Particle Filters, a Quasi-Monte-Carlo-Solution for Segmentation of Coronaries

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

[1] Imaging & Visualization Department,
Siemens Corporate Research, Princeton, NJ, USA
[2] CERTIS - Ecole Nationale des Ponts et Chaussees,
Champs-sur-Marne, France

Abstract. In this paper we propose a Particle Filter-based approach for the segmentation of coronary arteries. To this end, successive planes of the vessel are modeled as unknown states of a sequential process. Such states consist of the orientation, position, shape model and appearance (in statistical terms) of the vessel that are recovered in an incremental fashion, using a sequential Bayesian filter (Particle Filter). In order to account for bifurcations and branchings, we consider a Monte Carlo sampling rule that propagates in parallel multiple hypotheses. Promising results on the segmentation of coronary arteries demonstrate the potential of the proposed approach.

1 Introduction

Cardio-vascular diseases are the leading cause of deaths in the USA (39%) and therefore there is a demand for constant improvement of diagnostic tools to detect and measure anomalies in the cardiac muscle. Coronary arteries are thin vessels that feed the heart muscle in blood. Therefore, their segmentation provides a valuable diagnostic tool for clinicians interested in detecting calcifications, and stenosis. Because of the low contrast conditions, and the coronaries vicinity to the blood pool, segmentation is a difficult task. Computer Tomography (CT) and Magnetic Resonance (MR) imaging of the heart have became standard tools to medical diagnosis resulting to a substantial number of patients being imaged.

On one hand, vessel segmentation techniques consist of model-free and model-based methods. Vessel enhancement approaches [6] and differential geometry-driven methods [11] do not segment vessels per se, but allow a better visualization. Region growing [21], flux maximization [2], morphological operators [5] and skeleton-based techniques [19] are more advanced vessel segmentation techniques. On the other hand, model-based techniques use prior knowledge and features to match a model with the input image and extract the vessels. Prior knowledge refer either to the whole structure, or to the local vessel model. Tracking approaches recover the vessel centerline - given a starting condition - through processing information on the vessel cross section [10]. Vessel template matching [17], generalized cylindrical models [15] as well as parametric/geometric deformable models [18] are alternatives to vessel tracking and seek to minimize an objective function computed along the model.

J. Duncan and G. Gerig (Eds.): MICCAI 2005, LNCS 3749, pp. 246–253, 2005.
© Springer-Verlag Berlin Heidelberg 2005

Level sets [16] is an established method to address such minimization [12]. One can refer to the fast marching algorithm and its variant for vessel segmentation using the minimal path principle [1]. To discourage leaking, a local shape term that constrains the diameter of the vessel was proposed in [14].

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.

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. 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. In the proposed technique, unlike standard techniques where the most probable hypothesis is maintained, a discrete number of states (possible solutions) remain active and are associated with a probability density function. The final paradigm consists of a fast multiple hypothesis adaptive propagation technique where the vessel structure and its appearance are successfully recovered. Such a framework allows to naturally address the non-linearities 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 and introduce the concept of the proposed approach and Particle Filters, while vessel segmentation is presented in section 3. Implementation and validation are part of section 4, while discussion is part of the last section.

2 Preliminaries and Particle Filters

To explain our method at a concept level, let us assume that a segment of the vessel has been detected: a 2D shape on a 3D plane. Similar to region growing and front propagation techniques, our method aims to segment the vessel in adjacent planes. To this end, one can consider the hypotheses ω of the vessel being at a certain location ($\mathbf{x}$), having certain orientation ($\mathbf{\Theta}$), and referring to certain shape - an elliptic model is a common choice (ϵ) - with certain appearance characteristics ($\mathbf{p}_{vessel}$).

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

Then, segmentation consists in finding the optimal parameters of ω given the observed 3D volume. Let us consider a probabilistic interpretation of the problem with $\pi(\omega)$

being the posterior distribution that measures the fitness of the vector ω with the observation. Under the assumption that such a law is present, segmentation consists in finding at each step the set of parameters ω that maximizes $\pi(\omega)$. However, since such a model is unknown, one can assume an autoregressive mechanism that, given prior knowledge, predicts the actual position of the vessel and a sequential estimate of its corresponding states. To this end, we define:

- a state/feature vector ω,
- an iterative process to predict the next state and update the density function, that can be done using a Bayes sequential estimator and is based on the computation of the present state ω_t pdf of a system, based on observations from time 1 to time t $z_{1:t}$: $\pi(\omega_t|z_{1:t})$. Assuming that one has access to the prior pdf $\pi(\omega_{t-1}|z_{1:t-1})$, the posterior pdf $\pi(\omega_t|z_{1:t})$ is computed according to the Bayes rule:

$$\pi(\omega_t|z_{1:t}) \;=\; \frac{\pi(z_t|\omega_t)\pi(\omega_t|z_{1:t-1})}{\pi(z_t|z_{1:t-1})}.$$

  The recursive computation of the prior and the posterior pdf leads to the exact computation of the posterior density.
- a distance between prediction and actual observation, based on the observation.

Kalman filter is the most popular variant of this model, a well known linear approach able to track vessels with limited variation in appearance and geometry. Cardiac vessel trees are highly irregular. Random bifurcations, branches of variable width, non-linear visual properties because of the presence of calcifications, stents, stenosis and diseased vessel lumen are some examples demonstrating the non-linearity of the vessel tree as shown in [Fig. (1)].

Consequently simple parametric statistical models will fail to account for the statistical and geometric properties of the vessel leading to the consideration of more complex distributions. To this end, instead of one single prediction, a collection of hypotheses can be generated at each step and being evaluated using the distance between prediction and actual observation. Nevertheless, in practical cases, it is impossible to compute exactly the posterior pdf $\pi(\omega_t|z_{1:t})$, which is to be approximated. An elegant approach to implement such a technique refers to the use of particle filters where each given hypothesis is a state in the feature space (or particle), and the collection of hypothesis is a sampling of the feature space.

Particle Filters [3, 9] are sequential Monte-Carlo techniques that are used to estimate the Bayesian posterior probability density functions (pdf) [7, 20]. In terms of a mathematical formulation, such a method approximates the posterior pdf by M random measures $\{\omega_t^m, m = 1..M\}$ associated to M weights $\{\lambda_t^m, m = 1..M\}$, such that

$$\pi(\omega_t|z_{1:t}) \;\approx\; \sum_{m=1}^{M} \lambda_t^m \delta(\omega_t - \omega_t^m), \tag{1}$$

where each weight λ_t^m reflects the importance of the sample ω_t^m in the pdf.The samples ω_t^m are drawn using the principle of *Importance Density* [8], of pdf $q(\omega_t|x_{1:t}^m, z_t)$, and it is shown that their weights λ_t^m are updated according to

$$\lambda_t^m \propto \lambda_{t-1}^m \frac{\pi(z_t|\omega_t^m)\pi(\omega_t^m|\omega_{t-1}^m)}{q(\omega_t^m|\omega_{t-1}^m, z_t)}. \tag{2}$$

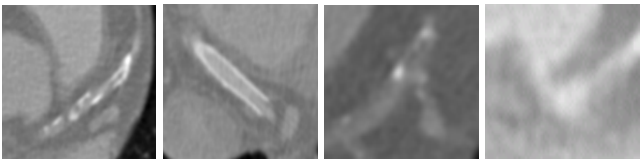

(a) (b) (c) (d)

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

Once a set of samples has been drawn, $\pi(\omega_t^m|\omega_{t-1}^m, z_t)$ can be computed out of the observation z_t for each sample, and the estimation of the posteriori pdf can be sequentially updated.

3 Vessel Segmentation

We now consider the application of such a 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 prior automatic procedure. Simple segmentation of that area can provide an initial guess on the statistical properties of the vessel appearance. It is reasonable to assume irregularity in the appearance $\mathbf{p}_{vessel}$ 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:

$$\omega = (\mathbf{x}, \Theta, \epsilon, (P_B, \mu_B, \sigma_B), (P_C, \mu_C, \sigma_C)) \tag{3}$$

The vessel state vector consists of the 3D location of the vessel $\mathbf{x}$, the tangent vector $\mathbf{\Theta}$, its shape model at a given cross-section (the model used here is an ellipse (α (major axis radius), β (minor axis radius), ϕ (orientation)) and the appearance $\mathbf{p}_{vessel}$, mixture of two gaussians.

Once such a recursive paradigm is built, the next and last issue to be addressed is the definition of a measure between a prediction and the actual observation. 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 expectation-maximization principle.

Let us now consider a random state vector ω, that refers to a certain segmentation hypothesis that is to be evaluated ($p(\omega|\mathcal{D})$) where $\mathcal{D}$ is the observed 3D volume. Such a hypothesis should refer to a region that has consistent visual properties with the ones expected ($\mathbf{p}_{vessel}$). While the separation of the vessels from the cardiac muscle is a rather tedious task (blood is present in both organs), one can claim that their separation

from the liquid of the vascular structure is possible and can be used to validate the goodness of a hypothesis.

For the vessel lumen pixels distribution $\mathbf{p}_{vessel}$, 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_{ap} between the model $p(\omega) = \mathbf{p}_{vessel}$ and the observation $q(\omega)$:

$$D_{ap} \;=\; \int p(\omega) log\left(\frac{p(\omega)}{q(\omega)}\right) \;+\; q(\omega) log\left(\frac{q(\omega)}{p(\omega)}\right) d\omega,$$

which have important values when the distance between these two distributions is significant. Therefore, one can consider the following measure $\left[p(\omega|D_{ap}) \;=\; e^{-\frac{|D_{ap}|}{\sigma_{ap}}}\right]$ where σ_{ap} is a normalization factor. Toward discriminating the vessel from the vascular liquid one can consider a ribbon measure

$$D_{rb} = \begin{cases} -\infty \;, & \mu_{int} \leq \mu_{ext} \\ \frac{\mu_{int}-\mu_{ext}}{\mu_{int}+\mu_{ext}}, & otherwise \end{cases}$$

where μ_{in} is the mean within the ellipse and μ_{ext} is the mean within a ring centered at the ellipse center with greater radius (the ring area is equal to the inner circle area). Such a measure aims at maximizing the distance between the mean values of the interior and the exterior region [13], based on the fact that the coronary arteries are brighter than the background, and can also be used to measure the fitness of the segmentation: $\left[p(\omega|D_{rb}) \;=\; e^{-\frac{|D_{rb}|}{\sigma_{rb}}}\right]$. We assume that the two conditions are independent and therefore one can multiply the two measures to determine the goodness of the hypothesis under consideration.

Given a starting point and a number of particles, one now performs random perturbations to each particle in the feature space. Once a perturbation has been applied, the corresponding hypothesis is evaluated using the visual matching and the ribbon measure introduced earlier. At each step of the process, segmentation refers to a weighted linear combination of the state vectors (the particles) [EQ. (1)].

4 Implementation

After certain iterations, 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. At the same time, in the presence of bifurcations, new hypotheses are to be introduced in order to capture the entire vessel tree. 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 [4].

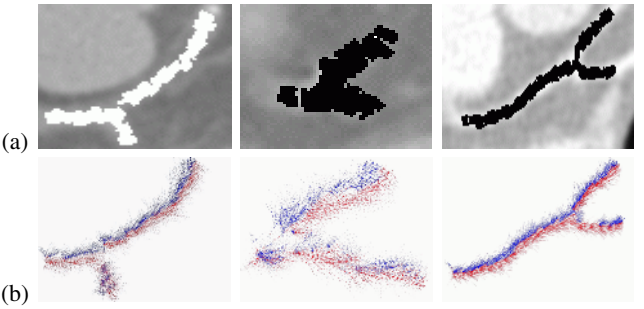

Fig. 2. (a) branching points between LCX and LAD for three patients with the particles' mean state overlaid, (b) the particles , clustered using K-means, follow up the two branches

The Sampling Importance Resampling (SIR) algorithm [7] consists in choosing the prior density $\pi(\omega_t|\omega_{t-1})$ as importance density $q(\omega_t|\omega^m_{1:t}, z_t)$. This leads to the following condition, from [EQ. (2)]: $\lambda^m_t \propto \lambda^m_{t-1}\pi(z_t|\omega^m_t)$.

The samples are updated by setting $\omega^m_t \propto \pi(\omega_t|\omega^m_{t-1})$, 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 $\pi(\omega_t|\omega_{t-1})$, 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.

Particular attention is also to be paid during the resampling process to address branching and bifurcations. When a branching occurs, the particles split up in the two daughter branches, and then track them separately (see [Fig. (2)]). Although Particle Filters track the two branches, experiments have shown a branching detection heuristics improves the results. To this end, a simple K-means approach 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.

Regarding the initial configuration, the use of approximatively $1,000$ particles gave sufficient results for our experiments. We perform a systematic resampling according to the Sampling Importance Resampling every time the effective sampling size $N_{eff} = \sum_i 1/\lambda_i^2$ (where λ_i is the weight of the i^{th} particle) falls below half the number of particles. The preference for SIR, compared to Stratified Resampling [4], is motivated by the robustness of the segmentation.

5 Discussion

In this paper, we have proposed a particle-filter based approach to vascular segmentation. Experiments were conducted on several healthy and diseased patients CTA data sets, segmenting both the *Left Main Coronary Artery* and the *Right Coronary Artery*.

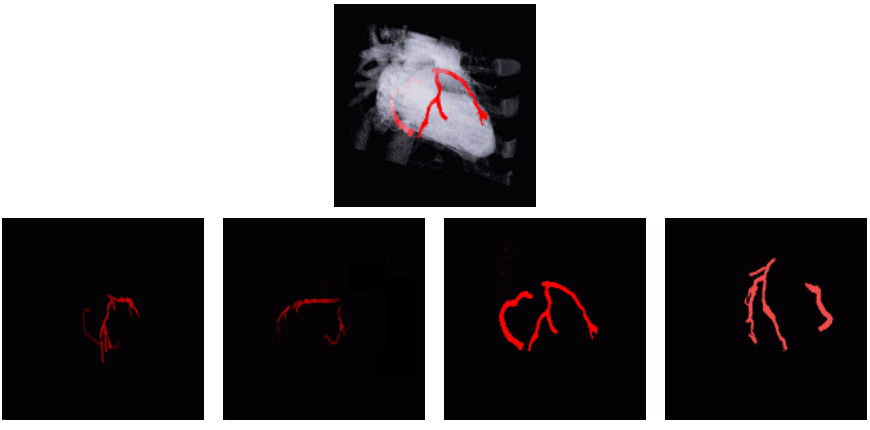

Fig. 3. Segmentation of the Left anterior descending coronary artery and Right coronary artery in CTA (in red) for four patients, and full view with heart

Validation is a challenging but required step 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

vessel name	RCA	Acute Marginal	LAD	First Septal	LCX	Obtuse Marginal
% of cases segmented	100%	85.3%	100%	94%	100%	94%

a small portion of visual results are also presented in [Fig. (3)]. The table indicates the number of branches (in percentage) successfully segmented.

These results were achieved with a one-click initialization. From the first point provided by the user, the initial direction is determined as the direction of minimal gradient variation. 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 successfully detects all 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.

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. The use of an image term and a statistical model makes the probability measure both robust to pathologies, and yet, drives the segmentation toward the most probable solution given the statistical prior.

Future work consists in learning the variation law that rules the feature space toward better tests for hypotheses validation, as well as the one that controls process noise, to better guide the resampling stage toward an intelligent reduction of the required number of particles.

References

1. 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.
2. 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.
3. A. Doucet, J. de Freitas, and N. Gordon. *Sequential Monte Carlo Methods in Practice*. Springer-Verlag, New York, 2001.
4. P. Fearnhead and P. Clifford. Online inference for well-log data. *Journal of the Royal Statistical Society*, 65:887–899, 2003.
5. 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.
6. 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.
7. N. Gordon. Novel Approach to Nonlinear/Non-Gaussian Bayesian State Estimation. *IEE Proceedings*, 140:107–113, 1993.
8. N. Gordon. On Sequential Monte Carlo Sampling Methods for Bayesian Filtering. *Statistics and Computing*, 10:197–208, 2000.
9. 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.
10. M. Hart and L. Holley. A method of Automated Coronary Artery Trackin in Unsubtracted Angiograms. *IEEE Computers in Cardiology*, pages 93–96, 1993.
11. 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.
12. R. Malladi and J. Sethian. A Real-Time Algorithm for Medical Shape Recovery. In *International Conference on Computer Vision*, pages 304–310, 1998.
13. D. Mumford and J. Shah. Boundary detection by minimizing functionals. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 22–26, 1985.
14. D. Nain, A. Yezzi, and G. Turk. Vessel Segmentation Using a Shape Driven Flow. In *Medical Imaging Copmuting and Computer-Assisted Intervention*, 2004.
15. 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.
16. S. Osher and N. Paragios. *Geometric Level Set Methods in Imaging, Vision and Graphics*. Springer Verlag, 2003.
17. 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.
18. 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.
19. 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.
20. W. West. Modelling with mixtures. In J. Bernardo, J. Berger, A. Dawid, and A. Smith, editors, *Bayesian Statistics*. Clarendon Press, 1993.
21. 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.