker than the vascular lumen, and the vascular pathway can be approximately determined by following the local maxima of the intensity. However, to achieve consistent orientations of the axis, we designed a method in which the continuity and smoothness of the axis are controlled by the elasticity and the flexibility of a model (see section 3.4) called extensible skeleton [19]. Starting from a user-selected point the skeleton grows by iteratively adding subsequent axis points within a prediction-estimation scheme (see section 4.1) controlled by a multi-scale analysis of the image moments (see section 3.2.3). The location of the next point is predicted according the local orientation of the tubular structure. The local orientation is estimated by calculating the eigen-vectors of the inertia matrix, where the intensities of the voxels play the role of elementary masses. The eigen-analysis is carried out at multiple scales, in order to cope with varying diameters of the vessel. The coordinates of the predicted point are corrected under the influence of image forces and of prior model shape constraints.

Thus extracted axes were exploited to generate cross-sectional images in which an automatic boundary extraction was performed using adaptive iso-contours. The goal was a quantitative assessment of arterial stenoses, based on cross-sectional diameters and areas [28, 32]. The robustness and accuracy of the method have been evaluated on MRA data of 5 reference phantoms and of 17 patients' carotid arteries. 97% of the centerlines were exploitable in the carotid arteries (100% in the phantoms). Mean difference (\pm standard deviation) between stenosis percentages semi-automatically measured and visually estimated by radiologists was $0.23\% \pm 7.89\%$. The reproducibility of the semi-automatic method was significantly better.

Another application of the extracted axes is the optimization of the acquisition plane for blood flow measurement by phase contrast MRA. Accuracy of 2-dimensional imaging of blood flow depends on the alignment between the vessels and the imaging plane. This is challenging when the flow in several vessels is to be evaluated with one acquisition. The optimal plane is determined by minimizing misalignment between its normal vector and the axes' tangent vectors [29]. The method was evaluated on a phantom and on 35 patients, by seeking the optimal plane for cerebral blood flow quantification simultaneously in internal carotids and vertebral arteries. In the phantom, difference of orientation and of height between known and calculated plane was 1.2 degrees and 2.5 mm respectively. In the patients, all but one axe were correctly extracted and the misalignment of the plane was within 12 degrees per artery.

2.2.2. Computed tomography angiography [26]

Owing to bright structures (bones, calcifications) in CTA, the above-described method is not directly applicable. Based on an empirical study of the CTA image intensities, we proposed a pre-segmentation method using adaptive thresholding. The median axis of thus binarized image is then extracted. As this axis is expected to be located within the binary object, far from its boundaries, a 3D distance map is calculated. The axis is constructed by iteratively adding points corresponding to local maxima of the distance map. The iterative prediction/estimation scheme (see section 4.1) is similar to that described in section 2.2.1. Boundary location is refined by analyzing radial intensity profiles in the vicinity of their intersections with the thresholds (see section 3.1). The axis extraction was evaluated qualitatively in images of 60 patients. While using fixed thresholds only 50% of axes were exploitable, with adaptive thresholds 80% of previous failures were transformed into exploitable axes. The accuracy of boundary extraction was not rigorously assessed yet.

2.3. Quantification of lumen and wall in cross-sectional images

While the focus of the projects described in sections 2.2.1 and 2.2.2 was the extraction of the axes, and the lumen contours were extracted by rather basic methods, the sections 2.3.1 and 2.3.2 report more advanced methods devised for the segmentation of both vascular lumen and wall. Recent evidence indicates that the type of atherosclerotic plaque, *i.e.* the composition and structure of the pathological vascular wall, rather than the degree of stenosis, is an important indicator of the risk of thromboembolic vascular events. It determines the plaque vulnerability to disruption.

2.3.1. High-resolution MRI of vascular wall [24]

Depending on the antenna, acquisition sequence, strength of the magnetic field and other parameters, the resolution of MR images may fall below 0.1 mm. Such images are likely to provide evidence of the vessel wall composition. Although such resolutions are not yet available in clinical routine, several research groups anticipate the development of appropriate image processing tools. We addressed the problem of the extraction of the internal and external boundary of the vessel wall in high-resolution cross-sectional MR images as a preliminary step toward the analysis of the wall composition.

We use an original implementation of active contours (see section 3.4.1) designed to permit an initialization with a single point. Its deformations are calculated on a normalized-length copy of the actual contour, which avoids costly re-discretization and re-computation of the inverse of the system matrix during the growth of the contour. The first contour is initialized by a point within the lumen. When its growth stops, the second contour is generated beyond the first (internal) boundary and grows until the external boundary of the vascular wall is reached.

Accuracy and reproducibility of this method were mainly assessed on high resolution MR images of rabbit aorta. Compared to contours traced by experts, the differences in endoluminal areas and in wall thickness were within the inter- and intra-variability of the experts, while the intra-variability of the method was significantly smaller. Ongoing work attempts to classify the components of the vessel wall, based on multi-spectral MR images, as well as on dynamic images obtained after administration of a contrast agent that progressively penetrates into some components of the wall.

2.3.2. Calcified CTA image analysis [35]

The goal of one of our projects is the generation of patient-specific models of carotid arteries, for computational fluid dynamics simulations. We use CTA images in attempt to extract both wall thickness and calcifications, which is challenging due to current limitations of spatial resolution and of contrast in the vicinity of the vessel wall. However, the processing time does not need to be compatible with clinical routine and user intervention is allowed in order to provide an extra cue in the segmentation method.

First, a search-ray scheme [30] is used to detect discontinuities (see section 3.1). From a point belonging to the lumen, a number of rays is cast and the radial intensity profiles are analyzed. Automatic differentiation between lumen and plaque is based on an empirical study of CTA image intensities [26]. At the end of this phase, each region is represented by a set of points defining a polygon. Then, the polygonal boundaries are refined, in order to improve continuity and to recover possibly missing parts, since in case of concave lumen some of its parts may be "hidden" behind the calcifications. This is done by use of iso-contours. The adaptively calculated iso-values may be interactively adjusted at any moment, as the iso-contours are displayed in real time and the user can check their correctness. Lastly, iso-contours are also used to separate the inner and outer boundary of the wall.

Pre-clinic evaluation of the method is ongoing. It can be noted however that, unlike the lumen and calcification boundaries, the semi-automatic extraction of the outer boundary is not based on image discontinuity evidence, but on radiologists' experience.

2.4. Three-dimensional surface segmentation of cylinders

While the extraction of the axis and of cross-sectional contours only provides measurements for conventional treatment planning and follow-up, the following two methods were devised to extract a complete surface of a given vessel, which is exploitable for the pre-operative simulation of the consequences of various therapeutic choices. They respectively use a discrete model and a continuous model of cylindrical surfaces.

2.4.1. Stent-pose planning [25, 37]

Arterial stenoses and aneurysms are increasingly treated using stents (intravascular expandable grids). We attempted to facilitate the pre-operative choice of the stent's length

and diameter by providing appropriate measurements and by interactively simulating the deployment of the stent within the segmented artery. An angiographic 3D image is first segmented using an active surface, actually a simplex mesh (see section 3.5.1), to create a patient-specific vascular model. A model of a folded stent is then slid along the vessel axis until an interactively chosen delivery location, and the model's geometry is modified to simulate the unfolding of the stent. Lastly, the vessel model is re-meshed to fit the shape of the unfolded-stent model and thus to simulate the shape of the arterial lumen boundaries after stenting. The simulation is purely geometric and does not yet include mechanical interactions between the stent, the vessel wall and the surrounding tissues.

Currently, both segmentation and simulation have only been visually evaluated. In general, the simulation is very useful to assess the positioning of the stent and if the stent diameter fits to the vessel's healthy segments. However, the segmentation is sometimes inaccurate when two vessels are too close and cannot be distinguished because of an insufficient resolution of the images. In some cases we also noticed that the vessel axis is not well centered within the vessel lumen. This occurs when the vessel bends strongly.

2.4.2. Volumetric quantification of lumen in CTA [21, 33, 36, 38]

This method is based on a parametric model of a generalized cylindrical surface [21] (see section 3.5.2), using the state representation of systems. The model consists of an axis and of a surface obtained by continuously sliding, along the axis, a planar contour that remains locally perpendicular to it. Actually, the contour is represented in a reference frame, the origin of which slides along the axis, while one of its constituent vectors remains tangent to the axis. The state variables and matrices describe the evolution of this frame (location and orientation) and of the contour shape between the ends of the axis. Our method identifies the parameters of the model from a series of contours, using a Kalman state estimator (KSE).

The method is initialized by computing an image of intensity gradient magnitude $|\nabla(f(\mathbf{x}))|$ and a line coarsely corresponding to the vessel pathway between a pair of seed points interactively chosen inside the vascular lumen. This line is obtained as follows: 1) a binary image is computed from the input dataset, using two adaptive thresholds (see section 2.2.2), 2) a distance map is calculated within thus extracted approximate vascular lumen, and 3) a minimum-cost path is computed between the seed points [20]. Then the KSE (see section 4.1) is used to sequentially control the extraction of planar contours in consecutive planes orthogonal to the approximate pathway, and to add cylinder pieces (comprised between each pair of contours) to the final model. The contours are described by a Fourier series (FS) decomposition, in which each harmonic of order n corresponds to an ellipse. Each cylinder piece is expected to have constant values of state variables, which is equivalent to locally constant axial curvature and torsion, constant rotation angle of the reference frame with respect to the Frenet frame, and linear variation of the FS coefficients. With constant axial curvature and torsion the axis piece is a helix. Using




Fig. 2. Cross-sectional lumen contours and whole cylinders reconstructed and superimposed onto the original datasets (rendered with shaded surface display). The bright spots are calcifications. The upper and lower cross-sections respectively correspond to the datasets on the left and on the right.

this property, an observation vector is calculated from two consecutive contours, *i.e.* the values of the state variables are deduced from the locations, orientations and shapes of these contours [33]. The observation is then filtered by the KSE to produce an estimate of the cylinder piece. The KSE adapts the progression speed to the complexity of the local vascular shapes (axial and superficial): in complex shapes (*e.g.* high local changes of curvature) the predicted length of the cylinder piece is automatically decreased, in order to reduce the difference between predictions and observations.

The accuracy of the segmentation depends on the accuracy of the contour extraction. Furthermore, as numerous contours are to be extracted, the extraction method has to be fast. We have chosen the fast-marching implementation of the implicit deformable contours (see section 3.4.2). The contour is represented by a front that propagates outwards from the intersection point between the current plane and the curve representing the approximate pathway. Its displacement is driven by the Eikonal equation (see Eq. 17, section 4.2), in which $F(\mathbf{x})$ is a monotonic propagation speed. This monotonicity leads to problems that we attempted to solve in our work [38]. We describe hereafter in details one of its variants that was only presented in a short poster [36].

- *Speed function.* Usual speed functions are exponential: $F(\mathbf{x}) = \exp(-\alpha|\nabla(f(\mathbf{x}))|)$, or sigmoid: $F(\mathbf{x}) = 1 - 1/(1 + \exp(-(|\nabla(f(\mathbf{x}))| - \alpha)/\beta))$. However, these functions are

not appropriate when the contrast is strongly variable, like in vascular cross-sections with calcifications. Indeed, when the front reaches the low-contrast boundaries, it remains relatively far from the boundaries in the highly contrasted zones, and when it reaches the high-contrast boundaries, it already leaks far away through the low-contrast breaches. Instead, we work with a "brake" function $B(\mathbf{x}) = 1/F(\mathbf{x})$ obtained by multiplying the magnitude of the image intensity gradient by three exponential factors $C(\cdot)$: $B(\mathbf{x}) = |\nabla(f(\mathbf{x}))|C(d(\mathbf{x}), \alpha_1, \beta_1)C(f(\mathbf{x}), \alpha_2, \beta_2)C(f(\mathbf{x}), \alpha_3, \beta_3)$. The function $C(\cdot)$ of a scalar u is defined as follows: $C(u, \alpha, \beta) = 1 + \exp(-(u - \alpha)/\beta)$, where α is a shift and β is a signed sharpness parameter. In the first factor, $d(\mathbf{x})$ is a mono-dimensional derivative of $|\nabla(f(\mathbf{x}))|$ in the normal direction with respect to the front, $\alpha_1 = 0$ and $\beta_1 > 0$. In the remaining factors, α_2 and α_3 are respectively equal to the lower and upper limit of the typical range of luminal intensities in CTA images [26], $\beta_2 > 0$ and $\beta_3 < 0$. Thus, the front slows down not only for large values of $|\nabla(f(\mathbf{x}))|$, but also after passing the discontinuity ($d(\mathbf{x}) < 0$) and beyond the typical range of luminal intensities.

- *Stopping criterion.* We want to detect the moment when the contour begins to "leak" out of the lumen, so as to avoid wasting time for further propagation. The criterion is heuristic, based on experimental observations of the FS representation of the contours. The shape of a front propagating within the vascular lumen cross-section is initially circular, which can be represented by the fundamental harmonic alone ($n = 1$). The other harmonics ($n > 1$) increase slowly, as long as the front is confined within the lumen, but they "explode" when the front begins to leak out of the lumen, since the shape becomes very irregular. At that moment, the standard deviations σ_n of these harmonics become larger than their mean values μ_n . We track the evolution of the harmonics up to the order 3, and the front propagation is stopped when both the 2nd and 3rd order harmonics meet the criterion: $\sigma_n > \mu_n$.
- *Backtracking.* With the above described modifications, we can expect that a front fitting to the actual vessel boundary exists and corresponds to a propagation time T_{opt} slightly smaller than the stopping time T_{stop} . As all the front locations covered during the propagation are stored, we perform a backtracking, in order to select the front that maximizes the average gradient magnitude on its perimeter. If we note $\Omega(T)$ the set of points \mathbf{x} that compose the front at time T , and $S(T) = \sum_{\mathbf{x} \in \Omega(T)} |\nabla(f(\mathbf{x}))|$, then: $T_{opt} = \operatorname{argmax}_T \{S(T)/\operatorname{card}(\Omega(T))\}$.

We evaluated the proposed method on CTA images representing 6 physical phantoms (having circular reference sections 6 mm in diameter, and two stenoses of variable shape, position, eccentricity and known severity) and on 76 CTA datasets representing the carotid arteries of patients (fig. 2). In the reference sections of the phantoms the mean error and standard deviation of the measured diameters were $0.14 \text{ mm} \pm 0.13 \text{ mm}$, while the stenosis percentage error was equal to $3.30\% \pm 5.06\%$. In the images of patients, the cylinders were visually inspected and scored from 0 (false) to 3 (perfect). 53 cylinders

obtained the score 3, 13 cylinders had perfect shape in healthy regions and some anomalies in diseased regions mainly owing to remaining difficulties to cope with calcifications, 10 cylinders were false, due to a bad shape of the initial minimal path, which occurs when the external carotid artery is locally "stuck" to the internal carotid artery due to insufficient image resolution. The current implementation of the KSE, dependent on the initial approximate pathway, is responsible for this limitation.

2.5. Vascular tree segmentation

The extraction of the vascular tree is helpful to remove the surrounding tissues and thus improve the display. It is also the first step toward an automated search of possible pathologies, in order to attract the attention of the physician and thus speed-up the analysis of huge datasets generated by modern scanners. The methods described in the following two sections focus on the skeleton of the vascular tree, while the volumetric segmentation is very rough. Both share the same framework, recursively extracting the branches (see section 4.1) by use of a spherical volume of interest (cell), which moves step-by-step along the vessels. At each step, the cell is initially centered in a predicted axis point. Its location and radius are iteratively modified, so as to encompass the vessel. The content of the cell is then segmented, in order to coarsely separate the vessel from the background. Then the connected components (CC) of the segmented sub-volume on the surface of the cell are analyzed. The number of CCs is interpreted as follows: 1 - end of the branch, 2 - continuation of the current branch, 3 or more - bifurcation. New candidate points are predicted in consequence at the centers of the CCs, and the current center of the cell is added to a graph representing the vascular tree axes. Similarly, the segmented content of the cell is added to the union of the sub-volumes segmented in the previous iterations, so as to construct an approximate outline of the vascular lumen.

2.5.1. Well-contrasted images [27]

By "well-contrasted" we mean that the vessel can be coarsely separated from the surrounding tissues by thresholding (possibly adaptive and local, *e.g.* half the local maximum intensity in contrast-enhanced subtracted MRA). The adaptation of the cell size and position is based on the analysis of the surface CCs. Typically, at the beginning of its growth, the cell is contained within the vessel, so only one "vessel" CC and zero "background" CC is found. Then it begins to protrude from the vessel on one side, thus creating a "background" CC. The location of the cell is then corrected by pushing its center in the opposite direction. The number of CCs may vary during the growth, until it stabilizes after a number of iterations, when the cell encompasses the vessel.

Applied to 16 MRA images from patients, representing the neck (carotid, vertebral, basilar arteries, etc.) and the aorto-iliac region (with renal, mesenteric arteries, etc.), the method detected all the principal branches of the vascular tree and the bifurcations in which the secondary branch had at least 5% of the cross-sectional area compared to

the main branch. Some branches were not extracted owing to a strong local signal drop in severe stenoses. Conversely, a few false bifurcations have been detected, where two vessels were very close to each other and the assumption of good contrast did not hold. On average, the processing speed was 50 axis points per second.

2.5.2. Coronary arteries [34]

In the datasets representing the coronary arteries, the surrounding structures may have both lower and higher intensities (respectively parenchyma and calcifications), as well as the same intensity range (heart ventricles filled with blood). Thresholding is therefore not applicable. Both the cell size adaptation and the detection of the vessel end are more difficult. To cope with these difficulties, we proposed two main innovations.

- Thresholding is replaced by K-means clustering with an original metric. While the metric in the background cluster remains the intensity difference between the point and the mean value of this cluster, the metric of the foreground cluster is a combination of the intensity difference and the inertia moment with respect to the cluster's principal inertia axis (see section 3.2.3) in the XYZ-space, weighted by a coefficient φ . This enforces the cylindrical shape of the segmented volume by including in the foreground cluster the voxels that are close to its axis even if their intensity sensibly differs from the mean intensity of the vessel, and by excluding the voxels that are located further from the axis even if they have similar intensity. In this way, we expect to separate the arteries of interest from other bright structures that appear "stuck" to them in the images, such as veins or heart cavities.
- The shape of the segmented sub-volume is checked against a geometric model (see section 3.3), in order to validate it as part of a vessel. Otherwise, *i.e.* when it is misshapen, we stop the extraction of the current branch.

We also changed the implementation of the correction of the cell location and size. A distance map is calculated within the segmented volume and the position of this map's maximum is used to update the candidate's location if necessary. Then the ratio of the segmented volume to the total volume of the cell is compared with the theoretic ratio for a cylinder having a radius equal to half the radius of the cell. We increase or decrease the cell radius accordingly. The last step of the method is the local improvement of the axis. Since the cell must locally encompass the vessel and the predicted axis points are located on the cell surface, the distances between consecutive points are quite large. An advantage is a quick analysis of the vascular tree. A drawback however is that straight line segments joining two points are sometimes not entirely located within the segmented volume. To resolve this problem, a minimum cost path [20] is computed between the previous and current point, using a cost function based on the distance map. Thus calculated axis always lays within the vessel and closer to its center.

The method is sensitive to the variation of the parameter φ . When its value is too small the adjacent structures are not eliminated, so the number of spurious branches is

higher. When the value is too high the bifurcations are considered as adjacent structures and they are eliminated. The missed branches may be recovered by use of additional seed points. On average, the processing speed was 1 axis point per second.

3. Models

In medical image processing, underlying models are necessary to cope with incomplete, low-contrast and noisy data. This section gives an overview of the models used in the above-described methods. These models are classified according to increasing complexity. We begin with brightness patterns and simple geometric models useful (under some conditions) to identify points located either on the vessel boundary or on the vessel axis. Then we follow with active contours used to extract the cross-sectional boundaries. At the end, we describe complex models capable of representing a complete cylindrical shape. From now on, critical comments are provided for each model and algorithm.

3.1. Intensity profiles [26, 36, 35]

The intensity profile is the simplest model, mainly used in the analysis of cross-sectional images. It consists of a straight-line segment (often called search ray) which can be manually traced by the user or generated in a ray casting scheme [30]. Depending on the algorithm, rays may have different line width and may be cast following different patterns (radial, spherical, cylindrical, normal to a more complex curve or surface, etc.), and the intensity values along the ray may be interpolated or smoothed.

The profile models embed prior knowledge of intensities, contrasts and gradients, dependent on the imaging modality. In the methods described in sections 2.2.2 and 2.3.2 we used radial profiles, while in section 2.4.1 the rays were cast along the normal vectors from the nodes of the simplex surface mesh. In application to CTA images, these methods rely on the typical profiles experimentally determined in a previous study [26].

The main advantages of search ray based methods are the simplicity of implementation, the reduced computational time and the directional nature of the rays: the search rays can be defined according to the expected shape of the object, so that the image data be explored in an object-adapted way. However, several problems arise: 1) to interpret the profiles when the starting point of the rays is not guaranteed to lay within the vessel lumen and when their direction is not orthogonal to actual boundaries, 2) to manage the spatial continuity of the interfaces, 3) to reach certain parts of concave regions, when casting rays from one point, or 4) to avoid self-crossing of the rays cast orthogonally to an evolving complex interface. Nevertheless, this simple approach may be used to initialize a more sophisticated model.

3.2. Brightness patterns

Almost all the existing methods in vascular image processing are based on the assumption that the vessel lumen is brighter than its neighborhood, which is consistent with the efforts done in the angiography imaging techniques. A few acquisition techniques however, particularly in MRA, provide "black-blood" images. This can probably be overcome by a global inversion of the image dynamics. Nevertheless, in some cases this assumption is locally false, *e.g.*: in CTA images of the head, several arteries locally pass through holes in the skull and the bone is much brighter than the vessel (vascular wall is often too thin, compared with the image resolution, to give rise to a darker interface between the lumen and the bone). Similarly, concentric calcified atherosclerotic plaques, as well as stents, appear as direct neighborhood locally brighter than the lumen. These exceptions limit the application of the models described below.

3.2.1. Morphological model [6, 12]

As mentioned in section 2.1.1, the model consists of 9 parallel line segments (sticks) of fixed length. If the sticks are aligned with a vessel, the longitudinal variation of intensity in each stick is expected to be small, while the difference between the central stick and its neighbors should be large. Let \bar{f}^k be the truncated mean intensity within the k -th stick ($k = 0$ corresponds to the central stick), and let $\tilde{f} = 8\bar{f}^0 - \sum_{k=1}^8 \bar{f}^k$ be the output of the enhancement filter. Ideally, if the respective mean intensities of the vessel and of the background are f_V and f_B , we should have $\tilde{f} = 8(f_V - f_B) \gg 0$ when the central stick is aligned on a thin vessel, and $\tilde{f} = 0$ when it is entirely contained in the background, which leads to vessel enhancement with background suppression. The optimal orientation is sought by an exhaustive exploration of the discrete space of possible solutions, *i.e.* by rotating the sticks in all possible discrete directions.

According to our experience, this procedure leads to a good alignment of the model with the structures perceptible in the image: axes of small vessels and boundaries of larger objects. In the absence of such structures, the model aligns itself on random noise configurations. The main drawback of this model is an inevitably time-consuming implementation of the search of the optimal orientation. For this reason, although a multi-scale generalization would be straightforward (by use of sticks having variable length and spacing), only a mono-scale version is used, so that only small vessels can be enhanced. Furthermore, this model hardly fits the vessels with locally strong curvature. The model does not sufficiently distinguish "line-like" structures from "plate-like" structures and edges. Hence, while the background is well suppressed, thin objects and the edges of highly contrasted objects remain in the filtered image.

3.2.2. Second derivatives

Another way of characterizing the intensity variations, small along the vessel and large in the cross-sectional plane, is the use of second derivatives (see section 2.1.2). Unlike

the above-described model, this approach does not require an exhaustive search of the optimal orientation, although such implementation has been proposed in the past [4]. Instead, the Hessian matrix is built with 6 partial derivatives that are usually computed by convolving the image with the appropriate derivatives of a Gaussian kernel, the standard deviation (σ) of which determines the scale ($f_{xx} = f \otimes G_{xx}^\sigma, f_{xy} = f_{yx} = f \otimes G_{xy}^\sigma, f_{xz} = f_{zx} = f \otimes G_{xz}^\sigma, f_{yy} = f \otimes G_{yy}^\sigma, f_{yz} = f_{zy} = f \otimes G_{yz}^\sigma$ and $f_{zz} = f \otimes G_{zz}^\sigma$):

$$\mathbf{H} = \begin{bmatrix} f_{xx} & f_{xy} & f_{xz} \\ f_{yx} & f_{yy} & f_{yz} \\ f_{zx} & f_{zy} & f_{zz} \end{bmatrix}. \quad (1)$$

Its eigen-vector associated with the eigen-value λ_1 (having the smallest absolute value) determines the axial orientation. A "vesselness" criterion is computed, based on the eigen-values, *e.g.* the Frangi's criterion [9]:

$$\tilde{f} = [1 - \exp(-R_A^2/2a^2)] \exp(-R_B^2/2b^2) [1 - \exp(-S^2/2c^2)], \quad (2)$$

measures the cross-sectional circularity $R_A = |\lambda_2/\lambda_3|$, axial elongation $R_B = |\lambda_1|/\sqrt{|\lambda_2\lambda_3|}$ and local non-uniformity $S = \sqrt{\lambda_1^2 + \lambda_2^2 + \lambda_3^2}$ (a, b and c are weighting parameters).

Since a computationally efficient implementation is possible (see section 2.1.2), multi-scale analysis is used, and the optimal scale (with largest response of the filter) provides an estimate of the local vessel radius, which is an interesting feature. However, this model may locally be inappropriate, when singularities occur owing to stenoses and bifurcations. When implementing a "vesselness" criterion, the choice of weighting parameters may be difficult, *e.g.*: in Frangi's filter, while the parameters a and b may remain fixed, c depends on the image dynamics. Furthermore, strongly contrasted "plate-like" objects often remain in the filtered image.

3.2.3. Inertia moments [19, 28, 32]

In the method used in section 2.2.1, but also in [22], the underlying model is the so-called equivalent ellipsoid, *i.e.* elongated object that revolves with minimum inertia moment $I_{\mathbf{a}}$ around its principal axis \mathbf{a} . Equations from mechanics are used with image intensities instead of mass density. The matrix of inertia:

$$\mathbf{I} = \begin{bmatrix} I_{xx} & I_{xy} & I_{xz} \\ I_{yx} & I_{yy} & I_{yz} \\ I_{zx} & I_{zy} & I_{zz} \end{bmatrix} \quad (3)$$

with respect to a reference frame, the origin of which is the gravity center, is defined by the moments of inertia I_{xx}, I_{yy}, I_{zz} with respect to the axes of the reference frame, and

by the products of inertia $I_{xy} = I_{yx}$, $I_{xz} = I_{zx}$ and $I_{yz} = I_{zy}$:

$$\begin{aligned}
 I_{xx} &\equiv \sum_{\mathbf{x}_j \in S(\mathbf{x}, \rho)} f(\mathbf{x}_j)(y_j^2 + z_j^2), & I_{yy} &\equiv \sum_{\mathbf{x}_j \in S(\mathbf{x}, \rho)} f(\mathbf{x}_j)(z_j^2 + x_j^2), \\
 I_{zz} &\equiv \sum_{\mathbf{x}_j \in S(\mathbf{x}, \rho)} f(\mathbf{x}_j)(x_j^2 + y_j^2), \\
 I_{xy} &\equiv - \sum_{\mathbf{x}_j \in S(\mathbf{x}, \rho)} f(\mathbf{x}_j)x_j y_j, & I_{xz} &\equiv - \sum_{\mathbf{x}_j \in S(\mathbf{x}, \rho)} f(\mathbf{x}_j)x_j z_j, \\
 I_{yz} &\equiv - \sum_{\mathbf{x}_j \in S(\mathbf{x}, \rho)} f(\mathbf{x}_j)y_j z_j,
 \end{aligned} \tag{4}$$

calculated within a spherical cell $S(\mathbf{x}, \rho)$ centered in $\mathbf{x} = (x, y, z)$ and with radius ρ . The moment of inertia with respect to \mathbf{a} is: $I_{\mathbf{a}} = \mathbf{e}_{\mathbf{a}}^T \mathbf{I} \mathbf{e}_{\mathbf{a}}$, where $\mathbf{e}_{\mathbf{a}}$ is the unit vector in the direction \mathbf{a} . The orientation of \mathbf{a} can be inferred from the eigen-analysis of \mathbf{I} . Let $\lambda_1 \leq \lambda_2 \leq \lambda_3$ denote the eigenvalues of \mathbf{I} and $\mathbf{e}_1, \mathbf{e}_2, \mathbf{e}_3$ be the associated eigenvectors. Like in section 3.2.2, the axis orientation is given by \mathbf{e}_1 , while \mathbf{e}_2 and \mathbf{e}_3 define the orthogonal plane. When the cell contains a thin cylinder with a circular cross-section, then $\lambda_1 \ll \lambda_2 \cong \lambda_3$. Conversely, when the content of the cell is isotropic (no privileged orientation) the inertia moments are similar for any orientation ($\lambda_1 \cong \lambda_2 \cong \lambda_3$).

The advantage of this approach is its lower sensitivity to noise, since integrals (sums) are computed instead of derivatives. Its drawback is high sensitivity to other bright structures that may be contained within the cell, particularly in CTA images. For this reason, in section 2.5.2 the model based on inertia moments is combined with a kind of outliers rejection: the points located far from the currently estimated inertia axis and/or from the current mean intensity of the "vessel" cluster are discarded. However, the latter combination is uneasy to balance.

3.3. Elementary geometric shapes [34]

While the models from the sections 3.2.1 and 3.2.2 represent directly the intensity patterns likely to correspond to vascular lumen in angiography images, the method described in section 2.5.2 checks the locally segmented sub-volume against a higher-level geometric model. The model is adaptively constructed from the data as a union of a sphere, with radius equal to the distance map value at the cell center, and of several truncated cones. Each conical frustum corresponds to one CC on the cell surface. Its axis is the line between the centers of the cell and of the CC. One of its radii is the radius of the central sphere and the other radius is the value of the distance map at the center of the CC. Once the model is constructed, its intersection with the segmented volume is computed, as well as their union, and the ratio of these values is compared to a threshold.

This model has two advantages: 1) it is adaptive, and 2) it is capable of approximately representing both simple vessels (possibly curved) and bifurcations. Its drawbacks however are also two: 1) it can hardly represent pathologic shapes, and 2) the choice of the threshold value becomes uneasy for small vessels represented by few voxels.

3.4. Deformable contours

The models described in this section represent an interface between regions in a N -dimensional image. This interface moves until it adheres to image discontinuities. Its motion can be analyzed in two ways: by placing "markers" on the interface and tracking them through the time (Lagrangian approach) or by placing, on the image grid, "observers" that memorize the time when the interface passes through their location. Hereafter, we describe closed planar (2-dimensional) active contours used to extract the cross-sectional boundaries of the vessels. Nevertheless, the extensible skeleton model of the vessel axis (see section 2.2.1) is an open contour having the same properties as the conventional snake, and only its algorithmic implementation differs. Furthermore, the same models can be extended to the 3D space. However, the specificity of the cylindrical shapes requires an additional axial constraint, which justifies a separate section (3.5).

3.4.1. Normalized snake - Lagrangian approach [24]

A snake is a parametric curve $\mathbf{g}(s, t) = (x(s, t), y(s, t))^T$ ($s \in [0, 1]$ is the arc length), which evolves in time t and space, under the action of external forces and the reaction of internal forces, until equilibrium is reached. The deformation is achieved via solving Euler-Lagrange equations in attempt to minimize an energy objective function: $E(\mathbf{g}) = E_{int}(\mathbf{g}) + E_{ext}(\mathbf{g})$. This concept was introduced in image processing [1], in order to ensure continuity and smoothness of the detected boundaries, by use of an internal energy that imposes constraints on the first and second derivatives of the curve:

$$E_{int}(\mathbf{g}) = E_{elast}(\mathbf{g}) + E_{flex}(\mathbf{g}) = \alpha \int_0^1 \left| \frac{\partial \mathbf{g}(s, t)}{\partial s} \right|^2 ds + \beta \int_0^1 \left| \frac{\partial^2 \mathbf{g}(s, t)}{\partial s^2} \right|^2 ds, \quad (5)$$

where α and β respectively control the elasticity and the flexibility of the resulting curve. External forces are represented by a potential energy $E_{ext}(\mathbf{g}) = \int_{s=0}^1 P(\mathbf{g}(s, t)) ds$.

These are mainly image forces usually designed to attract the snake toward strong intensity gradients. Balloon force was added [2], in order to inflate the contour outwards, so that it can grow when initialized far from the boundaries (*i.e.* in a region with gradients close to zero). The link between potentials $P(\mathbf{g})$ and forces is given by: $\mathbf{f}(\mathbf{g}) = -\nabla P(\mathbf{g})$. The energy minimization is a temporal and spatial discrete process, using a finite number L of snake points $\mathbf{g}(i, k)$ that approximate the snake as a polygon. The partial derivatives in eq.(5) are approximated by finite differences $\partial \mathbf{g}(s, t) / \partial s \approx (\mathbf{g}((i+1), k) - \mathbf{g}((i-1), k)) / 2h$, with h the discretization step of the contour. The snake is iteratively deformed, explicitly using the external forces $\mathbf{f}_{ext}(\mathbf{g})$ at each point of the snake, while implicitly minimizing the internal energy. The associated Euler equations are solved by matrix inversion, which leads to the evolution equation

(where \mathbf{U} stands for a diagonal unit matrix):

$$\mathbf{g}(i, k) = [\gamma \mathbf{U} + \mathbf{A}]^{-1} [\gamma \mathbf{g}(i, (k-1)) + \mathbf{f}_{ext}(\mathbf{g}(i, (k-1)))] . \quad (6)$$

The matrix \mathbf{A} represents a discretized formulation of the internal energy. The damping parameter γ controls the deformation magnitude at each iteration. The discretization step h depend on the number of snake points L and, initially, is equal to the distance between the snake points. The evolution equation (6) is only valid if h remains unchanged. After each iteration however, the length of a snake grows due to the external force, and so does the distance between the snake points. The internal energy, especially $E_{elast}(\mathbf{g})$ associated with the contour's tension, grows with the snake length and may stop the snake before the boundary is reached. Hence, the model needs to be resampled with a new (higher) number of points or/and a new step h . The $L \times L$ matrix \mathbf{A} , as well as its inverse have to be recomputed. This is a time-consuming task which limits the applications of the classical snake model.

We proposed a new numerical scheme for the snake energy minimization, which preserves the validity of eq.(6) despite the evolution of the snake's size. The matrices $[\lambda \mathbf{U} - \mathbf{A}]$ and $[\lambda \mathbf{U} - \mathbf{A}]^{-1}$ are computed only once, and a fixed number of points is used in all the iterations. The idea resides in scaling the snake size at each iteration, in order to normalize the internal energy of the model with a fixed discretization step $h' = 1$. The scaled snake $\mathbf{g}'(i, (k-1))$ has a normalized length equal to the number of its points. It is deduced from the actual snake by use of the average distance \bar{h}_{k-1} between its points:

$$\mathbf{g}'(i, (k-1)) = \mathbf{g}(i, (k-1)) / \bar{h}_{k-1} . \quad (7)$$

The scaled snake is deformed according to eq.(6), using the external forces \mathbf{f}_{ext} from the non-scaled snake. Its deformation vectors are computed with a normalized internal energy, and the deformation is directly applied on the non-scaled snake. Lastly, the snake points are evenly redistributed along the snake polygonal outline.

Unlike the basic snake, which is strongly dependent on its initialization, the addition of the balloon force permits an initialization with a single point, but adds an additional parameter. Our implementation reduced both the difficulty of parameter tuning and the computational time. However, in its current version, the contour can only grow, which may lead to false results if locally weak boundaries are missed.

3.4.2. Geodesic contour - Eulerian approach [36]

Here an interface is represented by the zero-level of a higher dimensional level-sets function ψ (3D function for 2D contours) describing a $(N+1)$ -dimensional hypersurface. One can easily imagine the 2D contour as an intersection between a horizontal plane and a 3D surface. The shape of the contour depends on the height where the plane is located, and it can split into several independent sub-contours if the shape of the surface is a

little bit complex. This possibility of topology changes is one of the most important features of the level-sets approach. In our application however, the most important is the possibility of representing irregular pathologic cross-sectional shapes, with possibly sharp edges, while the level-sets function is continuous and differentiable. This property clearly differs from snakes that include an explicit smoothness constraint. However, the lack of this constraint becomes a drawback in noisy and poorly contrasted data.

The motion of the zero-level is embedded in the displacements of the hypersurface, and locally follows the direction normal to the contour. The normals, as well as curvatures can easily be inferred from the level-sets function. Detailed description with all formulas can be found in [10]. In section 4.2 we only quickly comment two types of algorithms for the contour evolution.

3.5. Complex cylinders

Conventional deformable surfaces tend to the spherical shape in the absence of external constraints. This is not the most natural solution when the model is expected to adhere to a cylindrical surface. In order to preserve the cylindrical shape, the surface can be attached to an axis.

3.5.1. Deformable simplex cylinder (DCS) [25, 37]

Both the centerline and the surface of the vessel can be represented using N -dimensional k -simplex meshes [11]. In this discrete representation each N -dimensional vertex has exactly $k+1$ neighbor vertices. To represent the surface of the vessel we use a 3D 2-simplex mesh $\mathcal{S}_v = \{\mathbf{g}_j \in \mathbf{R}^3, 0 \leq j < J\}$, while a 3D 1-simplex mesh with radius r_i associated to each vertex \mathbf{a}_i is used to represent the vessel axis $\mathcal{C}_v = \{(\mathbf{a}_i, r_i) \in \mathbf{R}^3 \times \mathbf{R}^+, 0 \leq i < I\}$. The complete model includes \mathcal{C}_v , \mathcal{S}_v and the spatial relationships between \mathbf{a}_i and \mathbf{g}_j .

According to the classical deformable models framework, the mesh deformation is controlled by two types of forces: 1) internal $\mathbf{f}_{int}(\mathbf{g}_j)$, the computation of which benefits from the specificity of the simplex-meshes formalism and uses a criterion enforcing the regularity of the surface curvature, and 2) external $\mathbf{f}_{ext}(\mathbf{g}_j)$ that attract the model towards the vessel boundary in the image. Each surface vertex \mathbf{g}_j is associated with the 3 closest axis vertices. When the surface undergoes a deformation, the axis bends accordingly through an external force resulting from the surface forces reported onto the axis. Conversely, the centerline bending is reported onto the surface to ensure a homogeneous deformation behavior of the cylinder surface. The corresponding external force acting on the surface is a weighted sum of an axial component $\mathbf{f}_{axial}(\mathbf{g}_j)$ (each vertex tends to follow the axis global motion) and of a radial component $\mathbf{f}_{radial}(\mathbf{g}_j)$ (each vertex tends to align on a circle around the centerline). The vertices iteratively evolve according to an equation similar to eq.(6) but involving all these forces.

The simplex-meshes composing the DCS model have a number of useful features: fast computation, possibility of managing the surface and the centerline within the same

formalism, control of the 3D continuity and regularity of the surface, duality with triangulation (easy display). However, like many deformable models, the DCS model requires an initialization close to the final solution. In particular, the initial axis should fall within the vessel lumen. Furthermore, the parameter tuning is difficult and the implementation of image forces based on a search-ray strategy (see section 3.1) sometimes leads to instabilities. This model is very well suited to the simulations, but much more difficult to use in segmentation.

3.5.2. Right Generalized Cylinder (RGC) models [21]

A RGC is defined by its axis and by a stacking of (possibly complex-shaped) planar contours orthogonal to this axis ("right" refers to the angle between the axis and the contour plane). Such models, in association with tracking algorithms (see section 4.1), are frequently used in 3D vessel segmentation, since they directly respond to the medical need of the cross-sectional display and measurements. We also used it in the interactive simulation of stent deployment (see section 2.4.1). Usually, the stack of contours is discrete [5, 7, 13] without formal link between the shapes of consecutive contours.

However, the RGC state model used in section 2.4.2 describes a continuous surface \mathcal{S} attached to the generator axis \mathcal{H} . While the state formalism usually describes the system evolution across time, in our case t is a longitudinal parameter describing the evolution along the axis \mathcal{H} . Each contour composing the surface \mathcal{S} is represented by $2N + 1$ complex coefficients of its FS decomposition, N is the allowed number of harmonics, which determines the level of details [3]. The complete RGC model is constructed by a finite set of M pieces, each one defined by the state vector \mathbf{x}_j ($0 \leq j < M$):

$$\mathbf{x}_j = [\kappa_j \cos(\nu_j), \kappa_j \sin(\nu_j), \tau_j, \Delta_j, \{[z_l, \zeta_l]\}_j]^T, \quad (8)$$

where $\kappa_j \in \mathbf{R}^+$, $\tau_j \in \mathbf{R}^+$, $\nu_j \in [-\pi, \pi]$ and $\Delta_j \in \mathbf{R}^+$ are respectively the piecewise constant curvature, torsion, rotation and length parameters describing the piece's axis; and $\{[z_l, \zeta_l] \in \mathbf{C} \times \mathbf{C}\}_j$ is the set of ordered complex pairs representing the first contour's FS coefficients, and the corresponding linear evolution coefficients $\zeta_l = dz_l/dt$. The last contour of the j -th cylinder piece becomes the first contour of the $(j+1)$ -th piece. According to the state formalism, transition operators are also defined:

$$\mathbf{Tr}(t_0, t_1) = \mathbf{\Gamma}^{-1}(t_0) \cdot (\mathbf{h}(t_1) - \mathbf{h}(t_0)) \quad (9)$$

grants access to the translation transformation between axis points $\mathbf{h}(t_0)$, $\mathbf{h}(t_1)$, and

$$\mathbf{\Phi}(t_0, t_1) = \mathbf{\Gamma}^{-1}(t_0) \cdot \mathbf{\Gamma}(t_1) \quad (10)$$

expresses the transformation between orthonormal frames $\mathbf{\Gamma}(t_0)$, $\mathbf{\Gamma}(t_1)$ that orient contour planes orthogonal to \mathcal{H} at points $\mathbf{h}(t_0)$, $\mathbf{h}(t_1)$. These operators are linked with the parameters of a given cylinder piece as follows:

$$\mathbf{\Phi}_j(\theta_1, \theta_2) = \exp(\mathbf{\Psi}_j(\theta_2 - \theta_1)), \quad (11)$$

$$\Psi_j(\theta) = \begin{bmatrix} 0 & -\kappa_j(\theta) \cos(\nu_j(\theta)) & \kappa_j(\theta) \sin(\nu_j(\theta)) \\ \kappa_j(\theta) \cos(\nu_j(\theta)) & 0 & -\tau_j(\theta) - \frac{d\nu_j(\theta)}{d\theta} \\ -\kappa_j(\theta) \sin(\nu_j(\theta)) & \tau_j(\theta) + \frac{d\nu_j(\theta)}{d\theta} & 0 \end{bmatrix}, \quad (12)$$

$$\begin{aligned} \mathbf{Tr}_j(\theta_1, \theta_2) &= [\mathbf{U}(\theta_2 - \theta_1) + \\ &\quad \frac{\Psi_j}{|\mathbf{p}_j|^2} (1 - \cos(|\mathbf{p}_j|(\theta_2 - \theta_1))) + \\ &\quad \frac{\Psi_j}{|\mathbf{p}_j|^3} (|\mathbf{p}_j|(\theta_2 - \theta_1) - \sin(|\mathbf{p}_j|(\theta_2 - \theta_1)))] [1 \ 0 \ 0]^T, \end{aligned} \quad (13)$$

$$\mathbf{p}_j = [\tau_j, \kappa_j \sin(\nu_j), \kappa_j \cos(\nu_j)]^T. \quad (14)$$

The state formalism perfectly corresponds to the idea of tracking (see section 4.1). The continuous model is helpful in exploiting the underlying anatomical continuity, and provides closed-form expressions for the computation of cross-sectional and volumetric measurements. However, its capability of representing pathologic vessels is limited, due to its monotonic evolution along t and to the uniqueness of the contour attached to $\mathcal{H}(t)$ for each value of t . Clearly, large aneurysms (vessel wall protrusions) and arterial dissections with more than one channels where blood can circulate, cannot be fully represented.

4. Algorithms

Our methods make use of standard mathematical techniques and algorithms, such as the eigen-analysis, optimization (particularly in energy minimization), minimum-cost path search, etc. We mentioned these techniques in the appropriate sections. Hereafter, we only describe and criticize the algorithms into which we brought some contributions.

4.1. Tracking [19, 26, 27, 28, 29, 32, 33, 34]

Many existing methods for axis extraction in tubular objects (*e.g.* [13, 15, 16, 17]) follow an iterative process, which starts from a (frequently) user-selected point and, in every iteration, extends the axis by one point within a prediction-estimation scheme. We proposed a generic object-oriented programming framework [27], in which these methods can be efficiently implemented by modifying some components, mainly the prediction and the correction. New algorithms can thus easily be derived (see sections 2.2.1, 2.2.2 and 2.5). The tracking process can be summarized as follows (fig. 3):

- For each candidate taken from a stack of points to be processed, the correction of its location is done if the point is not close enough to the center of the vessel axis defined as the center of the local cylindrical structure. The corrected candidate is pushed back onto the stack. The first candidate may be manually selected by the user.
- If the candidate does not need to be corrected it is checked against a "vesselness" criterion. If it does not pass the "vesselness" proof it is discarded. Otherwise it is added to a graph that represents the vascular tree axis. This test is optional.

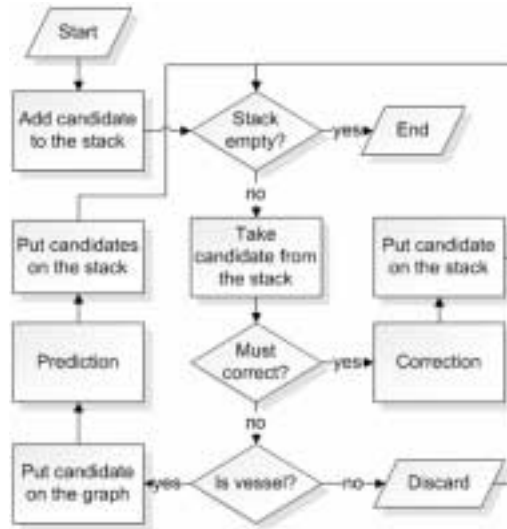


Fig. 3. General framework for recursive axis extraction

- Several (typically between zero and two) new candidate points are predicted and pushed into the stack to be processed later. Hence, the algorithm recursively tracks the branches and detects the bifurcations (when two candidates are predicted). The detection of bifurcations is optional. Frequently, the algorithm implementation is designed to track one single vascular segment and only one candidate is predicted at each iteration. The process finishes when there are no more points in the stack.

The Kalman state estimator (KSE) can also be used in the context of tracking (see section 2.4.2), since it manages the prediction/estimation strategy in a rigorous way. Based on a M -dimensional discrete-time measurement: $\mathbf{m}_i = \mathbf{H}\mathbf{x}_i + \mathbf{v}_i$, it addresses the general problem of the estimation of an N -dimensional state of a process that is governed by the linear stochastic difference equation:

$$\mathbf{x}_i = \mathbf{A}\mathbf{x}_{i-1} + \mathbf{B}\mathbf{u}_i + \mathbf{w}_{i-1}, \quad (15)$$

where \mathbf{u} is an optional control input. The random variables \mathbf{w} and \mathbf{v} respectively represent the process and measurement noises assumed to be statistically independent between them, white and with normal probability distributions: $p(\mathbf{w}) \sim \mathcal{N}(0, \mathbf{Q})$ and $p(\mathbf{v}) \sim \mathcal{N}(0, \mathbf{R})$, where \mathbf{Q} and \mathbf{R} are the respective covariance matrices.

While the algorithm is standard, the key to the use of the KSE is an appropriate definition of the following matrices and vectors \mathbf{A} , \mathbf{B} , \mathbf{H} , \mathbf{Q} , \mathbf{R} , \mathbf{u} , \mathbf{m} , and of the initialization of the state $\hat{\mathbf{x}}_0$ and noise \mathbf{P}_0 . In our application (see section 2.4.2) we made the

following choices: $\mathbf{m} \equiv \mathbf{x}$ is the parameter vector of the RGC state model (eq. 8, section 3.5.2); $\mathbf{B} = \mathbf{U}$, $\hat{\mathbf{x}}_0 = \mathbf{m}_0$ means that in the absence of noise the measured cylinder pieces are expected to be good estimates of the state; $\mathbf{A} = \mathbf{U}$, *i.e.* the prediction is actually a simple projection; $\mathbf{B} = \mathbf{0}$, $\mathbf{u} = \mathbf{0}$, *i.e.* there is no control input. As for \mathbf{P}_0 and the noise covariance matrices \mathbf{Q} , \mathbf{R} , they were determined experimentally: after an initial guess, we kept track of all the parameters in each run of the algorithm on real images, which allowed us to update the covariance matrices by learning. It can be noted that the high confidence in the initial measurement vector is justified by the application: the user is expected to select the starting point in a healthy part of the artery of interest where the cross-sectional shape is circular, the axial curvature is low and the contrasts are high.

The tracking strategy can be seen as a seeded region-growing technique adapted to the particular case of thin branching structures. It has two main advantages: its potential to manage bifurcations, and its potential to yield fast algorithms, since the processing is limited to a small part of the image volume. Unlike the minimum-cost path method also used to extract the vessel axis, the tracking algorithms do not necessarily need to be initialized by seed points located at the ends of each branch. In an ideal case, an entire vessel tree may be extracted starting from a single initial point. However, tracking also has the main drawback of the region-growing approach: the difficulty to define an efficient stopping criterion. Either the tracking stops on the first discontinuity (usually stenosis), or it "leaks" out beyond the ends of the vessels. The most difficult to cope with are the severe stenoses located at bifurcations, which unluckily is a frequent location of such pathologies. This difficulty can be overcome in the minimum-cost path approach.

4.2. Fast marching [36]

As mentioned in section 3.4.2, the level-sets formalism transforms a 2D contour detection problem into a 3D surface evolution problem. This dramatically increases the computational burden, which is the cost of the advantages of this approach. Let $F(\mathbf{x})$ be the local speed of the hypersurface in the normal direction. $F(\mathbf{x})$ is a function that may depend both on data and on surface geometry. Based on Hamilton-Jacobi formulation, the evolution equation for the moving hypersurface can be written as follows:

$$\psi_t + F|\nabla\psi| = 0. \quad (16)$$

Direct implementation of eq.(16) is computationally expensive, but general: it includes the case where some parts of the contour have positive speeds, while others have negative speeds. The active contour can thus move back and forth, until convergence. The fast-marching algorithm uses a monotonic restriction of $F(\mathbf{x})$ and the Eikonal equation:

$$|\nabla T(\mathbf{x})|F(\mathbf{x}) = 1, \quad (17)$$

where the function $T(\mathbf{x})$ represents the arrival time, *i.e.* time when the surface crosses a given point \mathbf{x} in space. It was demonstrated that one possible implementation of the

solution of this equation can be based on the Dijkstra's minimum-cost path search. We have chosen this version. The advantage of the fast-marching approach is a considerable speed-up. However, with a monotonic speed function, the contours may only grow (or only shrink if the speed is negative), which leads to difficulties mentioned in section 2.4.2. To some extent, this behavior can be compared with our implementation of snake with balloon force. In particular, we carried out some experiments with the fast-marching method directly applied to extract the 3D surface of a vessel, starting from one initial point within the lumen. Except in very well contrasted vessels, the surface tends to pass beyond the boundaries in the vicinity of the initial point before it reaches the boundary in the distant parts of the vessel. Furthermore, sub-pixel accuracy cannot be achieved with the Dijkstra's implementation, since it is based on a graph search using image pixels as nodes, and the points of the final contour are placed in the centers of the pixels.

5. Conclusion

Our work in the field of 3D vascular image processing has been mainly motivated by the wish of helping the clinicians in their task of making therapeutic decisions. After an overview of the methods developed in the last decade, in order to deal with various applications, we have highlighted the models and algorithms they make use of (see summary in table 1) and we provided their critical analysis.

Over the years, we have progressively enlarged the scope of the problems tackled. When maximum intensity projection (MIP) display was almost the only alternative of the tedious exploration of the native slices, we dealt with vessel enhancement as a pre-processing step before MIP computation, in order to improve the visualization of small vessels. With rapid advances in acquisition protocols, 3D images became more than a means to visualize the vasculature: the clinicians began to make measurements of diameters and lengths. In order to extract precise and reproducible quantitative morphological information, segmentation tools are required. Consequently, our work has then been focused on the accurate generation of the vessel axis, followed by boundary-extraction in image planes orthogonal to the axis. To overcome the limitations of this approach, owing to independent extraction of 2D cross-sectional contours, we implemented a 3D cylindrical deformable model. It was devised to construct patient-specific geometrical representation of vascular segments and to deduce from them the measurements useful in pre-operative stent-graft planning. Despite its numerous advantages in geometrical modeling, this method showed strong limitations in segmentation. Whereas the 3D deformable model is a discrete surface, we also developed a method directly providing a continuous representation of the vessel surface based on a parametric model of a generalized cylindrical surface. The latter model provides closed-form expressions to calculate several measurements, both conventional (such as the degree of stenosis) and volumetric that may be better correlated with the hemodynamic impact of the pathologies. How-

Tab. 1. Summary of the models and algorithms used in our methods

APPLICATION	MODALITY	MODELS	ALGORITHMS
Denoising and enhancement	subtracted MRA	morphological (sticks)	exhaustive search
Enhancement	MRA and CTA	second derivatives	multi-scale eigen-analysis of the Hessian
Acquisition optimization and stenosis quantification	subtracted MRA	inertia moments	tracking and multi-scale eigen-analysis of the inertia matrix
Stenosis quantification	CTA	intensity profiles	adaptive thresholding, tracking, ray casting
Quantification of lumen and wall	high resolution MRI	normalized snake	energy minimization (Lagrangian approach)
Quantification of lumen and wall	CTA	intensity profiles	ray casting
Stent-pose planning	MRA and CTA	simplex (DCS) and basic RGC	energy minimization (Lagrangian approach)
Volumetric quantification of lumen	CTA	RGC state model and geodesic contour	tracking, Kalman state estimator and fast-marching level-sets
Vascular tree extraction	MRA and CTA	elementary geometric shapes	tracking and adaptive thresholding or K-means clustering

ever, recent evidence indicates that the composition of the atherosclerotic plaque, rather than the degree of stenosis, determines the risk of acute vascular events. Therefore, one focus of our ongoing projects is the segmentation of both vascular lumen and wall, in attempt to quantitatively support the assessment of plaque vulnerability.

Despite the first results, the automated image-based assessment of the vessel wall composition remains an open question, which also requires further improvement of image acquisition techniques. Another open question is the segmentation of the vascular tree. It has many potential applications in visualization, detection of pathologies, model construction for the simulation of blood flow patterns to predict the outcome of surgical intervention, etc. Until now, most methods were devised to extract only one vascular segment at a time, while ignoring bifurcations. Our current objective, which we consider as the first step toward the vessel tree segmentation, is to achieve the extraction of the vascular tree skeleton. We pay a particular attention to the coronary arteries, which are of major interest in public health and represent a challenge in image processing.

We believe that the methods developed in near future will take into account the

dynamic character of the most recent acquisitions of thoracic CTA datasets, with the focus on cardiac motion. These datasets represent huge spatio-temporal volumes over several cardiac cycles. The recovery of the motion parameters is likely to improve both the reconstruction process (in order to provide a high quality display) and the extraction of clinically pertinent information via spatio-temporal segmentation. In the past, the evolution of the image processing methods followed the path of increasing abstraction level of the models used. If this logic is continued, we can expect the development of new more powerful mathematical models. In particular, the active surface models may probably be still improved. In the field of medical image analysis, the models should also include anatomic *a priori* knowledge via appropriate formal models.

Acknowledgements

This work was supported by the following projects: Colciencias #12040416468, ECOS-Nord #C03S02 and #C07M04, CIFI Uniandes #24 and PAI Polonium #05733RD. The authors are grateful to Alfred Anwander, Jacques Azencot, Michael Baltaxe Milwer, Christine Bresson, Juan Francisco Carrillo, Eduardo E. Dávila, Catherine Desbleds Mansard, Isabelle E. Magnin, Johan Montagnat, Piotr Orłowski and María Alejandra Zuluaga Valencia for their contributions in development, and to Loïc Boussel, Olivier Champin, Philippe C. Douek, Carolina Gómez, Albinka Maghiar, Jean-Michel Serfaty and Luis Felipe Uriza for their participation in the medical evaluation of the algorithms.

References

- 1987**
- [1] Kass M., Witkin A. and Terzopoulos D. : Snakes: Active contour models. *Int J Comput Vision*. 1(4), 321-331.
- 1991**
- [2] Cohen L.D. : On active contour models and balloons. *Comp Vision Graph Image Process: Image Underst*, 53(2), 211-218.
- 1992**
- [3] Staib L.H. and Duncan J.S. : Boundary finding with parametrically deformable models. *IEEE Pattern Anal Mach Intell*, 14, 1061-1075.
- 1995**
- [4] Du Y.P., Parker D.L., Davis W.L. : Vessel enhancement filtering in three-dimensional MR angiography. *J Magn Reson Imaging*, 5, 151-157.
- 1996**
- [5] Verdonck B., Bloch I. and Maître H. : Accurate segmentation of blood vessels from 3D medical images. In: *Int Conf Image Proc*, Lausanne, Switzerland, 3, 311-314.
- 1997**
- [6] Orkisz M., Bresson C., Magnin I.E., Champin O., Douek P.C. : Improved vessel visualization in MR Angiography by non-linear anisotropic filtering. *Magn Reson Med*, 37, 914-919.

- [7] Swift R.D., Ramaswamy K. and Higgins W.E.: Adaptive axes generation algorithm for 3D tubular structures. In: *Int Conf Image Proc, Sta Barbara, CA, USA, 2*, 136-139.

1998

- [8] Sato Y., Nakajima S., Shiraga N., Atsumi H., Yoshida S., Koller T., Gerig G. and Kikinis R.: Three-dimensional multi-scale line filter for segmentation and visualization of curvilinear structures in medical images. *Med Image Anal*, 2(2), 143-168.

1999

- [9] Frangi A.F., Niessen W.J., Hoogeveen R.M., Walsum T. and Viergever M.A.: Model-based quantitation of 3-D magnetic resonance angiographic images. *IEEE Trans. Med. Imaging*, 18(10), 946-956.
- [10] Sethian J.A.: *Level set methods and fast marching methods. Evolving interfaces in computational geometry, fluid mechanics, computer vision, and materials science*, Cambridge University Press, Cambridge.
- [11] Delingette, H.: General Object Reconstruction based on Simplex Meshes. *Int J Comput Vision*, 32(2), 111-146.

2000

- [12] Orkisz M., Hernández Hoyos M., Douek P.C., Magnin I.E.: Advances of blood vessel morphology analysis in 3D magnetic resonance images. *Mach Graph Vision*, 9(1/2), 463-471.
- [13] Wink O., Niessen W.J. and Viergever M.A.: Fast delineation and visualization of vessel in 3D angiography images. *IEEE Trans. Med. Imaging*, 19(4), 337-346.

2001

- [14] Orkisz M., Hernández Hoyos M.: Models for 3D vascular image analysis. *J Med Informatics Technol*. 2(1), IP13-IP22.
- [15] Antiga L., Ene-Iordache B., Remuzzi G., Remuzzi A.: Automatic generation of glomerular capillary topological organization. *Microvascular Research* 62, 346-354.
- [16] Flasque N., Desvignes M., Constans J.M., Revenu M.: Acquisition, segmentation and tracking of the cerebral vascular tree on 3D magnetic resonance angiography images. *Med Image Anal* 5, 173-183.
- [17] Toumoulin C., Boldak C., Dillenseger J.L., Coatrieux J.L., Rolland Y.: Fast detection and characterization of vessels in very large data sets using geometrical moments. *IEEE Trans Biomed Eng* 48, 604-606.

2002

- [18] Suri J.S., Liu K., Reden L., Laxminarayan S.: A review on MR vascular image processing: skeleton versus nonskeleton approaches: part II. *IEEE Trans Information Technol Biomed* 6, 338-350.
- [19] Hernández Hoyos M., Orkisz M., Puech P., Mansard-Desbleds C., Douek P.C., Magnin I.E.: Computer assisted analysis of 3D MRA images. *RadioGraphics*, 22(2), 421-436.
- [20] Wink O., Frangi A.F., Verdonck B., Viergever M.A., Niessen W.J.: 3D MRA coronary axis determination using a minimum cost path approach. *Magn Reson Med*, 47, 1169-1175.

2003

- [21] Azencot J., Orkisz M.: Deterministic and stochastic state model of right generalized cylinder (RGC-sm): application in computer phantoms synthesis, *Graph Mod*, 65, 323-350.
- [22] Boldak C., Rolland Y., Toumoulin C., Coatrieux J.L.: An improved model-based vessel tracking algorithm with application to Computed Tomography Angiography. *J Biocybern Biomed Eng*, 3(1), 41-64.

2004

- [23] Kirbas C., Quek F.K.H.: A review of vessel extraction techniques and algorithms. *ACM Comput Surv* 36, 81-121.

- [24] Desbleds Mansard C., Canet Soulas E.P., Anwander A., Chaabane L., Neyran B., Serfaty J.-M., Magnin I.E., Douek P.C., Orkisz M.: Quantification of multi-contrast vascular MR Images with the NLSnake, an active contour model: in vitro validation and in vivo evaluation, *Magn Reson Med*, 51(2), 370-379.
- [25] Flórez Valencia L., Montagnat J., Orkisz M.: 3D graphical models for vascular-stent pose simulation, *Mach Graph Vision*, 13(3), 235-248.
- [26] Flórez Valencia L., Vincent F., Orkisz M.: Fast 3D pre-segmentation of arteries in computed tomography angiograms. In: *Int Conf Comput Vision Graph*, Warsaw, Poland, Springer Verlag, 361-366.
- 2005**
- [27] Carrillo J.F., Orkisz M., Hernández Hoyos M.: Extraction of 3D vascular tree skeletons based on the analysis of connected components evolution. In: *CAIP'2005 - 11th Int IAPR Conf Computer Analysis of Images and Patterns*, Versailles, France, Springer Verlag LNCS 3691, 604-611.
- [28] Hernández Hoyos M., Orkisz M., Douek P.C., Magnin I.E.: Assessment of carotid artery stenoses in 3D contrast-enhanced magnetic resonance angiography, based on improved generation of the centerline. *Mach Graphics Vision* 14(4), 349-378.
- [29] Hernández Hoyos M., Orłowski P., Piątkowska-Janko E., Bogorodzki P., Orkisz M.: Vascular centerline extraction in 3D MR angiograms for phase contrast MRI blood flow measurement, *Int J Comp Assisted Radiol Surg*, 1(1), 51-61 (DOI: 10.1007/s11548-006-0005-0).
- [30] Lorenz C., von Berg J.: Fast automated object detection by recursive casting of search rays". In: *Comp Assisted Radiol Surg*, Berlin, Elsevier-Verlag, 230-235.
- 2006**
- [31] Lell M.M., Anders K., Uder M., Klotz E., Ditt H., Vega-Higuera F., Boskamp T., Bautz W.A., Tomandl B.F.: New techniques in CT angiography. *Radiographics*, 26, S45-S62.
- [32] Hernández Hoyos M., Serfaty J.M., Maghiar A., Mansard C., Orkisz M., Magnin I.E., Douek P.C.: Evaluation of semi-automatic arterial stenosis quantification. *Int J Comp Assisted Radiol Surg*, 1(3), 167-175 (DOI 10.1007/s11548-006-0049-1).
- [33] Flórez Valencia L., Azencot J., Vincent F., Orkisz M., Magnin I.E.: Segmentation and quantification of blood vessels in 3d images using a right generalized cylinder state model, *Int Conf Image Process*, Atlanta, GA, USA, 2441-2444.
- 2007**
- [34] Carrillo J.F., Hernández Hoyos M., Dávila E.E., Orkisz M.: Recursive tracking of vascular tree axes in 3D medical images, *Int J Comp Assisted Radiol Surg*, 1(6), 331-339 (DOI 10.1007/s11548-007-0068-6).
- [35] Zuluaga Valencia M.A., Dávila E.E., Uriza L.F., Hernández Hoyos M.: Carotid artery segmentation and characterization in 3D Computed Tomography (CT) images. In: *Comp Assisted Radiol Surg*, Berlin, Elsevier-Verlag, accepted.
- [36] Flórez Valencia L., Baltaxe Milwer M., Hernández Hoyos M., Vincent F., Douek P.C., Magnin I.E., Orkisz M.: Fast marching level-sets for the extraction of vascular cross-sectional contours in CT angiography images. In: *Comp Assisted Radiol Surg*, Berlin, Elsevier-Verlag, accepted.
- [37] Flórez Valencia L., Montagnat J., Orkisz M.: 3D models for vascular lumen segmentation in MRA images and for artery-stenting simulation, *ITBM-RBM Innov Technol Biol Med*, accepted (DOI 10.1016/j.rbmret.2007.04.001).
- [38] Baltaxe Milwer M., Flórez Valencia L., Hernández Hoyos M., Magnin I.E., Orkisz M.: Fast-marching contours for the segmentation of vessel lumen in CTA cross-sections. In: *EMBC - 29th Annual Int Conf IEEE Eng Med Biol Soc*, Lyon, France, accepted.