Monte-carlo Sampling, Particle Filters and Segmentation of Coronaries

Charles Florin
Nikos Paragios
Jim Williams

Research Report 05-03
January 2005

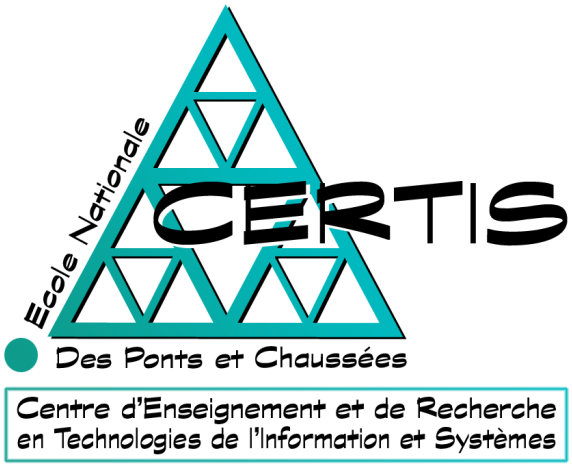

CERTIS, ENPC,
77455 Marne la Vallee, France,
http://www.enpc.fr/certis/

Monte-carlo Sampling, Particle Filters and Segmentation of Coronaries

Echantillonage de Monte Carlo, filtres particulaires et segmentation de coronaires

Charles Florin[1,2]
Nikos Paragios[2]
Jim Williams[1]

[1]Siemens Corp. Research, 755 College Road, Princeton 08540 NJ http://www.scr.siemens.com
[2]CERTIS, ENPC, 77455 Marne la Vallee, France, http://www.enpc.fr/certis/

Abstract

In this report, we propose a particle filter-based propagation approach for the segmentation of vascular structures in 3D volumes. In our approach, successive planes of the structure are modeled as states of a particle filter. Such a state consists in the 3D structure, orientation, position, geometric form and appearance in statistical means. In order to account for bifurcations and branchings, we consider a Monte Carlo sampling rule that propagates in parallel multiple hypotheses. In order to account for prior knowledge, notions of a linear (Kalman) filter are incorporated within the proposed approach. The prior knowledge constrains the vessel detection, which combines edge-driven and region-based metrics. Promising results on the segmentation of coronary arteries demonstrate the potential of the proposed approach.

Résumé

Dans ce rapport, les auteurs proposent une segmentation basee sur un filtre particulaire pour les structures vasculaires dans des volumes 3D. Suivant leur approche, les plans successifs de la structure sont modeles comme les etats d'un filtre particulaire. Un tel etat est constitue d'une structure 3D, une orientation, une position, une forme geometrique et un model d'apparence, en termes statistiques. Afin de prendre en compte les bifurcations et branchements potentiels, les auteurs considerent une regle d'echantillonage de Monte Carlo qui se propoage en suivant de multiples hypotheses en parrallele. Afin de tenir compte de connaissance a priori, un filtre lineaire (Kalman) est incorpore a la methode. Des resultats prometteurs sur la segmentation des coronaires demontrent le potentiel de l'approche presentee.

Contents

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, the segmentation is a difficult task.

Since Computer Tomography (CT) and Magnetic Resonance (MR) imaging is 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 heart diseases, there effects on the heart function, to optimize the time of surgical operation, and the effectiveness of treatments. All this results in a large amount of information to process. To automatize the process, one has to segment the coronary arteries from the rest of the data first. Many techniques have been developed recently for this task.

One may distinguish model-free segmentation techniques from model-based methods. Among the model-free techniques, skeleton-based techniques [31] aim at detecting skeletons, from which the whole vessel tree is reconstructed. Also, vessel enhancement using a multiscale-structural term derived from the image intensity Hessian matrix [29, 12], and differential geometry-based methods [19], have been widely used. They both consist in characterizing tubular structures using 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.

Region growing methods [35] consist in progressively segmenting 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, 32] 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.

The maximization of flux was introduced in [33] and was exploited for vessel segmentation in [6] in low contrast conditions using vessel measures introduced in [12]. The vectors normal to the vessel boundaries are collected using the Hessian matrix eigenvectors collected on the points satisfying the vessel measures. The geometric maximizing flux algorithm is then applied to recover the vessel boundaries.

In [2], Bouix, Siddiqi and Tannenbaum apply a method that relies on the average outward flux of the gradient vector field of the Euclidian distance from the vessel boundary to recover skeleton points.

In [1], Armande et al. introduce a multiscale method to segment thin nets (line where the gray level is locally extremum). First, the image is filtered by a gaussian at a certain scale. For each scale, the image maximum curvature is computed and based on differential properties, the points that belong to the vessels centerline are kept.

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*) [26] were 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 [24] 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 rigid models. Deformable models can either be parametric or geometric. Parametric deformable models [28] can be viewed as elastic surfaces (often called *snakes*), and cannot handle topological changes. Geometric deformable models [4, 30], on the contrary, preserve topology and are well fitted for vessels segmentation. Like snakes, deformable models aim at minimizing the energy computed along the model. Level sets [25] are a way to consider deformable model for non-linear problems, such as vessel segmentation [22]. The implicit function is defined all over the input image, and the zero-level set determines the deformable model. To discourage leaking, a local shape term that constrains the diameter of the vessel was proposed [23]. One should also mention the method introduced in [21], where the optimization of a co-dimension

two active contour was considered to segment brain vessels.

To account for the snakes sensitivity to initialization, [10] introduces snakes determined after a learning process based on a non-parametric estimator. The learning function uses a Parzen window estimator, with a Gaussian kernel. The Parzen window estimator relies on feature value observations, and compare these values with model values. The objective function is then chosen to maximize the probability distribution of these observable quantities.

The Fast Marching algorithm was introduced as an implementation of front propagation to recover isosurfaces for any given Riemannian metric. The Minimal Path algorithm [5] takes advantage of the Fast Marching algorithm to determine the path of minimal weight between two points, backtracking from one point toward the other crossing the isosurfaces perpendicularly. This method can be considered to track vessels centerline provided that two points on the centerline are know prior to the process.

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. On the other hand, 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 loose all the local implicit function properties.

To improve segmentation results, a new method must account for non-linearities coming from two origins: branchings, and pathologies. This excludes any type of parametric models, or linear models, which would require a particular 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. The supposition that would math both the prior knowledge and the information provided by the input image would receive the highest probability.

In this paper, we propose a particle-based approach to vessel segmentation. Our approach (i) combines edge-driven and region-based tracking metrics, (ii) accounts for the structural and appearance non-linearity of the vessel through the maintenance of multiple hypotheses, (iii) can address pathological cases, and (iv) can incorporate prior local knowledge on the vessel structure through constraints derived from a linear (Kalman) filter. The final paradigm consists of a fast mul-

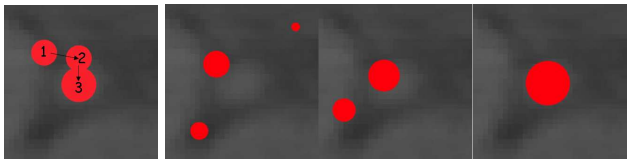

(a) (b)

Figure 1: (a) one single hypothesis is iteratively improved, (b) Collection of hypothesis, iteratively selected and improved

tiple hypothesis propagation technique where the vessel structure as well as its appearance are successfully recovered.

The remainder of this paper is organized as follows. In section 2, we motivate vessel segmentation and introduce the concept of the proposed approach and the features space. Random sampling and particle filters for tracking are introduced in section 3 while vessel segmentation with notions of local prior knowledge are presented in section 4. Experimental results and discussion are part of the last section.

2 Preliminaries

Cardio-vascular diseases are the leading cause of deaths in USA (39%) and therefore there is an imminent need for automated diagnostic tools to detect anomalies in the proper operation of coronaries. Such tools could lead to early diagnostics of the problem and therefore prevention that eventually will 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 can lead to the complete reconstruction 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, one should define

- a state/feature vector,
- an iterative process to update the density function,
- a distance between prediction and actual observation.

2.1 Particle Filters: basic concept

To explain the concept of Particle Filters, let us start with one single hypothesis about the vessel. This hypothesis describes the location, orientation, shape and

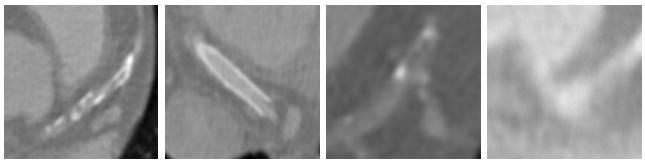

(a) (b) (c) (d)

Figure 2: (a) calcification, (b) stent (high intensity prosthesis), (c) branching with obtuse angles, (d) stenosis (sudden reduction of vessel cross section diameter).

appearance of the vessel, and a probability function can measure the likelihood of this given hypothesis. If any parameter is modified in this hypothesis, another probability is measured and the best of these two hypothesis can be kept to describe the vessel. The best hypothesis can thus be iteratively modified to improve the model (see [FIG. (1-a)]).

Now, instead of one single hypothesis, a collection of hypothesis can be used. Since a probability can be associated to each hypothesis, a probability density function can be drawn over the feature space. The less probable hypothesis can then be iteratively replaced by the most probable, slightly modified (see [FIG. (1-b)]).

A Particle Filter work essentially the same: each hypothesis is actually a state of the feature space (or particle), and the collection of hypothesis is a sampling of the feature space.

2.2 The State/Feature Vector

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

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

where the vessel state vector consists of the 3D location of the vessel $\mathbf{x}$, the tangent vector $\mathbf{\Theta}$, its exact position at a given cross-section (segmentation is done through an ellipse (α (major axis radius), β (minor axis radius), ϕ (orientation)) and the parameters required for the pdf estimation of the appearance of the vessel $\mathbf{p}_{vessel}$, as a mixture of two gaussians:

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

One can assume non-linearity in the appearance of the vessel because of the presence of calcifications, stents, stenosis and diseased vessel lumen [FIG. (2)]. 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 blood

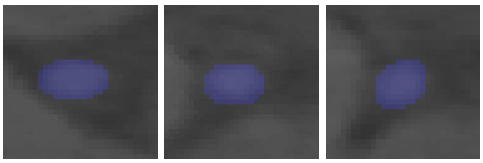

Figure 3: Three vessels cross sections detected using the ribbon measure. (P_B, μ_B, σ_B) and the calcification (P_C, μ_C, σ_C) subject to the constraint $[P_C + P_B = 1]$ leading to the following state vector:

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

2.3 Bayes Sequential Estimator

The Bayesian tracking problem can be simply formulated as the computation of the present state x_t pdf of a system, based on observations from time 1 to time t $z_{1:t}$: $p(x_t|z_{1:t})$. Assuming that one can have access to the prior pdf $p(x_{t-1}|z_{1:t-1})$, the posterior pdf $p(x_t|z_{1:t})$ can be computed according to the Bayes rule:

$$p(x_t|z_{1:t}) \;=\; \frac{p(z_t|x_t)p(x_t|z_{1:t-1})}{p(z_t|z_{1:t-1})},$$

where the prior pdf is computed via the Chapman-Kolmogorov equation

$$p(x_t|z_{1:t-1}) \;=\; \int p(x_t|x_{t-1})p(x_{t-1}|z_{1:t-1})dx_{t-1},$$

and

$$p(z_t|z_{1:t-1}) \;=\; \int p(z_t|x_t)p(x_t|z_{1:t-1})dx_t$$

The recursive computation of the prior and the posterior pdf leads to the exact computation of the posterior density. Nevertheless, in practical cases, it is impossible to compute exactly the posterior pdf $p(x_t|z_{1:t})$, which must be approximated.

2.4 Prediction & Observation: Distance

Once such a recursive paradigm was built, the next and last issue to be addressed is the definition of a measure between the distribution and the data through statistical means. 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 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 shape S and appearance A models, given the observation z:

$$p(S, A|z) \;=\; p(S|z) * p(A|z) \tag{3}$$

assuming statistical independence of shape and appearance (which is obviously not true in practice).

- Probability measure for shape

  The vessel cross section is modeled by an ellipse ϵ, for which the *ribbon measure* R is computed:

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

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

  where μ_{int} is the intensities mean value for the voxels in the ellipse, and μ_{ext} is the intensities mean value for the voxels in the ribbon around the ellipse, such that the ribbon and the ellipse have the same area.

  Since the coronary arteries are brighter than the background, the ellipse that best matches the vessel's cross section maximizes R (see [FIG. (3)]).

- 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:

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

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

The combination of edge-driven and region-based metrics can measure the fitness of the observation to the prior knowledge included in the state vector. Linear models are a reasonable approximation to sequential motion estimation. The Kalman filter is the most notable example among them.

2.5 Linear Case: The Kalman Filter

The Kalman filter [18] is a set of mathematical equations that provides efficient computational (recursive) mean to estimate the state of a process, in a way that minimizes the mean of the squared error. The filter is very powerful in several aspects: it supports estimations of past, present, and even future states, and it can do so even when the precise nature of the modeled system is unknown.

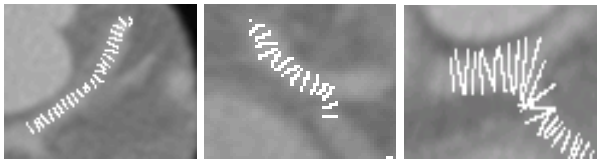

Figure 4: *Kalman filters & Vessel Segmentation. The non-linearity of the vessel geometric structure (in particular branching) is reflected on the Kalman filter results.*

Such a filter assumes that the posterior density is Gaussian at each time step, and that the current state x_t and observation z_t are linearly dependent on the past state x_{t-1}. Such assumptions simplify the Bayesian equations to the following form:

$$\begin{cases} x_t = F_t x_{t-1} + v_{t-1} \\ z_t = H_t x_t + n_t, \end{cases}$$

where v_{t-1} and n_t refer to zero mean Gaussian noise with covariance matrices Q_{t-1} and R_t that are assumed to be statistically independent. The matrix F_t is considered known and relates the former state x_{t-1} to the current state x_t. The matrix H_t is also known and relates the state x_t to the observation z_t. Then, the pdfs can be computed recursively according to the following formulas:

$$\begin{cases} p(x_{t-1}|z_{1:t-1}) = \mathcal{N}(x_{t-1}; m_{t-1|t-1}, P_{t-1|t-1}) \\ p(x_t|z_{1:t-1}) = \mathcal{N}(x_t; m_{t|t-1}, P_{t|t-1}) \\ p(x_t|z_{1:t}) = \mathcal{N}(x_t; m_{t|t}, P_{t|t}) \end{cases}$$

with

$$\begin{cases} m_{t|t-1} = F_t m_{t-1|t-1} \\ P_{t|t-1} = Q_{t-1} + F_t P_{t-1|t-1} F_t^T \\ m_{t|t} = m_{t|t-1} + K_t(z_t - H_t m_{t|t-1}) \\ P_{t|t} = P_{t|t-1} - K_t H_t P_{t|t-1} \end{cases}$$

where

$$K_t = P_{t|t-1} H_t^T \left(H_t P_{t|t-1} H_t^T + R_t\right)^{-1}$$

Some (negative) examples of the application of such a linear model to vessel segmentation are shown in [FIG. (4)], using the state space earlier introduced and the Kullback-Leibler information criterion to measure the distance between prediction and observation. Kalman filters have been considered to track vessels in retinal images [27] where particular handling to treat branching is considered.

However, the strong assumption of Gaussian noise and in particular the linearity case make their use in vessel segmentation quite problematic. One can claim

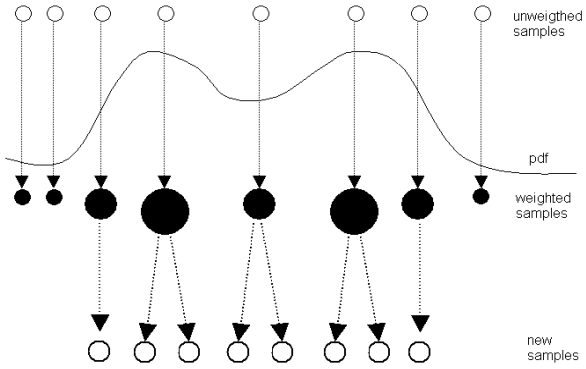

Figure 5: *The resampling process: a random selection chooses the samples with the highest weights where a local perturbation is applied.*

that neither the observation space [FIG. (2)], nor the structure/geometric space are linear [FIG. (4)] and such a method will fail to account for pathological cases where such linearity is absent.

3 Particle Filters

Particle filters [7, 15] are sequential Monte-Carlo techniques that can be used to estimate the Bayesian posterior probability density function (pdf) with a set of samples [13, 34]. 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

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

where each weight w_t^m reflects the importance of the sample x_t^m in the pdf, as shown in [FIG. (5)].

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

$$w_t^m \propto w_{t-1}^m \frac{p(z_t|x_t^m)p(x_t^m|x_{t-1}^m)}{q(x_t^m|x_{t-1}^m, z_t)}. \tag{4}$$

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

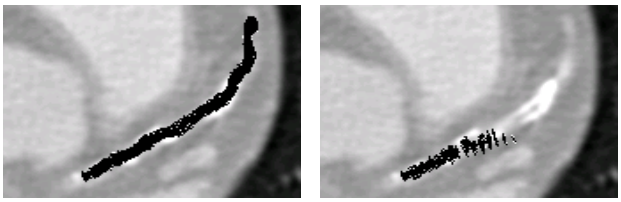

(a) (b)

Figure 6: *Particle filters & resampling; (a) least possible states are tolerated in SIR, leading to a more robust segmentation in very hard non-linear cases (b) only the most probable states are preserved in SRR*

ability to track significant changes on the pdf; therefore a resampling procedure has to be considered 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 consider two of them, the most prominent ones,

- SAMPLING IMPORTANCE RESAMPLING [13] where the most probable samples are randomly selected and are perturbed to create new samples,
- STRATIFIED RESAMPLING [9] where the samples of heavy weights are preserved and the rest of them are re-sampled.

3.1 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 $q(x_t|x^m_{1:t}, z_t)$. This leads to the following condition

$$w^m_t \propto w^m_{t-1} p(z_t|x^m_t). \tag{5}$$

The samples are updated by setting $x^m_t \propto p(x_t|x^m_{t-1})$, and perturbed according to a random noise vector.

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.

3.2 Stratified resampling

Stratified resampling [9] (SRR) is an alternative to the Sampling Importance Resampling. The central idea is (i) to threshold the samples according to their

weights, (ii) to preserve the samples with important weights, and (iii) to re-sample the least expected samples.

Such an operation is performed according to a constant factor $1/c$, where c is the unique solution of

$$\sum_{m=1}^{M} min(c\, w_t^m, 1) \;=\; N,$$

where $N\ (< M)$ is the number of particles of non-zero weight. It is then proved that the use of this threshold results into an optimum resampling process, over all unbiased resampling algorithms, in terms of minimizing

$$\sum_{m=1}^{M} E((W_t^m - w_t^m)^2),$$

where W_t^m is a random variable that describes all possible values that w_t^m can take. A resampling is said unbiased if $E(W_t^m) = w_t^m$.

The basic algorithm for computing c is first to order the set of weights $\{w_t^m\}_{m=1..M}$. Then, when the smallest weight $w_t^{m_0}$ that satisfies

$$\sum_{m=1}^{M} min(\frac{w_t^m}{w_t^{m_0}}, 1) \;\leq\; N,$$

is determined, we set $c = (N - A_{m_0})\,/B_{m_0}$, where A_{m_0} is the number of weights that are larger than $w_t^{m_0}$ ($M - m_0$ for an ordered set of weights), and B_{m_0} is the sum of the remaining elements. Finally, the set of samples can be split into two sets: the samples with weights above $1/c$ will be kept for the next time step, and the samples with weights below $1/c$ will be re-sampled.

4 Particle Filters & 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

$$(\ (P_B, \mu_B, \sigma_B)\ ,\ (P_C, \mu_C, \sigma_C)\ )$$

using an expectation/maximisation 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 $\omega_\tau = (\mathbf{x}_\tau, \Theta_\tau, \epsilon_\tau)$, which belong to the feature space (see Section 2.2) where we use the notion of particle filters.

In other words, given a current position of a given particle, we perform a random perturbation on its state vector ω_τ therefore creating a new configuration $\omega_{\tau+1}$ that refers to a new cross-section, as well as to the a new segment of the vessel. Once such a configuration is available, one can recover the distance between the prior state of the vessel and the new potential state as:

$$v(p|q) = e(-D(p|q)) + \delta \oint_{\epsilon} e^{-g(|\nabla I|)} d\epsilon$$

that is used within the sequential estimator to determine the strength of the new configuration and to update the pdf of this particular particle. In practice, given the root of the vessel, a certain number of possible (random) states are introduced that are propagated to perform segmentation. The number of particles as well as the resampling within such an approach are the most challenging issues to be addressed. Once a resampling step is considered, a Kalman filter is run over several time steps to obtain the statistical distribution model (q) that refers to prior knowledge or prior most prominent state that is then used to update the particle's weight and estimate the posterior pdf.

4.1 Resampling

One should mention that several particles will converge to highly improbable configurations and therefore have to be eliminated. Furthermore, the coronaries consist of a tree with increasing complexity, therefore the number of particles has to be adaptive, in particular when situations like branching arise. Regarding the initial configuration, the use of approximatively $1,000$ particles gave the sufficient results for our experiments. Based on the time needed to observe an anatomical change on the vessel, we perform a systematic resampling according to the Sampling Importance Resampling every 10 cross sections. The number of time steps between two resamplings was a compromise between low complexity and ability to capture the anatomy of the vessel [3]. The preference for SIR, compared to SRR, is motivated by the robustness of the segmentation, see [FIG. (6)].

4.2 Branching detection

When a branching occurs, the particles split up in the two daughter branches, and then track them separately (see [FIG. (7)]). In order to detect branchings and

[3] One can assume that branching will be observable in at at least such a number of cross sections.

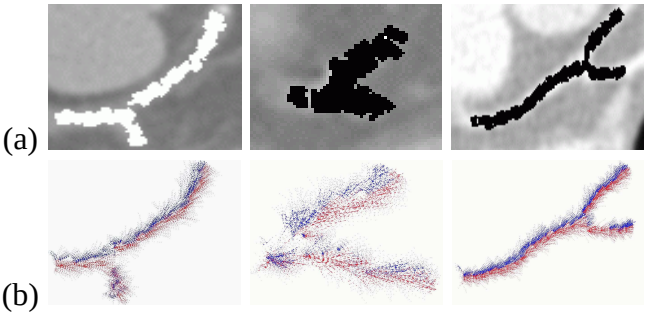

Figure 7: *(a) branching points between LCX and LAD for three patients with the particles' mean state overlaid, (b) the particles , clustered using K-means, follow up the two branches.*

eventually introduce new particles, simple clustering techniques are considered. To this end, a simple K-means approach on the joint space (position+orientation) of the particles can be 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. (4)], with a bi-modal distribution.

The K-means algorithm [8] partitions N points into K disjoint clusters, minimizing the sum-of-squares

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

4.3 Using Kalman Filter for Prior Knowledge

One of the reasons for medical imaging being among the most successful application areas of computer vision is the domain specific knowledge that is available to the user. In most cases, the structures to be recovered follow a given anatomy that, even if not inherently linear, can be so modeled to improve algorithms. To this end, one can consider linear models for a precise prior probability estimation in the particle filter framework. Such information can be used to determine the number of particles, constrain their propagation and improve the resampling strategy.

Vessels can be considered locally linear both in their shape and appearance; therefore Kalman filters [18] track properly small vessel segments. Using this characteristic, these filters have been used to provide prior knowledge. Run over few steps, the Kalman filters give a priori knowledge that consists of statistics in the feature space of the vessel segment:

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

This prior knowledge, computed and updated as the particle filter runs, is used in two different ways: (i) to compute the observation pdf $p(z_t|x_t)$, (ii) in the im-

portance density $q(x_t|x^m_{1:t}, z_t)$, to select the particles that are closer to the prior knowledge in the feature space.

- The pixel intensity distribution $((P_B, \mu_B, \sigma_B), (P_C, \mu_C, \sigma_C))$, corrected by the Kalman filter, is used as a model; the particle's probability measure is then equal to the distance in the distribution space between the model and the observed distribution. That way, the statistics model is regularly updated, and takes into account any non-linear state transition (such as diseases or prosthesis for the case of coronary arteries, see [FIG. (2)]).

- The Kalman filter prior gives a precious piece of information about the position / orientation of the vessel. When the linear filter is run right before the resampling process, the importance density $q(x_t|x^m_{1:t}, z_t)$ ($= p(x_t|x_{t-1})$ for SIR) takes account of the prior knowledge in terms of position/orientation. Therefore, while being resampled, the particles belonging to the vessel (and thus responding positively to $p(x_t|x_{t-1})$, given the prior knowledge) are more likely to be selected. The number of particles can be reduced in this way.

The pdf $p(z_t|x^m_t)$ in [EQ. (5)] is computed from the distance between the Kalman shape appearance prediction q and the actual observation p. The distance we use is the symmetrized Kullback-Leibler distance:

$$\int p(x)log(\frac{p(x)}{q(x)}) \; + \; q(x)log(\frac{q(x)}{p(x)})dx.$$

4.4 Implementation and Validation

One should mention that several particles will converge to highly improbable configurations and therefore have to be eliminated. Furthermore, the coronaries consist in a tree of increasing complexity, therefore the number of particles has to be adaptive, in particular when situations like branching arise. Regarding the initial configuration, the use of approximatively $1,000$ particles gave the sufficient results for our experiments. Based on the time needed to observe an anatomical change on the vessel, 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 ith particle) falls below half the number of particles. The number of time steps between two resamplings was a compromise between low complexity and ability to capture the anatomy of the vessel [4]. The preference for SIR, compared to Stratified Resampling [20], is motivated by the robustness of the segmentation.

[4] One can assume that branching will be observable in at least such a number of cross sections.

Validation is a difficult part for any coronary segmentation method. The algorithm has been run on 28 patients, and has successfully recovered all the main arteries (RCA, LCA, LCX) for each patient (see [FIG. (8)]). These results have been achieve 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.

Although the main branchings were correctly detected, some of the 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.

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 diseased patients CTA data sets, segmenting the *Left Anterior Descending* and the *Right Coronary Artery* [FIG. (8)].

Introducing further prior knowledge in the segmentation process is the most prominent future direction. Current efforts consist of the use of linear models to account for prior information. In parallel to that, the use of non parametric, non-linear approximations of prior knowledge, in terms of anatomy as well as appearance, could significantly increase the performance of particle filters for vessel segmentation.

References

[1] N. Armande, P. Montesinos, O. Monga, and G. Vaysseix. Thin nets extraction using a multi-scale approach. *Computer Vision and Image Understanding: CVIU*, 73(2):248–257, 1999.

[2] S. Bouix, K. Siddiqi, and A. R. Tannenbaum. Flux driven automatic centerline extraction. In *Medical Image Analysis*, 2004.

[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] A. Doucet, J. de Freitas, and N. Gordon. *Sequential Monte Carlo Methods in Practice*. Springer-Verlag, New York, 2001.

[8] R Duda and P. Hart. *Pattern Classification and Scene Analysis*. John Wiley and Sons, 1973.

[9] P. Fearnhead and P. Clifford. Online inference for well-log data. *Journal of the Royal Statistical Society*, 65:887–899, 2003.

[10] S. D. Fenster, C. G. Kuo, and J. R. Kender. Nonparametric training of snakes to find indistinct boundaries. In *Mathematical Methods in Biomedical Imaging Analysis*, Dec. 2001.

[11] M. Figueiredo and J. Leitao. A nonsmoothing approach to the estimation of of vessel controus 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] R. Kalman. A New Approach to Linear Filtering and Prediction Problems. *ASME–Journal of Basic Engineering*, 82:35–45, 1960.

[19] K. Krissian, G. Malandain, N. Ayache, R. Vaillant, and Y. Trousset. Model based detection of tubular structures in 3d images. *Computer Vision and Image Understanding*, 80:130–171, 2000.

[20] Jun S. Liu and Rong Chen. Sequential Monte Carlo methods for dynamic systems. *Journal of the American Statistical Association*, 93(443):1032–1044, 1998.

[21] L. Lorigo, O. Faugeras, E. Grimson, R. Keriven, R. Kikinis, A. Nabavi, and C. Westin. Codimension-Two Geodesic Active Controus for the Segmentation of Tubular Structures. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages I:444–451, 2000.

[22] R. Malladi and J. Sethian. A Real-Time Algorithm for Medical Shape Recovery. In *International Conference on Computer Vision*, pages 304–310, 1998.

[23] D. Nain, A. Yezzi, and G. Turk. Vessel Segmentation Using a Shape Driven Flow. In *Medical Imaging Copmuting and Computer-Assisted Intervention*, 2004.

[24] T. O´ Donnell, T. Boult, X. Fang, and A. Gupta. The Extruded Generalized Cylider: A Deformable Model for Object Recovery. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 174–181, 1994.

[25] S. Osher and N. Paragios. *Geometric Level Set Methods in Imaging, Vision and Graphics*. Springer Verlag, 2003.

[26] R. Petrocelli, K. Manbeck, and J. Elion. Three Dimentional Structue Recognition in Digital Angiograms using Gauss-Markov Models. *IEEE Computers in Radiology*, pages 101–104, 1993.

[27] P. Quelhas and J. Boyce. Vessel Segmentation and Branching Detection using an Adaptive Profile Kalman Filter in Retinal Blood Vessel Structure. In *Pattern Recognition and Image Analysis: First Iberian Conference*, pages 802–809, 2003.

[28] 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.

[29] Y. Sato, S. Nakajima, H. Atsumi, T. Koller, G. Gerig, S. Yoshida, and R. Kikinis. 3D Multiscale 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.

[30] J. Sethian. A Review of the Theory, Algorithms, and Applications of Level Set Methods for Propagating Interfaces. *Cambridge University Press*, pages 487–499, 1995.

[31] E. Sorantin, C. Halmai, 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.

[32] K. Toyama and A. Blake. Probabilistic Tracking in a Metric Space. In *IEEE International Conference in Computer Vision*, pages 50–59, 2001.

[33] A. Vasilevskiy and K. Siddiqi. Flux Maximizing Geometric Flows. In *IEEE International Conference in Computer Vision*, pages I: 149–154, 2001.

[34] W. West. Modelling with mixtures. In J. Bernardo, J. Berger, A. Dawid, and A. Smith, editors, *Bayesian Statistics*. Clarendon Press, 1993.

[35] 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.

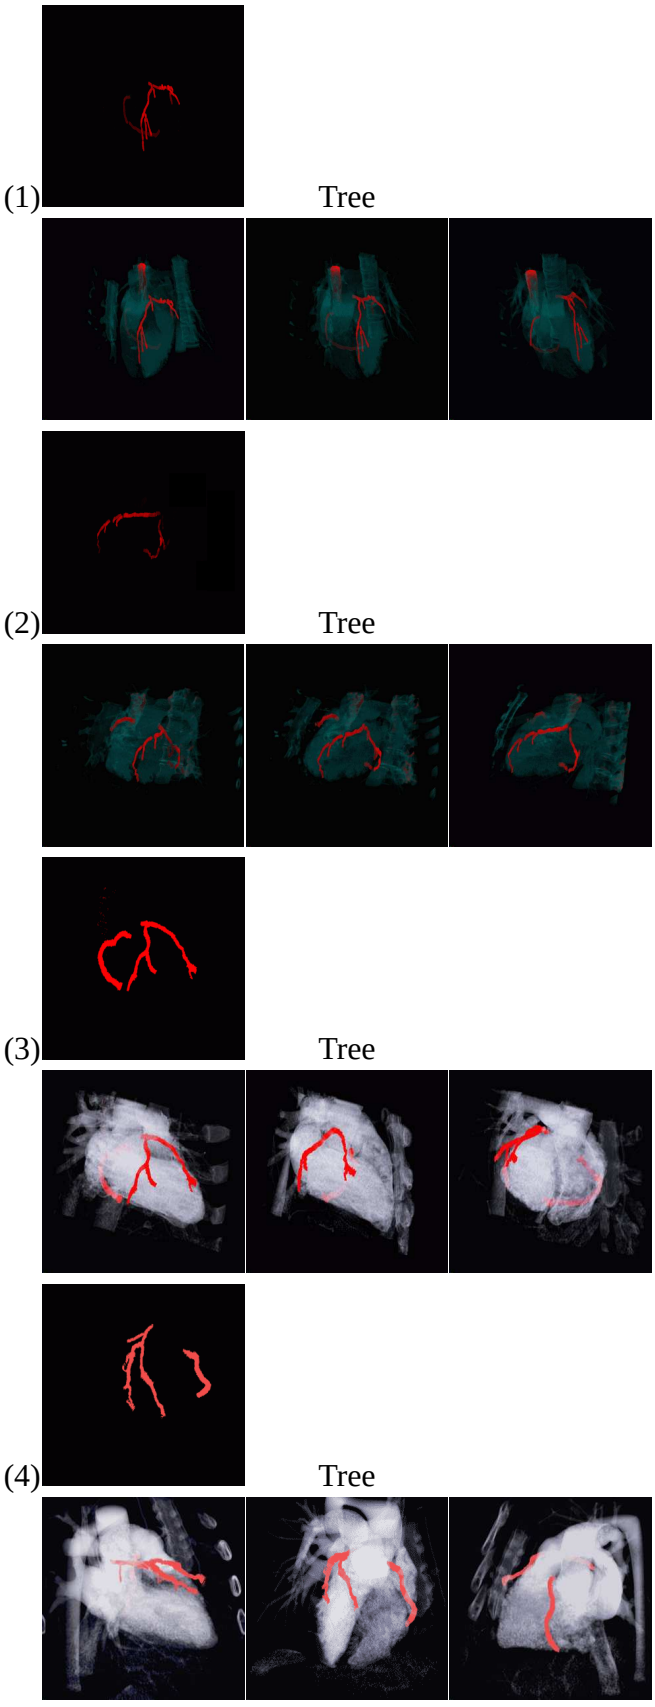

Figure 8: *Segmentation of the Left anterior descending coronary artery and Right coronary artery in CTA (in red) for four patients (1) & (2) & (3) & (4); Different* $3D$ *views super-imposed to the cardiac volume are presented.*