

Globally Optimal Active Contours, Sequential Monte Carlo and On-Line Learning for Vessel Segmentation

Charles Florin¹, Nikos Paragios², and Jim Williams¹

¹ Imaging & Visualization Department,
Siemens Corporate Research, Princeton, NJ, USA
755 College Road East, Princeton, NJ 08540, USA

² MAS - Ecole Centrale de Paris, Grande Voie des Vignes,
F-92 295 Chatenay-Malabry Cedex, France

Abstract. In this paper we propose a Particle Filter-based propagation approach for the segmentation of vascular structures in 3D volumes. Because of pathologies and inhomogeneities, many deterministic methods fail to segment certain types of vessel. Statistical methods represent the solution using a probability density function (pdf). This pdf does not only indicate the best possible solution, but also valuable information about the solution's variance. Particle Filters are used to learn the variations of direction and appearance of the vessel as the segmentation goes. These variations are used in turn in the particle filters framework to control the perturbations introduced in the Sampling Importance Resampling step (SIR). For the segmentation itself, successive planes of the vessel are modeled as states of a Particle Filter. Such states consist of the orientation, position and appearance (in statistical terms) of the vessel. The shape of the vessel and subsequently the particles pdf are recovered using globally active contours, implemented using circular shortest paths by branch and bound [1] that guarantees the global optimal solution. Promising results on the segmentation of coronary arteries demonstrate the potential of the proposed approach.

1 Introduction

Segmentation of vascular structures is a problem that arises in numerous situations in medical imaging, in particular for cardiac applications. Coronary arteries are thin vessels responsible for feeding the heart muscle in blood, and their segmentation provides a valuable tool for clinicians to diagnose diseases such as calcifications, and stenosis. Because of the low contrast conditions, and the coronaries vicinity to the blood pool, segmentation is a difficult task.

Since Computer Tomography (CT) and Magnetic Resonance (MR) imaging of the heart are now widely available, the number of patients imaged has significantly increased these past few years. Clinicians are now interested in periodically getting new images from the same patients to measure the development and severity of vascular diseases and their effects on the heart function. Such information is used to optimize the time of surgical operation and the effectiveness of treatments.

Vessel segmentation techniques consist of model-free and model-based methods. Skeleton-based techniques are the most primitive among the model-free [29] and aim

at detecting the vessel skeletons, from which the whole vessel tree is reconstructed. Vessel enhancement using a multiscale-structural term derived from the image intensity Hessian matrix [27, 12] and differential geometry-driven methods [19] refer to a different class of model-free approaches that characterize tubular structures using the ratios between the Hessian matrix eigenvalues. Voxels that best fit the characterization are rendered brighter than the others, and the resulting image enhance tubular structures.

In [3], an anisotropic filtering technique, called *Vesselness Enhancement Diffusion*, is introduced that can be used to filter noisy images preserving vessels boundaries. The diffusivity function relies on the *vesselness* function introduced in [12] to filter along the vessel principal direction and not across. In the resulting image, the background is smoothed, whereas the vessel remains unchanged. The flux maximization criterion, a step forwards, was introduced in [31] and was exploited for vessel segmentation in [6] in low contrast conditions using vessel measures introduced in [12].

Region growing methods [33] progressively segment the vessels from a seed point, based on intensity similarity between adjacent pixels. These methods work fine for homogeneous regions, but not for pathological vessels, and may leak into other structures of similar intensity. Morphological operators [11] can be applied to correct a segmentation, smooth its edges or eventually fill holes in the structure of interest, but fail to account for prior knowledge. Tracking approaches [17, 30] are based on the application of local operators to track the vessel. Given a starting condition, such methods recover the vessel centerline through processing information on the vessel cross section [16]. Various forms of edge-driven techniques, similarity/matching terms between the vessel profile in successive planes, as well as their combination, were considered to perform tracking.

On the other hand, model-based techniques use prior knowledge and features to match a model with the input image and extract the vessels. The knowledge may concern the whole structure, or consist in modeling locally the vessel. Vessels template matching techniques (*Deformable Template Matcher*) [25] have been investigated. The structure model consists of a series of connected nodes that is deformed to best match the input image. Generalized Cylindrical models are modified in Extruded Generalized Cylinders in [23] to recover vessels in angiograms. For curvy vessels, the local basis used for classical generalized cylinders may be twisted, and a non-orthogonality issue may occur. This problem is solved keeping the vessel cross section orthogonal to the centerline, and the two normal vectors always on the same side of the tangent vector spine, as the algorithm moves along the vessel.

Nevertheless, since vessels vary enormously from one patient to another, deformable models are preferred to template models. Deformable models can either be parametric or geometric. Parametric deformable models [26] can be viewed as elastic surfaces (often called *snakes*), and cannot handle topological changes. Geometric deformable models [4, 28], on the contrary, can change their topology during the process and therefore are well suited to vessel segmentation. Like snakes, deformable models aim at minimizing the energy computed along the model. Level sets [24] are a way to apply deformable model to non-linear problems, such as vessel segmentation [21]. One can refer to the fast marching algorithm and its variant for vessel segmentation using the minimal path principle [2, 5] to determine the path of minimal length between two points, backtracking from one point toward the other crossing the isosurfaces

perpendicularly. To discourage leaking, a local shape term that constrains the diameter of the vessel was proposed in [22]. One should also mention the method introduced in [20], where the optimization of a co-dimension two active contour was presented to segment brain vessels.

One can claim that existing approaches suffer from certain limitations. Local operators, region growing techniques, morphological filters as well as geometric contours might be very sensitive to local minima and fail to take into account prior knowledge on the form of the vessel. Parallel to that, cylindrical models, parametric active contours and template matching techniques may not be well suited to account for the non-linearity of the vessel structure, and require particular handling of branchings and bifurcations. Tracking methods can often fail in the presence of missing and corrupted data, or sudden changes. Level sets are very computational time-consuming and the Fast Marching algorithm loses all the local implicit function properties.

To improve segmentation results, a new method must account for non-linearities coming from branchings, pathologies, and acquisition artifacts, such as motion blur or CT beam hardening. This excludes any type of parametric models, or linear models, which would require special handling for bifurcations and non-linearities. Furthermore, the low contrast condition that features the coronaries drove the authors toward a method that would handle multiple hypotheses, and keep only the few most probable. The segmentation result would not be a deterministic result, but rather the most probable state of a vessel among several suppositions. Last, but not least, medical imaging is a field with vast prior knowledge; therefore, the new method must account for prior knowledge - if available -.

In this paper, we propose a particle-based approach to vessel segmentation where we re-formulate the problem of recovering successive planes of the vessel in a probabilistic fashion with numerous possible states. To this end, given an initial state for the vessel position, several hypotheses are generated uniformly in the feature space, and evaluated according to the observed data. From these hypothesis, a probability density function (pdf) can be defined, and used as a prior for a more efficient distribution of the hypothesis. Such an approach:

- combines edge-driven and region-based tracking metrics,
- recovers at each plane the optimal segmentation solution, that is the global minimum of the designed cost function,
- accounts for the structural and appearance non-linearity of the vessel,
- addresses pathological cases, and can incorporate prior local knowledge on the vessel structure.

The final paradigm consists of a fast multiple hypothesis propagation technique where the vessel structure as well as its appearance are successfully recovered. Such a framework allows to naturally address the non-linearity of the geometry and the appearance of coronaries and is compared in a favorable fashion with the existing approaches. The remainder of this paper is organized as follows. In section 2, we motivate vessel segmentation, introduce the feature space, and describe the measure used to quantify the quality of a given hypothesis. Random sampling and Particle Filters for tracking are introduced in section 3 while section 4 presents the overall system actually used to track vessels. Experimental results and discussion are part of the last section.

2 Vessel Segmentation

Cardio-vascular diseases are the leading cause of deaths in the USA (39%) and therefore there is a constant demand for improvement of diagnostic tools to detect and measure anomalies in the coronary tree. Such tools aid early diagnosis of the problem and therefore prevention that can significantly decrease the mortality rate due to cardiac diseases. One can consider the problem of vessel segmentation as a tracking problem of tubular structures in 3D volumes. Thus, given a starting position, the objective is to consider a feature vector that, upon its successful propagation, provides a complete segmentation of the coronaries. The statistical interpretation of such an objective refers to the introduction of a probability density function (pdf) that uses previous states to predict possible new positions of the vessel and image features to evaluate the new position. To this end, we define

- a state/feature vector, that defines the local geometry of a coronary artery
- an iterative process to update the density function, to predict the next state
- a distance between prediction and actual observation, to measure the quality of a given feature vector with respect to the image data.

2.1 The State/Feature Vector

One can define the state of the vessel at a given time as follows:

$$\mathbf{x} = \underbrace{(x_1, x_2, x_3)}_{\text{position}}, \underbrace{\Theta = (\theta_1, \theta_2, \theta_3)}_{\text{orientation}}, \underbrace{\mathbf{p}_{vessel}}_{\text{appearance}}$$

where the vessel state vector consists of the 3D location of the vessel \mathbf{x} , the tangent vector Θ , and the parameters required for the pdf estimation of the appearance of the vessel \mathbf{p}_{vessel} , as a mixture of two gaussians:

$$\mathbf{p}_{vessel} = ((P_B, \mu_B, \sigma_B), (P_C, \mu_C, \sigma_C)) \tag{1}$$

It is reasonable to assume irregularity in the appearance of the vessel because of the presence of calcifications, stents, stenosis and diseased vessel lumen [FIG. (1)]. Therefore simple parametric statistical models on the appearance space will fail to account for the statistical properties of the vessel and more complex distributions are to be considered. We consider a Gaussian mixture model that consists of two components to represent the evolving distribution of the vessel, the contrast enhanced blood (P_B, μ_B, σ_B) and the high density components, such as calcifications or stent, (P_C, μ_C, σ_C) subject to the constraint $[P_C + P_B = 1]$ leading to the following state vector:

$$\omega = (\mathbf{x}, \Theta, (P_B, \mu_B, \sigma_B), (P_C, \mu_C, \sigma_C)) \tag{2}$$

Such a state vector is to be recored for subsequent planes leading to complete reconstruction of the vessel tree. However, neither the planes position and orientation, nor the actual position of the vessel within this plane is known. In order to recover the most prominent plane position, a constrained multiple hypotheses framework will be used




Fig. 1. (a) calcification, (b) stent (high intensity prosthesis), (c) branching with obtuse angles, (d) stenosis (sudden reduction of vessel cross section diameter)

according to a particle filter implementation. Such a framework will be explained at a later section.

Let us assume for the moment that the plane position is known as well as its orientation. Vessel segmentation consists of recovering the area of image within this plane that corresponds to the vessel. Snakes [18] as well as their geometric alternatives have been popular techniques to address such a demand. Despite numerous improvements, such methods often converge to local minimum. Such a limitation was addressed in [1] - known as circular shortest path algorithm by branch and bound - once appropriate initial conditions have been given to the process, that in our case could be satisfied.

2.2 Circular Shortest Paths and 2D Vessel Segmentation

The *Circular Shortest Paths by Branch and Bound* (CSP) [1] is a binary search-tree technique to recover the globally optimal active contour, given a point inside the contour and a potential map. First of all, let us note that the problem of finding the globally optimal active contour is equivalent to computing the minimal *weight* path (given a Riemannian metric) that connects a point at angle $\theta = 0$ to its equivalent at $\theta = 2\pi$ across the log-polar transform of the original image, see [FIG. (2)]. Given a Riemannian metric g (usually equal to the image gradient), the *weight* W of a path \mathbf{P} is defined as:

$$W(\mathbf{P}) = \int_{\mathbf{P}} g(\mathbf{P}(s)) ds. \quad (3)$$

Given a start point p_0 at $\theta = 0$, the end point $p_{2\pi}$ at $\theta = 2\pi$ of the minimal *noncircular* path is defined as

$$p = \operatorname{argmin}_{\mathbf{P}(2\pi)=p} W(\mathbf{P}). \quad (4)$$

This end point $p_{2\pi}$ is very quickly found using the well-known Dijkstra [7] algorithm, with the Riemannian metric g ([EQ. (3)]) playing the role of potential map. To demonstrate the use of a binary search-tree, a property needs to be stated at that point, whose proof is straightforward (see [1]):

for two subsets $S_1 \subseteq S_2$, the minimal path \mathbf{P}_2 of S_2 has a lower weight than the minimal path \mathbf{P}_1 of S_1 , otherwise stated as $W(\mathbf{P}_2) \leq W(\mathbf{P}_1)$.

A corollary is:

for any point set S , the weight of the minimal path \mathbf{P} (circular or not) is a lower bound of the minimal *circular* path weight. Therefore, if $\{S_1, S_2\}$ is a partition of S , and $W(\mathbf{P}_1) \leq W(\mathbf{P}_2)$, the minimal *circular* path of S has its starting point p_0 (and obviously ending point $p_{2\pi}$ as well) in the subset S_1 .




Fig. 2. Discreet grid of a log-polar image, with a circular path (dashed, subset S_1) and a noncircular path (plain, subset S_2)

Consequently, a binary search-tree is used in the CSP algorithm. First, any set of initial points $S = \{p_0\}$ is divided into two subsets $S_1 = \{p_0\}_1$ and $S_2 = \{p_0\}_2$; second, the minimal *noncircular* paths P_1 and P_2 are computed using [EQ. (4)] for the two subsets. This procedure is then repeated with the subset of minimal path until the subsets are reduced to a single point. At the bottom of the binary search-tree, the subsets are reduced to singletons, and their minimal path are naturally *circular*. The Globally Optimal Circular Shortest Path is obtained that way. The low cost complexity (for width u and height v , $O(u^{1.6}v)$ average time, or less than a milisecond for 15x15 pixels cross section profile, see [FIG. (4)]) makes this method very attractive for repetitive testings, such as the particle filters presented in [SEC. (3)]. It also reduces the dimensionality of the feature space, compared to model-based methods (elliptic models, tubular models,...).

The CSP algorithm is an efficient technique to image segmentation for closest structures under the assumption that a point is given in the structure interior. Since segmentation in our case is approached as a multiple hypotheses testing, one can assume that each hypotheses generation could provide a start point to the CSP that is a necessity for the construction of the log-polar image. The multiple hypotheses generation could be done in a number of fashions. Sequential Monte Carlo is the prominent technique that associates evolving densities to the different hypotheses, and maintains a number of them. Particle filters is the most prominent technique to implemented such a strategy.

3 Particle Filters

3.1 Particle Filters: Generalities

Particle Filters [8, 15] are a sequential Monte-Carlo technique that is used to estimate the Bayesian posterior probability density function (pdf) with a set of samples [13, 32]. In terms of a mathematical formulation, such a method approximates the posterior pdf by M random measures $\{x_t^m, m = 1..M\}$ associated to M weights $\{w_t^m, m = 1..M\}$, such that

$$p(x_t|z_{1:t}) \approx \sum_{m=1}^M w_t^m \delta(x_t - x_t^m). \tag{5}$$

where each weight w_t^m reflects the importance of the sample x_t^m in the pdf, given the observations sequence $z_{1:t}$, as shown in [FIG. (3)]. Using Bayes rule, one can sequentially estimate $p(x_t|z_{1:t})$ from $p(x_{t-1}|z_{1:t-1})$, knowing $p(x_t|x_{t-1})$ and measuring $p(z_t|x_t)$:




Fig. 3. The resampling process: a random selection chooses the samples with the highest weights where a local perturbation is applied

$$\begin{aligned}
 p(x_t|z_{1:t}) &\propto p(z_t|x_t)p(x_t|z_{1:t-1}) \\
 &\propto p(z_t|x_t) \int p(x_t|x_{t-1})p(x_{t-1}|z_{1:t-1})dx_{t-1}
 \end{aligned}$$

$p(z_t|x_t)$ is discussed in [SEC. (4.1)], while a novel method to locally estimate $p(x_t|x_{t-1})$ is presented in [SEC. (3.3)]. The samples x_t^m are drawn using the principle of *Importance Density* [14], of pdf $q(x_t|x_{1:t}^m, z_t)$, and it is shown that their weights w_t^m are updated according to

$$w_t^m \propto w_{t-1}^m \frac{p(z_t|x_t^m)p(x_t^m|x_{t-1}^m)}{q(x_t^m|x_{t-1}^m, z_t)}. \tag{6}$$

Once a set of samples has been drawn, $p(x_t^m|x_{t-1}^m, z_t)$ can be computed out of the observation z_t for each sample, and the estimation of the posteriori pdf can be sequentially updated. Such a process will remove most of the particles and only the ones that express the data will present significant weights. Consequently the model will lose its ability to track significant changes on the pdf; therefore a resampling procedure has to be executed on a regular basis. Such a process will preserve as many samples as possible with respectful weights. One can find in the literature several resampling techniques. We chose the most prominent one, Sampling Importance Resampling, for its simplicity to implement, and because it allows more hypothesis with low probability to survive, compared to more selective techniques such as Stratified Resampling [10].

3.2 Sampling Importance Resampling

The Sampling Importance Resampling (SIR) algorithm [13] consists of choosing the prior density $p(x_t|x_{t-1})$ as importance density $q(x_t|x_{1:t}^m, z_t)$. This leads to the following condition, from [EQ. (6)]

$$w_t^m \propto w_{t-1}^m p(z_t|x_t^m). \tag{7}$$

The samples are updated by setting $x_t^m \propto p(x_t|x_{t-1}^m)$, and perturbed according to a random noise vector ϵ , so that $x_t^m \propto p(x_t|x_{t-1}^m)$. The SIR algorithm is the most widely

used resampling method because of its simplicity from the implementation point of view. Nevertheless, the SIR uses mostly the prior knowledge $p(x_t|x_{t-1})$, and does not take into account the most recent observations z_t . Such a strategy could lead to an overestimation of outliers. On the other hand, because SIR resampling is performed at each step, fewer samples are required, and thus the computational cost may be reduced with respect to other resampling algorithms. Finally, in practice, the estimation of ϵ 's law is difficult, and prior knowledge is usually required. A novel method is proposed in the following section [SEC. (3.3)] to circumvent this issue, by locally estimating $p(x_t|x_{t-1}^m)$.

3.3 Reinforced SIR: The State Transition Noise Pdf

After a particle x_{t-1}^m has been selected by the SIR algorithm, a random noise vector ϵ is added (see previous section [SEC. (3.2)]). A straightforward solution consists in using prior knowledge to estimate the law of ϵ once for all. This method presents two difficulties: first, prior knowledge may be limited and/or hard to obtain, second, vessels are linear structures only very locally, therefore the law of ϵ may greatly vary from one patient to another. In the technique presented in this paper, the distribution of ϵ is updated at every time step. At a given time step, each particle x_{t-1}^m selected by the SIR generates N offsprings by adding a random noise vector, uniformly distributed, and moving it forward (in the direction of the vessel, given by the particle hypothesis). Once their probability is estimated, these N offsprings particles provide a pdf ($p(x_t|x_{t-1}^m)$) which is then used for the distribution of the random vector ϵ .

The final paradigm for resampling follows the procedure:

1. first, particles are selected randomly according to their probability, as in any SIR procedure
2. second, the selected particles generates N new offsprings uniformly distributed
3. these offsprings probabilities are estimated, and a pdf is then drawn for each SIR selected particle
4. finally, this pdf is used to generate a random noise vector ϵ that perturbs the SIR selected particles

In other words, once the SIR selected a particle x_{t-1}^m to be resampled, $p(x_t|x_{t-1}^m, z_t)$ is estimated in a way similar to [EQ. (5)]:

$$p(x_t|x_{t-1}^m, z_t) \approx \sum_{i=1}^N w_t^i \delta(x_t - x_t^i), \tag{8}$$

where the x_t^i are generated from $x_{t-1}^m + \epsilon_i$, with the ϵ_i uniformly distributed. The weights w_t^i are estimated from the observation z_t .

This method presents two main advantages. First, as the noise vector ϵ is random, the advantages of SIR over exhaustive search are preserved. Second, the distribution of ϵ is updated at every time step, and for every particle, avoiding the disadvantages of having a noise distribution that would be determined once for all from prior knowledge. Vessels can be straight and suddenly become tortuous, or can have a very homogeneous shape/appearance before encountering a very inhomogeneous region. This *Reinforced SIR* captures the conditions change and adapts the noise vector distribution.

4 Particle Filters and Vessel Tracking

We now consider the application of such non linear model to vessel segmentation and tracking. Without loss of generality one can assume that the root of a coronary is known, either provided by the user or through some automatic procedure. Simple segmentation of that area can provide an initial guess on the statistical properties of the vessel

$$\mathbf{p}_{vessel} = ((P_B, \mu_B, \sigma_B), (P_C, \mu_C, \sigma_C)) \quad (9)$$

using an expectation/maximization process. Then, one can consider the problem of vessel segmentation equivalent to the recovery of successive cross-sections, along with the position of the vessel at any given cross-section. Such an approach is equivalent to finding a deterministic number of sequential states $\omega_\tau = (\mathbf{x}_\tau, \Theta_\tau, \mathbf{p}_{vessel})$, which belong to the feature space (see [SEC. (2.1)]) where we use the notion of Particle Filters.

The multiple hypotheses nature of the method requires a metric definition to validate their correctness. Given, the current state and the perturbation law we produce a number of new states following this law. Such states refer to a new plane, as well as a center point for the elliptic structure and therefore the CSP algorithm can be used to provide the most prominent area for the vessel given these initial conditions. We use this area and two metrics that aim to account for the shape and appearance of the vessel toward validation of the considered hypotheses.

4.1 Prediction and Observation: Distance

To this end, we are using mostly the image terms, and in particular the intensities that do correspond to the vessel in the current cross-section. The observed distribution of this set is approximated using a Gaussian mixture model according to the expectancy-maximization principle. Each hypothesis is composed by the features given in [EQ. (2)], therefore, the probability measure is essentially the likelihood of the observation z , given the appearance A model. The following measures (abusively called probabilities) are normalized so that their sum over all particles is equal to one.

- Probability measure for shape

Once the vessel's edge is detected using *Circular Shortest Path* [SEC. (2.2)], a measure of contrast, called the *ribbon measure*, R is computed:

$$\begin{cases} R = -\infty & , \quad \mu_{int} \leq \mu_{ext} \\ R = \frac{\mu_{int} - \mu_{ext}}{\mu_{int} + \mu_{ext}} & , \quad otherwise \end{cases}$$

while the correctness of the prediction is given by:

$$p(z|S) = e^{-\frac{|R|}{R_0}}$$

where μ_{int} is the mean intensity value for the voxels in the vessel, and μ_{ext} is the intensities mean value for the voxels in a band outside the vessel, such that the band and the vessel's lumen have the same area.




Fig. 4. Three vessels cross sections detected using the ribbon measure

Since the coronary arteries are brighter than the background, the best match maximizes R (see [FIG. (4)]).

- Probability measure for appearance

For the vessel lumen pixels distribution \mathbf{p}_{vessel} [EQ. (1)], the probability is measured as the distance between the hypothesized distribution and the distribution actually observed.

The distance we use is the symmetrized Kullback-Leibler distance D between the model $p(x) = \mathbf{p}_{vessel}$ and the observation $q(x)$, obtained from the CSP segmentation:

$$D = \int p(x) \log\left(\frac{p(x)}{q(x)}\right) + q(x) \log\left(\frac{q(x)}{p(x)}\right) dx,$$

$$p(z|A) = e^{-\frac{|D|}{D_0}}.$$

The combination of edge-driven and region-based metrics measures the fitness of the observation to the prior knowledge included in the state vector.

4.2 Branching Detection

When a branching occurs, the particles split up in the two daughter branches, and then track them separately (see [FIG. (5)]). As branchings are never perfectly balanced, one of the branches attracts most of the particles after few resampling steps. To avoid the collapse of one of the modes, two techniques are available: either to increase the number of particles in the weakest branch, or to treat the two branches separately. The second approach is preferred in this paper. To this end, a simple K-means clustering on the joint space (position+orientation) of the particles is considered. When the two clusters are well separated, the number of particles is doubled and equally dispatched in the two branches. The segmentation goes on, according to [EQ. (6)], with a bi-modal distribution.

The K-means algorithm [9] partitions N points, x_n , into K disjoint clusters, of centers μ_j , minimizing the sum-of-squares

$$J = \sum_{j=0}^K \sum_{n=0}^N |x_n - \mu_j|^2.$$

The K-mean procedure alternates two steps: first each point is associated to the nearest center μ_j , then each center is moved in the barycenter of the cluster.




Fig. 5. (a) branching points between LCX and LAD for three patients with the particles' mean state overlaid, (b) the particles , clustered using K-means, follow up the two branches

Table 1. Results table showing the percentage of branches correctly segmented, over a dataset of 34 patients, using Particle Filters (PF) and Front Propagation (FP)

vessel name	RCA	Acute Marginal	LAD	First Septal	LCX	Obtuse Marginal
% of branches, using PF	100%	85.3%	100%	94%	100%	94%
% of branches, using FP	64%	18%	53%	32%	39%	22%

4.3 Implementation and Validation

Regarding the initial configuration, the use of approximately 1,000 particles gave sufficient results for our experiments. We perform a systematic resampling according to the Sampling Importance Resampling every time the effective sampling size $N_{eff} = \sum_i 1/w_i^2$ (where w_i is the weight of the i th particle) falls below half the number of particles. As mentioned in Section 3.1, the preference for SIR, compared to Stratified Resampling [10], is motivated by the robustness of the segmentation. The *reinforced SIR* strategy exposed in [SEC. (3.3)] gives better results, for a constant number of particles.

Validation is a difficult part for any coronary segmentation method. The algorithm has been evaluated on 34 patients, and has successfully recovered all the main arteries (RCA, LAD, LCX) for each patient as shown in the following table, while a small portion of visual results are also presented in [FIG. (6)].

The percentage in the above table corresponds to the number of branches segmented by Particle Filters and identified by a human expert. For comparison purposes, the same test is performed using Front Propagation based on the image Hessian matrix [27]. These results were achieved with a one-click initialization; a method based on a PCA on the intensity volume gives the approximative initial direction. All patients presented some kind of artery pathologies in one, at least, of their coronary vessels. This means the Particle Filter successfully segmented both healthy and unhealthy coronaries. The method seems to outperform regarding the detection of the main branchings, while in some cases smaller branchings at the lowest parts of the vessel tree, have been missed. Nevertheless, one can argue that their clinical use is of lower importance. However, current studies focus on the issue of branchings for narrow vessels in very low contrast conditions. The comparative study demonstrate the Particle Filters capability to outperform deterministic hessian based methods in cases with corrupt data (pathologies).




Fig. 6. Segmentation of the Left anterior descending coronary artery and Right coronary artery in CTA (in red) for four patients (1) & (2); (a) coronary tree, (b,c,d) Different 3D views superimposed to the cardiac volume are presented.

5 Discussion

In this paper, we have shown that Particle Filters can be used for vascular segmentation. In the context of vascular segmentation, Particle Filters sequentially estimate the pdf of segmentations in a particular feature space. The case of coronary arteries was considered to validate such an approach where the ability to handle discontinuities on the structural (branching) as well as appearance space (calcifications, pathological cases, etc.) was demonstrated. The main advantage of such methods is the non-linearity assumption on the evolution of samples. Experiments were conducted on several healthy and diseased patients CTA data sets, segmenting the *Left Main Coronary Artery* and the *Right Coronary Artery* [FIG. (6)].

Introducing further prior knowledge in the segmentation process is the most prominent future direction. One can see such a contribution in two parallel paths. First, building better models that account for the appearance of the vessel seems to be a necessity toward capturing the coronaries at the lowest parts of the vessel tree. The current model is based on the global statistics of the appearance of the vessel and one can claim is a meaningful measure for vessel cross sections with a certain area.

References

1. B. Appleton and C. Sun. Circular shortest paths by branch and bound. 36(11):2513–2520, November 2003.
2. B. Avants and J. Williams. An adaptive minimal path generation technique for vessel tracking in cta/ce-mra volume images. In *MICCAI*, pages 707–716, 2000.
3. C. Cañero and P. Radeva. Vesselness enhancement diffusion. *Pattern Recognition Letters*, 24(16):3141 – 3151, 2003.
4. V. Caselles, F. Catté, B. Coll, and F. Dibos. A geometric model for active contours in image processing. *Numerische Mathematik*, 66(1):1–31, 1993.

5. T. Deschamps and L. D. Cohen. Fast extraction of minimal paths in 3D images and applications to virtual endoscopy. *Medical Image Analysis*, 5(4):281–299, December 2001.
6. M. Descoteaux, L. Collins, and K. Siddiqi. Geometric Flows for Segmenting Vasculature in MRI: Theory and Validation. In *Medical Imaging Computing and Computer-Assisted Intervention*, pages 500–507, 2004.
7. E. W. Dijkstra. A note on two problems in connexion with graphs. *Numerische Mathematik*, 1:269–271, 1959.
8. A. Doucet, J. de Freitas, and N. Gordon. *Sequential Monte Carlo Methods in Practice*. Springer-Verlag, New York, 2001.
9. R Duda and P. Hart. *Pattern Classification and Scene Analysis*. John Wiley and Sons, 1973.
10. P. Fearnhead and P. Clifford. Online inference for well-log data. *Journal of the Royal Statistical Society*, 65:887–899, 2003.
11. M. Figueiredo and J. Leitao. A nonsmoothing approach to the estimation of vessel contours in angiograms. *IEEE Transactions on Medical Imaging*, 14:162–172, 1995.
12. A. Frangi, W. Niessen, P. Nederkoorn, O. Elgersma, and M. Viergever. Three-dimensional model-based stenosis quantification of the carotid arteries from contrast-enhanced MR angiography. In *Mathematical Methods in Biomedical Image Analysis*, pages 110–118, 2000.
13. N. Gordon. Novel Approach to Nonlinear/Non-Gaussian Bayesian State Estimation. *IEE Proceedings*, 140:107–113, 1993.
14. N. Gordon. On Sequential Monte Carlo Sampling Methods for Bayesian Filtering. *Statistics and Computing*, 10:197–208, 2000.
15. N. Gordon. A Tutorial on Particle Filters for On-line Non-linear/Non-Gaussian Bayesian Tracking. *IEEE Transactions on Signal Processing*, 50:174–188, 2002.
16. M. Hart and L. Holley. A method of Automated Coronary Artery Trackin in Unsubtracted Angiograms. *IEEE Computers in Cardiology*, pages 93–96, 1993.
17. M. Isard and A. Blake. Contour Tracking by Stochastic Propagation of Conditional Density. In *European Conference on Computer Vision*, volume I, pages 343–356, 1996.
18. M. Kass, A. Witkin, and D. Terzopoulos. Snakes: Active Contour Models. In *IEEE International Conference in Computer Vision*, pages 261–268, 1987.
19. K. Krissian, G. Malandain, N. Ayache, R. Vaillant, and Y. Troussset. Model based detection of tubular structures in 3d images. *Computer Vision and Image Understanding*, 80:130–171, 2000.
20. L. Lorigo, O. Faugeras, E. Grimson, R. Keriven, R. Kikinis, A. Nabavi, and C. Westin. Codimension-Two Geodesic Active Contours for the Segmentation of Tubular Structures. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 1:444–451, 2000.
21. R. Malladi and J. Sethian. A Real-Time Algorithm for Medical Shape Recovery. In *International Conference on Computer Vision*, pages 304–310, 1998.
22. D. Nain, A. Yezzi, and G. Turk. Vessel Segmentation Using a Shape Driven Flow. In *Medical Imaging Computing and Computer-Assisted Intervention*, 2004.
23. T. O’Donnell, T. Boulton, X. Fang, and A. Gupta. The Extruded Generalized Cylinder: A Deformable Model for Object Recovery. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 174–181, 1994.
24. S. Osher and N. Paragios. *Geometric Level Set Methods in Imaging, Vision and Graphics*. Springer Verlag, 2003.
25. R. Petrocchi, K. Manbeck, and J. Elion. Three Dimensional Structure Recognition in Digital Angiograms using Gauss-Markov Models. *IEEE Computers in Radiology*, pages 101–104, 1993.
26. D. Rueckert, P. Burger, S. Forbat, R. Mohiadin, and G. Yang. Automatic Tracking of the Aorta in Cardiovascular MR images using Deformable Models. *IEEE Transactions on Medical Imaging*, 16:581–590, 1997.

27. Y. Sato, S. Nakajima, H. Atsumi, T. Koller, G. Gerig, S. Yoshida, and R. Kikinis. 3D Multi-scale line filter for segmentation and visualization of curvilinear structures in medical images. In *Conference on Computer Vision, Virtual Reality and Robotics in Medicine and Medial Robotics and Computer-Assisted Surgery*, pages 213–222, 1997.
28. J. Sethian. A Review of the Theory, Algorithms, and Applications of Level Set Methods for Propagating Interfaces. *Cambridge University Press*, pages 487–499, 1995.
29. E. Sorantin, C. Halmaj, B. Erbohelyi, K. Palagyi, K. Nyul, K. Olle, B. Geiger, F. Lindbichler, G. Friedrich, and K. Kiesler. Spiral-CT-based assesment of Tracheal Stenoses using 3D Skeletonization. *IEEE Transactions on Medical Imaging*, 21:263–273, 2002.
30. K. Toyama and A. Blake. Probabilistic Tracking in a Metric Space. In *IEEE International Conference in Computer Vision*, pages 50–59, 2001.
31. A. Vasilevskiy and K. Siddiqi. Flux Maximizing Geometric Flows. In *IEEE International Conference in Computer Vision*, pages I: 149–154, 2001.
32. W. West. Modelling with mixtures. In J. Bernardo, J. Berger, A. Dawid, and A. Smith, editors, *Bayesian Statistics*. Clarendon Press, 1993.
33. P. Yim, P. Choyke, and R. Summers. Grayscale Skeletonization of Small Vessels in Magnetic Resonance Angiography. *IEEE Transactions on Medical Imaging*, 19:568–576, 2000.