

MEDIAL-BASED BAYESIAN TRACKING FOR VASCULAR SEGMENTATION: APPLICATION TO CORONARY ARTERIES IN 3D CT ANGIOGRAPHY

David Lesage^{*†}, Elsa D. Angelini[†], Isabelle Bloch[†] and Gareth Funka-Lea^{*}

^{*}Siemens Corporate Research, Imaging and Visualization dept., Princeton NJ, USA.

[†]Institut Telecom, Telecom ParisTech, CNRS LTCI, Paris, France.

ABSTRACT

We propose a new Bayesian, stochastic tracking algorithm for the segmentation of blood vessels from 3D medical image data. Inspired by the recent developments in particle filtering, it relies on a constrained, medial-based geometric model and on an original sampling scheme for the selection of tracking hypotheses. A key property of this new sampling scheme is the ability to take into account a distribution of hypotheses broader than similar methods such as classical particle filters, while remaining computationally efficient. The proposed method was applied to the challenging and medically critical task of coronary artery segmentation from 3D cardiac computed tomography (CT) images. Prior knowledge, injected in the process, was learned from a manually segmented database of 19 cases. Qualitative and quantitative evaluation is presented on clinical data, including pathologies and local anomalies.

Index Terms— Bayesian tracking, Monte-Carlo method, Geometric model, Vascular segmentation, Cardiac CTA

1. INTRODUCTION

In biomedical applications, vascular structures are often of critical importance for diagnosis, treatment and surgery planning. Vessels are thin, elongated and complex structures embedded in increasingly large images. Manual delineation, although still heavily used in clinical routines, has become a considerable burden and automatic or semi-automatic segmentation remains challenging.

Vascular segmentation has received considerable attention in the literature [1]. A popular approach is to consider the segmentation as an iterative, *tracking* process. Classical region-growing techniques can be seen as primitive representatives of this class of methods. Front propagation techniques allow for refined analysis by imposing a structurally coherent exploration process. The robustness of local deterministic tracking is generally limited by the necessity of using low-level, causal criteria. In some settings, the tracking problem can be formulated as the extraction of globally minimal paths [2, 3]. Another approach, which is becoming increasingly popular, is the use of stochastic Bayesian tracking algorithms such as particle filters [4, 5]. Such algorithms have demonstrated particular robustness while allowing for high-level modeling.

In this paper, we propose a new Bayesian, stochastic tracking algorithm, inspired by recent developments in particle filtering design. It relies on a constrained, medial-based geometric model and on an original sampling scheme for the selection of tracking hypotheses. It was applied to coronary artery segmentation from 3D cardiac computed tomography (CT) images, a particularly challenging task in terms of anatomical variability and complexity. Prior knowledge was learned from a manually segmented database. Qualitative and quantitative validation is presented on clinical data.




Fig. 1. Geometric model and successor candidates.

The gray area indicates the potential successors of $x_t^{(i)}$, e.g. $x_{t+1}^{(j)}$.

2. GEOMETRIC MODEL

For the representation of vascular structures, we advocate the use of parametric geometric models for their conciseness and usability for further analysis. Our model follows directly from Blum's seminal work on medial representations [6]. A binary object is represented by a collection of maximally included hyper-spheres $\{S_k\}$ of centers $\{x_k\}$ and radiuses $\{r_k\}$. As illustrated in Figure 1, sphere centers form the *medial axis* of the object and associated radiuses allow for the reconstruction of the object surface [6].

Our model considers a subset of Blum's representation, defined by a geometric constraint. In our tracking framework, our geometric model is iteratively updated by adding new spheres centers which lie on the surface of the preceding sphere. As depicted in Figure 1, a new center x_t is at distance r_{t-1} of x_{t-1} . A non-branching structure is then represented by a *discrete* chain of spheres $\{S_0, \dots, S_T\}$.

Thanks to this connectivity constraint, our model is entirely described with the collection of centers $\{x_k\}$. The radius r_t and axis direction d_t appear as linked variables:

$$\vec{d}_t = x_{t+1} - x_t \quad \text{and} \quad r_t = |x_{t+1} - x_t|$$

In the following, we reduce all the notations to the centers $\{x_t\}$, keeping in mind that a center at time t sets radius and direction variables at time $t - 1$ through our geometric connectivity constraint.

Numerous tracking techniques for tubular structures rely on similar medial-based models. In the vascular segmentation field, one can cite particular efforts for centerline tracking based on cross-section detection [4, 5, 7, 8, 9]. Such techniques incorporate the direction and scale in their state vectors, resulting in large search spaces, especially in 3D. In contrast, our search space is limited to the geometric location of the successive centers, but we are intrinsically limited to tubular structures with *circular* cross-sections. This geometric model directly applies to elongated structures in 2D and 3D images. It can be extended to branching structures by using a tree or graph of spheres instead of a simple chain.

3. TRACKING SCHEME

The proposed tracking scheme is an iterative, population-based, stochastic Bayesian filter, largely inspired by recent developments in particle filtering for segmentation tasks. It aims at recovering a single vascular branch, from its root to its distal end, as a chain of spheres centers $X_T = \{x_0, \dots, x_T\}$ following the geometric model described previously and estimated given the observations, *i.e.*, the image, considered as stationary: $Y_t = Y, \forall t$.

Similarly to classical particle filters, we use a first order hidden Markov model and a Monte-Carlo approximation of the posterior distribution $p(X_t|Y)$ by a weighted population of N_t discrete samples $\{x_t^{(i)}, w_t^{(i)}\}_{i=1}^{N_t}$ [10]. Instead of using classical *update* and *prediction* steps, we introduce a new sampling scheme to recursively evolve this population.

At time-step t , we consider all the potential successors $\{x_{t+1}^{(i,j)}\}$ of $x_t^{(i)}$ as the discrete locations at distances $r_t^{(i,j)} \in [0, r_{max}]$ of $x_t^{(i)}$ (gray area in Figure 1). Importantly, a candidate $x_{t+1}^{(i,j)}$ *retrospectively* fixes radius $r_t^{(i,j)}$ and direction $\vec{d}_t^{(i,j)}$ according to our model. Associated weights $w_{t+1}^{(i,j)}$ are evaluated as follows:

$$w_{t+1}^{(i,j)} \propto w_t^{(i)} p(Y|x_{t+1}^{(i,j)}) p(x_{t+1}^{(i,j)}|x_t^{(i)})$$

where $p(Y|x_{t+1}^{(i,j)})$ and $p(x_{t+1}^{(i,j)}|x_t^{(i)})$ are the likelihood and dynamic prior of the candidate $x_{t+1}^{(i,j)}$, respectively. The *extended* population $\{x_{t+1}^{(i,j)}, w_{t+1}^{(i,j)}\}$ forms a distribution from which the new population of N_{t+1} samples $\{x_{t+1}^{(i)}, w_{t+1}^{(i)}\}_{i=1}^{N_{t+1}}$ is obtained by *importance* sampling [10]. When a candidate $(x_{t+1}^{(i,j)}, w_{t+1}^{(i,j)})$ is selected, it enters the new population as $(x_{t+1}^{(i)} = x_{t+1}^{(i,j)}, w_{t+1}^{(i)} = w_{t+1}^{(i,j)})$. Weights are then normalized before recursively proceeding to the next step.

This scheme can be seen as a systematic resampling step with *delayed* application of the dynamic prior. We emphasize that, in contrast with classical particle filters, our algorithm actually takes into account a wider range of hypotheses by evaluating *all* the potential successors $\{x_{t+1}^{(i,j)}\}$. Instead of selecting a limited number N_{t+1} of surviving samples through a prediction step, it rather considers a much larger pool ($N_t \times N_{r_{max}}$, with $N_{r_{max}}$ the number of discrete locations in the maximal sphere considered). The dynamic prior is applied retrospectively to these candidates. Another positive side effect of our geometric model and sampling scheme is the elimination of the need for a geometric step parameter. Our filter progresses adaptively, according to the sphere connectivity constraint.

3.1. Likelihood Estimation: $p(Y|x_{t+1}^{(i,j)})$

Our likelihood function is a conjunctive combination of three terms:

$$p(Y|x_{t+1}^{(i,j)}) = F(r_t^{(i,j)}) E(r_t^{(i,j)}) G(x_{t+1}^{(i,j)}) \quad (1)$$

The term $F(r_t^{(i,j)})$ measures the gradient flow through the surface of the sphere $S_t^{(i,j)}$ of radius $r_t^{(i,j)} = |x_{t+1}^{(i,j)} - x_t^{(i)}|$ centered on $x_t^{(i)}$:

$$F(r_t^{(i,j)}) = \oint_{S_t^{(i,j)}} \langle \nabla(I), \vec{n} \rangle dS_t^{(i,j)}$$

where $\nabla(I)$ is the image gradient and \vec{n} the unit normal to $S_t^{(i,j)}$. It quantifies the alignment of the sphere surface with the image gradient vector field. Similar terms can be found in previous works dedicated to the extraction or filtering of elongated structures [11, 12] or in the minimization scheme of geodesic active contours [13].




Fig. 2. Learned distributions of the dynamic priors. Left: radius variations. Right: direction variations.

The term $E(r_t^{(i,j)})$ corresponds to the response of the morphological erosion by the sphere $S_t^{(i,j)}$ at point $x_t^{(i)}$:

$$E(r_t^{(i,j)}) = \min_z \left\{ I(z) \mid |z - x_t^{(i)}| \leq r_t^{(i,j)} \right\}$$

where $I(z)$ is the image intensity at location z . This region-based term favors bright spheres and becomes especially useful for small vessels, when the computation of the gradient flow over small discrete spheres becomes unreliable.

The term $G(x_{t+1}^{(i,j)})$ is defined as:

$$G(x_{t+1}^{(i,j)}) = I(x_{t+1}^{(i,j)}) - E(r_t^{(i,j)})$$

This difference term counter-balances the shrinkage effect introduced by $E(r_t^{(i,j)})$ by favoring points with high intensity and low erosion response, *i.e.*, intuitively, locations far from $x_t^{(i)}$. We stress that this term depends on the individual location of $x_{t+1}^{(i,j)}$, not just on the sphere it lies on. It discriminates against locations along the border of the vessel and favors hyper-intense locations toward the centerline. The effects of each term are illustrated in Figure 3.

In addition, we constrained and normalized image intensities within a range of 0 to 300 Hounsfield units (typical intensities of contrast-enhanced blood pool in CT acquisitions), in order to reduce the influence of hyper-intense structures such as calcifications.

3.2. Dynamic Prior: $p(x_{t+1}^{(i,j)}|x_t^{(i)})$

Vessel dynamics are described in terms of radius and direction variations, considered as independent variables:

$$p(x_{t+1}^{(i,j)}|x_t^{(i)}) = p(r_t^{(i,j)}|r_{t-1}^{(i)}) p(\vec{d}_t^{(i,j)}|\vec{d}_{t-1}^{(i)})$$

Both distributions were learned from a pool of 19 manually segmented cardiac CT datasets. Radius variation was evaluated as the ratio $\frac{r_{t+1}}{r_t}$ and direction variation as the angle $\arccos\left(\frac{|\vec{d}_{t+1}|}{|\vec{d}_t|}\right)$. Samples from the manually segmented database allowed us to build detailed histograms of these distributions (see Figure 2). These non-parametric forms are used directly to evaluate $p(r_t^{(i,j)}|r_{t-1}^{(i)})$ and $p(\vec{d}_t^{(i,j)}|\vec{d}_{t-1}^{(i)})$.

4. EVALUATION AND RESULTS

Our method was evaluated on 19 cardiac CT datasets, for both left and right main coronary branches. Ground-truth segmentations were obtained by detailed manual delineation of the centerlines and associated cross-sections. For each dataset to be tested, dynamic priors




Fig. 3. Illustration of the likelihood terms.

(a) original image and sample position $x_t^{(i)}$; (b) gradient flow term $F(r_t^{(i,j)})$ for all potential successors of $x_t^{(i)}$ overlaid in red; (c) erosion term $E(r_t^{(i,j)})$; (d) “growing” term $G(x_{t+1}^{(i,j)})$; (e) likelihood image, conjunctive combination of the terms (b), (c) and (d).

were learned on the remaining 18 cases (Leave-One-Out rule) and quantitative results were averaged on 10 runs.

A result example is given in Figure 4. Besides the sample populations $\{\{x_t^{(i)}, w_t^{(i)}\}_{i=1}^{N_t}\}_{t=0}^T$, a key output of the algorithm is the maximum *a posteriori* (MAP) chain $\{x_0^*, \dots, x_T^*\}$. It is directly extracted by selecting x_T^* as the sample with highest weight at time-step T and by backtracking its ancestors. From this chain, an approximate *centerline* can be derived by interpolating a curve linking the spheres centers. Similarly, a volumic representation can be obtained thanks to the associated radius information. We emphasize the conciseness and usability of this MAP result, allowing, for instance, the generation of curved planar reformation views (see Figure 5(b)).

The algorithm is controlled by a threshold on the length of the extracted branch (9cm for coronaries) and a low threshold on the gradient flow. These constraints were sufficient to prevent “leaking” problems in all our experiments. Remaining parameters are the initial conditions (original seed point at the aortic root), the maximum vessel radius r_{max} (fixed to 6mm for coronaries) and the number of samples selected at each iteration (we used $N_t = 100, \forall t$).

Table 1 summarizes quantitative results for the 19 datasets. To evaluate the overall robustness of the algorithm, we compared the MAP results to the manual segmentations through the similarity index proposed in [14]: $SI = \frac{2TP}{2TP+FN+FP}$, where TP corresponds to true positive portions (matching segments), FP to false positives and FN to false negatives. We obtained a very high mean similarity factor of $94.8\% \pm 6.2\%$. Early stop was observed in one stenotic case. False positive segments were encountered at the distal end of some branches and directly depend on the stopping criterion.

In terms of accuracy, results for centerline positioning (see also Figure 4(c)) and radius estimation were satisfactorily in the order of the data resolution (intra-slice resolution of $0.33 \times 0.33\text{mm}$, interslice resolution from 0.33 to 0.6 mm). We stress that we focus more on robust tracking than accurate surface segmentation. Our algorithm output can serve as a close initialization for subsequent refinements, simplifying such tasks through the availability of the centerline and radius information.

Additionally, we observed that nearly optimal tracking performance could be obtained with a very limited number of samples ($N_t \simeq 20$). Increasing the number of samples improves the accuracy of the centerline and radius estimations. Nearly maximal accuracy is reached for $N_t = 100$ samples in all cases. Our original approach, which builds the *entire* distribution of potential successors prior to sampling, potentially explains the very low number of samples needed to attain acceptable tracking performance. This compares very favorably with similar particle-based works such as [4] and [5], which report the use of 1000 and 500 samples, respectively. Results also appeared to be satisfactorily stable to additive Gaussian noise. Until a noise standard of ~ 100 Hounsfield units




Fig. 4. Result sample.

(a) 3D rendering of the heart with the extracted coronary branch; (b) detail of the samples used by the filter, with the maximum *a posteriori* (MAP) chain (dark blue curve); (c) detail of the MAP centerline with the external envelope of the associated spheres, overlaid with the manually segmented centerlines (dark blue curves).

Measure	Mean	Standard Deviation
similarity index	0.948	0.062
centerline position error	0.4mm	0.33mm
radius estimation error	0.33mm	0.21mm

Table 1. Quantitative evaluation results (see text).

(H.U.), which can be considered as a reasonably high noise level by cardiac CT standards, similarity and accuracy measures are nearly unaffected. Then, increasing noise first impacts centerline and radius estimation accuracy. With a noise standard deviation superior to 200 H.U., the geometric information of the image is significantly altered and similarity performance deteriorates as the algorithm fails to track thin vessels disconnected by such a high noise level.

Overall, our Bayesian, multi-hypotheses framework ensures robustness by balancing prior knowledge and information extracted from the image. It proved to be robust to local anatomical anomalies such as aneurysms, stenoses and stents, as illustrated in Figure 5.

Computational efficiency is promising, especially given the high number of hypotheses evaluated by our sampling scheme. On average, the extraction of a main coronary branch is performed in less than 2 minutes on a Pentium 4 3Ghz computer. The computational time is directly proportional to the number of samples, providing the user with some degree of scalability between computational efficiency and centerline and radius accuracy.




Fig. 5. Result samples on data with local anomalies (indicated by arrows). (a) result in presence of a stenosis, an aneurysm and a stent; (b) associated curved planar reformation (CPR) view; (c) result with a stent, high curvature and thin distal end; (d) result with neighboring calcified branch and a stent; (e), (f) and (g) multi-planar reformation views of different results in presence of stents and calcifications.




Fig. 6. Preliminary result on peripheral CT data. Segmentation of the descending aorta and right iliac artery.

5. CONCLUSION AND PERSPECTIVES

In this paper, we presented a new tracking algorithm for the segmentation of vascular structures from medical images. Our Bayesian stochastic filter is inspired by recent developments in particle filtering for image segmentation. Novel aspects of our method include a constrained, medial-based geometric model, the use of prior dynamic knowledge learned from a ground-truth database and a new sampling scheme able to cover hypothesis distributions broader than classical particle filters while remaining computationally efficient. Validation of the method for 3D CT coronary artery extraction exhibits very promising results in terms of robustness and accuracy.

We are currently working on a multi-branch version of the algorithm, exploring original solutions to cope with classical problems of loss of diversity linked to importance sampling [10]. Other refinements include more elaborate dynamics and better stopping criteria. Last but not least, we emphasize the generality of our approach. By simply changing the maximum vessel radius r_{max} , we obtained promising preliminary results on peripheral CT data, as illustrated in Figure 6. More generally, our method can be directly applied in 2D and the geometric model is not specific to vascular structures. One could imagine, for instance, to adapt our algorithm for road extraction from 2D satellite images.

6. REFERENCES

- [1] C. Kirbas and F. Quek, "A review of vessel extraction techniques and algorithms," *ACM Computing Surveys*, vol. 36, no. 2, pp. 81–121, 2004.
- [2] T. Deschamps and L.D. Cohen, "Minimal paths in 3D images and application to virtual endoscopy," in *Proc. Sixth European Conference on Computer Vision*, 2000.
- [3] H. Li and A. Yezzi, "Vessels as 4D curves: Global minimal 4D paths to extract 3D tubular surfaces," in *Proc. IEEE Conference on Computer Vision and Pattern Recognition*, 2006, p. 82.
- [4] C. Florin, N. Paragios, and J. Williams, "Particle filters, a quasi-Monte Carlo solution for segmentation of coronaries.," in *Proc. Medical Image Computing and Computer Assisted Intervention*. 2005, pp. 246–253, Springer.
- [5] M. Schaap, I. Smal, C. Metz, T. van Walsum, and W. Niessen, "Bayesian tracking of elongated structures in 3D images," in *Proc. Information Processing in Medical Imaging*. 2007, vol. 4584, pp. 74–85, Springer.
- [6] H. Blum, "A transformation for extracting new descriptors of shape," *Models for the Perception of Speech and Visual Form*, pp. 362–380, 1967.
- [7] O. Wink, W. J. Niessen, and M. A. Viergever, "Fast delineation and visualization of vessels in 3D angiographic images.," *IEEE Trans. on Medical Imaging*, vol. 19, no. 4, pp. 337–346, 2000.
- [8] S. Aylward and E. Bullitt, "Initialization, noise, singularities, and scale in height ridge traversal for tubular object centerline extraction," *IEEE Trans. on Medical Imaging*, vol. 21, pp. 61–75, 2002.
- [9] Y. Fridman, S. M. Pizer, S. R. Aylward, and E. Bullitt, "Segmenting 3D branching tubular structures using cores.," in *Proc. Medical Image Computing and Computer Assisted Intervention*. 2003, pp. 570–577, Springer.
- [10] A. Doucet, N. de Freitas, and N. Gordon, *Sequential Monte Carlo Methods in Practice*, Springer-Verlag, 2001.
- [11] A. Vasilevskiy and K. Siddiqi, "Flux maximizing geometric flows.," *IEEE Trans. on Pattern Analysis and Machine Intelligence*, vol. 24, no. 12, pp. 1565–1578, 2002.
- [12] M. Holtzman-Gazit, R. Kimmel, N. Peled, and D. Goldsher, "Segmentation of thin structures in volumetric medical images," *IEEE Trans. on Image Processing*, vol. 5, pp. 354–363, 2006.
- [13] V. Caselles, R. Kimmel, and G. Sapiro, "Geodesic active contours," *International Journal of Computer Vision*, vol. 22, no. 1, pp. 61–79, 1997.
- [14] C. Metz, M. Schaap, A. Van Der Giessen, T. Van Walsum, and W. Niessen, "Semi-automatic coronary artery centerline extraction in computed tomography angiography data.," in *Proc. 4th IEEE International Symposium on Biomedical Imaging*, 2007, pp. 856–859.