Bayesian Maximal Paths for Coronary Artery Segmentation from 3D CT Angiograms

David Lesage^{1,2}, Elsa D. Angelini², Isabelle Bloch², and Gareth Funka-Lea¹

¹ Siemens Corporate Research, Imaging and Visualization dept., Princeton NJ, USA

² Institut Telecom, Telecom ParisTech, CNRS LTCI, Paris, France

Abstract. We propose a recursive Bayesian model for the delineation of coronary arteries from 3D CT angiograms (cardiac CTA) and discuss the use of discrete minimal path techniques as an efficient optimization scheme for the propagation of model realizations on a discrete graph. Design issues such as the definition of a suitable accumulative metric are analyzed in the context of our probabilistic formulation.

Our approach jointly optimizes the vascular centerline and associated radius on a 4D space+scale graph. It employs a simple heuristic scheme to dynamically limit scale-space exploration for increased computational performance. It incorporates prior knowledge on radius variations and derives the local data likelihood from a multiscale, oriented gradient flux-based feature. From minimal cost path techniques, it inherits practical properties such as computational efficiency and workflow versatility. We quantitatively evaluated a two-point interactive implementation on a large and varied cardiac CTA database. Additionally, results from the Rotterdam Coronary Artery Algorithm Evaluation Framework are provided for comparison with existing techniques. The scores obtained are excellent (97.5% average overlap with ground truth delineated by experts) and demonstrate the high potential of the method in terms of robustness to anomalies and poor image quality.

1 Introduction

In many biomedical applications, the segmentation of vascular structures is an important step towards diagnosis, treatment and surgery planning. Modern 3D angiographic modalities produce increasingly large and detailed acquisitions. This exacerbates the need for automated or semi-automated methods to reduce the burden of manual delineation while increasing repeatability.

Among the rich literature on vascular segmentation [1], *minimal path* methods are particularly popular. They classically extract a vessel as a path between two points on a regular lattice. The cost of a path is defined cumulatively by a vessel-dedicated metric and optimized over the lattice by graph-based methods (*e.g.*, Dijkstra-based and fast-marching algorithms for L_1 and L_2 metrics, respectively). Recent works [2,3] have shown that the lattice is not necessarily limited to the spatial positions of the vessel centerline and can be augmented with dimensions such as the vessel radius. Key properties of minimal paths are

their guarantee of global optimality, their computational efficiency and the control over boundary conditions. One can always find a path between two given points or even search exhaustively, *i.e.*, extract the paths from a single seed to *all* the other points. Most drawbacks, such as so-called “shortcut” issues, come from the practical difficulty of designing well-behaved cumulative metrics.

Another particularly active class of methods is the family of multi-hypotheses tracking methods [4,5,6,7]. They increase the robustness of the local tracking process by evolving several hypotheses in parallel. They differ on how hypotheses are selected and evolved. Bayesian formulations such as particle filters [5] and related schemes [6,7] have demonstrated their robustness and design versatility, in particular in the integration of prior model-based knowledge. Unlike minimal path techniques, they do not offer control over the exhaustiveness of the search, but are designed to explore large search spaces in a sparse manner, focusing on most promising areas. They are generally not dependent on the discretization of the search space but can be computationally expensive.

In this paper, we introduce a recursive Bayesian model related to those underlying multi-hypothesis probabilistic methods [5,6,7]. We discuss the use of minimal path techniques as an efficient optimization scheme propagating model realizations as paths on a discrete graph. We detail key design aspects such as the definition of the data likelihood and the derivation of a cumulative metric from the Bayesian formulation. Our approach was applied to the particularly challenging task of segmenting coronary arteries from 3D cardiac CTA. Qualitative and quantitative evaluation is given on clinical data.

2 Geometric and Bayesian Models

A vascular segment is modeled as a discrete series of states $X_{t \in [0:T]} = \{(p_t, r_t)\}$, composed of centerline points p_t and associated radii r_t (Fig. 1). Cross-sections are assumed to be circular, a reasonable approximation for thin vessels such as coronaries. Their local orientation is approximated by $d_t = \frac{p_t - p_{t-1}}{|p_t - p_{t-1}|}$.

We are interested in $X^* = \arg \max_X P(X)P(Y|X)$, the *maximum a posteriori* model realization given image observations $Y = \{Y^j\}$. An observation Y^j is obtained as the response of a multiscale oriented feature, computed at location p^j , radius r^j and direction d^j (Sec. 4). The likelihood $P(Y^j|X = X_{[0:t]}^i)$ depends on whether an observation Y^j is associated to the model realization $X_{[0:t]}^i$ or not. A particularity of our model is to distinguish two distributions, $P_v(Y^j)$ and

Fig. 1. Geometric model and parameterization

$P_{bg}(Y^j)$ for responses in vessels and in the background, respectively (similarly to [8] for road tracking). Assuming conditionally independent observations given a model realization, the joint likelihood of the observations is:

$$P(Y|X_{[0:t]}^i) = \prod_{Y^j \in X_{[0:t]}^i} P_v(Y^j) \prod_{Y^j \notin X_{[0:t]}^i} P_{bg}(Y^j) = \prod_{Y^j \in X_{[0:t]}^i} \frac{P_v(Y^j)}{P_{bg}(Y^j)} \prod_{Y^j \notin X_{[0:t]}^i} P_{bg}(Y^j)$$

where the (abusive) notation $Y^j \in X_{[0:t]}^i$ indicates that Y^j is associated to a certain state of $X_{[0:t]}^i$. We note $Y^j = Y_t^i$ if Y^j is associated with the particular state X_t^i . Assuming a 1st order Markovian prior and omitting the terms independent of the model realization, the recursive update of the *a posteriori* probability of a model realization (also referred to as an *hypothesis*) is given by Bayes' rule:

$$P(X_{[0:t]}^i | Y_{[0:t]}^i) \propto P(X_{[0:t-1]}^i | Y_{[0:t-1]}^i) P(X_t^i | X_{t-1}^i) \frac{P_v(Y_t^i)}{P_{bg}(Y_t^i)} \quad (1)$$

This Bayesian model is similar to those used in recent works on vascular probabilistic tracking [5,6,7]. The main difference is the integration of background information in the likelihood expression, similarly to [8].

3 Graph Optimization for Hypothesis Propagation

The theoretical maximum *a posteriori* problem is generally intractable as the search space grows exponentially with the length T of the model. In practice, probabilistic multi-hypothesis methods adopt *selection* schemes to limit the number of hypotheses, *e.g.*, stochastic resampling [5,7] or deterministic pruning [6].

As an alternative, we rely on the discretization of the search space, allowing the use of an efficient and robust minimal path-like optimization scheme. Spatial positions $\{p_t^i\}$ are discretized as the regular image grid and we use R different radius values (see Sec. 4). Our approach thus explores a 4D space+radius graph, where each node corresponds to a possible state $X^k = (p^i, r^j)$. To value the edges of the graph, we define our additive cost metric by noting that maximizing (1) is equivalent to minimizing its negative logarithm:

$$C(X_{[0:t]}^i) = C(X_{[0:t-1]}^i) - \log(P(X_t^i | X_{t-1}^i)) - \log(P_v(Y_t^i)) + \log(P_{bg}(Y_t^i)) - M \quad (2)$$

where $C(X_{[0:t]}^i) = -\log(P(X_{[0:t]}^i | Y_{[0:t]}^i))$ and M is a constant ensuring that $C(X_{[0:t]}^i)$ remains positive or null¹. This additive cost metric is directly suitable for Dijkstra-like optimization. Our algorithm sorts the hypotheses in a min-heap structure. When the heap root $X_t^i = (p_t^i, r_t^i)$ is popped, we consider its neighbor states, defined by the product set of the 26-neighboring positions $\{p_{t+1}^i\}$ of p_t^i and all possible scales. The cost of each neighbor is updated according to (2) (details in Sec. 4). Propagation continues until the heap is empty (exploration

¹ We set $M = \log(P_{bg}^{min})$, with $P_{bg}^{min} = 10^{-6}$ in our implementation.

of the entire grid) or when a given end position is popped (two-seed workflow). The result path (points+radiuses) is simply backtracked from the end state.

Exploring the full 4D graph can be very costly in time and memory. To alleviate that cost, we propose to propagate only the $H < R$ best hypotheses (of different radiuses) first reaching each spatial position. H controls the local, dynamic radius selection scheme limiting scale-space exploration. For $H = R$, the optimization would be globally optimal. With $H = 1$, only the locally best radius is propagated. Using $H > 1$ increases robustness to scale-related ambiguities.

Our approach can be viewed as a generalization of [9]. We also derive edge costs from an oriented, multiscale medialness feature, but our method uses the feature response indirectly to value the likelihood terms, incorporates radiuses in the optimization and controls the coherence of their variations through the prior term (Sec. 4). From its space+scale minimal path approach, our method is closely related to [3]. Where our Bayesian model yields an edge-based L_1 metric, authors in [3] prefer L_2 optimization, arguably less sensitive to discretization. We believe however that our metric is less parameter-dependent than the node-based potentials proposed in [3]. More importantly, our algorithm does not explore the 4D graph entirely and rather focuses on most promising radiuses, yielding considerable memory and time gains in practice. For $H < R$, our optimization scheme is heuristic and does not guarantee the extraction of the global maximum, but it is computationally efficient, based on a sound theoretical framework and particularly robust, as confirmed by the validation results from Sec. 5.

4 Implementation Details

Our implementation relies on an interactive workflow with user-provided start and end points. The heap is initialized with the set of seed states $X_0^i = (p_0, r^i)$, *i.e.* all possible radiuses r^i for the start position p_0 . To cover the typical radius range of coronary arteries, we use $R=16$ radius values from $r_{\min}=0.3$ to $r_{\max}=5.1\text{mm}$ with a fixed step of 0.3mm .

The prior term $P(X_t^i|X_{t-1}^i)$ in (2) constrains radius variations:

$$P(X_t^i|X_{t-1}^i) = P(r_t^i|r_{t-1}^i) = \begin{cases} (\sigma_- \sqrt{2\pi}) \mathcal{N}(r_t^i|r_{t-1}^i, \sigma_-^2) & \text{if } r_t^i \leq r_{t-1}^i, \\ (\sigma_+ \sqrt{2\pi}) \mathcal{N}(r_t^i|r_{t-1}^i, \sigma_+^2) & \text{if } r_t^i > r_{t-1}^i. \end{cases} \quad (3)$$

where $\mathcal{N}(\cdot|\mu, \sigma^2)$ is the Normal distribution of mean μ and variance σ^2 . We use an asymmetric formulation to penalize widening models more than shrinking models. We experimentally set $\sigma_- = 0.4\text{mm}$ and $\sigma_+ = 0.2\text{mm}$. These distributions are kept unnormalized, as discussed at the end of the section.

Observations Y_t^i are the responses of a multiscale oriented medialness feature:

$$Y_t^i = \text{MFlux}(p_t^i, r_t^i, d_t^i) = \frac{2}{N} \sum_{k=1}^{\frac{N}{2}} \min(\langle \nabla I(x_k), u_k \rangle, \langle \nabla I(x_k^\pi), u_k^\pi \rangle) \quad (4)$$

At cross-sectional contour points x_k , **MFlux**[10] accumulates the projections of the image gradient vectors $\nabla I(x_k)$ on inward radial directions $u_k = p_t^i - x_k$. The

Fig. 2. Flux-based feature and likelihood distributions. (a) Detail of a dataset with arteries running close to a heart chamber. (b) Response of the local cross-sectional gradient flux (maximum over scales and directions for each voxel). (c) Response of **MFlux**. (d) Examples of P_v and P_{bg} likelihood distributions.

sum of these projections corresponds to a measure of the local inward gradient flux through the cross-section, which is maximal when the model cross-section is well aligned with a vessel in the image. Instead of taking the sum, **MFlux** retains the minimum contribution between diametrically opposed points x_k and x_k^π (see Fig. 1). This non-linear combination dramatically reduces false positive step-edge responses caused by asymmetric flux contributions at locations such as the surface of heart chambers (Fig. 2). **MFlux** enjoys a high discriminative power over scales and positions while remaining computationally efficient. we use $N=8$ cross-section points and precompute gradient vectors using Gaussian derivatives with $\sigma_g=0.3\text{mm}$. Please refer to [10] for more in-depth discussion.

Observations are used to value the likelihood terms $P_v(Y_t^i)$ and $P_{bg}(Y_t^i)$ in (2). The vessel likelihood P_v was learnt as the response histogram on 10 datasets with ground truth segmentation (not included in the validation set of Sec. 5). The background likelihood P_{bg} is dataset-dependent and is estimated through feature responses at 10^6 randomly sampled parameters (p, r, d) . It thus encodes general information about the reliability of the data at hand, as it will vary as a function of the image quality and noise. P_v and P_{bg} are generally well separated thanks to the discriminative power of **MFlux**, as depicted in Fig. 2 (d).

We finally discuss the normalization of the distributions in (2), which can be considered as an implicit parameterization of our metric. Although (2) is defined up to an additive constant, the scaling of the distributions greatly influences how model realizations are compared. In particular, our algorithm deals with hypotheses of different lengths competing for the same position. Instead of classically normalizing each distribution so that it sums to unity, we scale them so that their respective maximum is always 1. Consequently, an ideal model realization will have a cumulative cost (log-probability) of 0 regardless of its length. By limiting length-related penalization to its minimum, we effectively reduce the risk of “shortcuts”, a classical issue of minimal path techniques.

5 Experiments and Validation

We first discuss some of the distinctive components of our approach. In general, the background likelihood $P_{bg}(Y_t^i)$ has a limited impact on the extracted result, but it markedly decreases the extent of exploration needed to reach the end state

Fig. 3. Effect of the background likelihood and scale prior terms. Left and middle: search space explored to extract the red vessel. Highlighted in red: *with* background likelihood P_{bg} (1.1×10^6 voxels visited). In dark: *without* P_{bg} (1.95×10^6 voxels). Right: difficult case extracted *without* scale prior (top) and *with* scale prior (bottom).

Table 1. Summary of quantitative results (see text for details)

(a) Internal validation

Measure	H=1	H=2	H=4
OV (avg.)	97.8%	99.1%	99.3%
AD (avg.)	0.375mm	0.365mm	0.365mm
AR (avg.)	0.195mm	0.192mm	0.191mm

(b) Results from the Rotterdam challenge [11]

Measure	% / mm			score		
	min.	max.	avg.	min.	max.	avg.
OV	76.0%	100.0%	97.5%	39.7	100.0	81.5
AD	0.20mm	3.78mm	0.37mm	22.4	55.9	36.3

Fig. 4. Result samples (extracted centerlines and mask from local radius estimation). (a) Full arterial tree extracted with one ostia seed and 9 distal seeds (in 57 sec. with $H = 4$). (b) 3D and multi-planar reformation (MPR) views of a curvy, small secondary vessel branching off the right coronary artery. (c) 3D and MPR views of a calcified left coronary artery. (d) Low image quality and occluded right coronary artery.

(up to $\sim 50\%$, see Fig. 3). By penalizing hypotheses in non-vascular areas, it yields significant speed gains. The scale prior $P(r_t^i | r_{t-1}^i)$ improves the robustness of the extraction in presence of local anomalies and/or low image quality (see Fig. 3, right) while still allowing smooth scale adaptations (Fig. 4 (b)).

Our algorithm was first validated on 51 CTA datasets of varying image quality and pathologies (Fig. 4). This database counts 858 coronary branches manually delineated by experts, including radius estimation. Table 1(a) shows quantitative

results for different values of H . OV is the average overlap with ground truth, *i.e.*, the portion of ground truth and extracted points coinciding within radius distance. The version with $H=1$ already obtains a high overlap score of 97.8%, improved up to 99.3% for $H=4$. These results demonstrate the high robustness of our approach and the relevance of keeping several hypotheses per spatial location. Increasing H over 4 did not bring significant improvement in our tests. AD , the average distance of the result centerline to the ground truth, is slightly larger than the typical intra-slice data resolution (0.33mm per voxel). This accuracy result is coherent with the discrete nature of our algorithm. The average radius estimation error AR of 0.191mm is satisfyingly subvoxelic.

We additionally submitted results (for $H = 4$) to the publicly open Rotterdam Coronary Artery Algorithm Evaluation Framework [11]. In Table 1(b), we report the average overlap OV and average distance AD results for 24 testing datasets (96 vessels). Other statistics are left out due to space restrictions [12]. Scoring in $[0, 100]$ is defined as follows: 50 for a result of the order of the inter-observer variability, 100 for a perfect result. The good performance of our method is confirmed, with 97.5% average overlap, exceeding inter-observer overlap (score of 81.5). The average distance AD is slightly larger than inter-observer variability (score of $36.3 < 50$). Our primary focus being robustness, we consider this accuracy level to be satisfactory for initial delineation before subsequent refinement. This evaluation also allows the direct comparison with other existing methods. For instance, our algorithm brings a noticeable robustness improvement (97.5% versus 91.9% OV) over the method from [13] which implements a classical two-seed minimal path technique based on image intensity and a Hessian-based vesselness feature. To date, the only publicly ranked method to outperform our approach in terms of robustness is [14], an adaptation of [4] supplemented with minimal paths, which obtained $OV = 98.5\%$. It is worth noticing that this method required more interaction, with the use of intermediate seed points, where we strictly limited ourselves to the provided start and end points.

Finally, we emphasize the computational efficiency of our approach. The effort required to extract a vessel depends on several factors, such as the image quality, the length and complexity of the target path and overall vascular network. In order to extract the longest vessel branch, the algorithm will basically explore the entire network (see Fig. 3). Consequently, the entire arterial tree can be extracted for the same computational effort as its longest branch, by specifying one seed at the ostium and one distal point per branch (see Fig. 4). With our C++ implementation, the average vessel branch is extracted in about one minute for $H = 4$ (less than 15 sec. for $H = 1$) on a 2.16GHz Core Duo CPU.

6 Conclusion and Perspectives

In this paper, we have presented a new algorithm for the segmentation of coronary arteries from 3D cardiac CTA data. Our approach relies on a recursive Bayesian model, from which it inherits its robustness, and is optimized by minimal path-like techniques, from which it inherits its computational efficiency and

workflow flexibility. In particular, our algorithm improves on classical minimal path techniques by providing a sound theoretical framework for the definition of the cumulative cost metric, whose different components were carefully studied. The high practical robustness of our technique was demonstrated through an extensive quantitative validation on clinical data.

The very promising results obtained by our proof-of-concept implementation open several high-potential perspectives. One lead is to exploit the reliability of the approach to improve overall clinical workflow for coronary disease assessment through fast, intuitive and reliable interactive tools to extend and correct centerlines for the most difficult cases. Another lead is further automation for full tree extraction. The main difficulty in this case is the design of robust stopping criteria for the propagation. We finally highlight the generality of the approach. Within the same framework, we are currently evaluating other vascular applications by adapting components such as the prior and likelihood distributions.

References

1. Kirbas, C., Quek, F.: A review of vessel extraction techniques and algorithms. *ACM Computing Surveys* 36(2), 81–121 (2004)
2. Wink, O., Niessen, W., Viergever, M.: Multiscale vessel tracking. *IEEE Trans. Med. Imaging* 23(1), 130–133 (2004)
3. Li, H., Yezzi, A.: Vessels as 4-D curves: Global minimal 4-D paths to extract 3-D tubular surfaces and centerlines. *IEEE Trans. Med. Imaging* 26, 1213–1223 (2007)
4. Friman, O., Hindennach, M., Peitgen, H.O.: Template-based multiple hypotheses tracking of small vessels. In: *Proc. Int. Symp. Biom. Imaging*, pp. 1047–1050 (2008)
5. Florin, C., Paragios, N., Williams, J.: Particle filters, a quasi-Monte Carlo solution for segmentation of coronaries. *Med. Image Comput. Assist. Interv.*, 246–253 (2005)
6. Schaap, M., Smal, I., Metz, C., van Walsum, T., Niessen, W.: Bayesian tracking of elongated structures in 3D images. In: *Inf. Process. Med. Imaging*, pp. 74–85 (2007)
7. Lesage, D., Angelini, E.D., Bloch, I., Funka-Lea, G.: Medial-based bayesian tracking for vascular segmentation: Application to coronary arteries in 3D CT angiography. In: *Proc. Int. Symp. Biom. Imaging*, pp. 268–271 (2008)
8. Geman, D., Jedynak, B.: An active testing model for tracking roads in satellite images. *IEEE Trans. Pattern Anal. Mach. Intell.* 18(1), 1–14 (1996)
9. Gulsun, M.A., Tek, H.: Robust vessel tree modeling. In: *Med. Image Comput. Assist. Interv.*, pp. 602–611 (2008)
10. Lesage, D., Angelini, E.D., Bloch, I., Funka-Lea, G.: Design and study of flux-based features for 3D vascular tracking. In: *Proc. Int. Symp. Biom. Imaging* (2009)
11. Rotterdam Coronary Artery Algorithm Evaluation Framework, <http://coronary.bigr.nl/>
12. Metz, C., Schaap, M., van Walsum, T., van Der Giessen, A., Weustink, A., Mollet, N., Krestin, G., Niessen, W.: 3D segmentation in the clinic: A grand challenge II - coronary artery tracking. *The MIDAS Journal* (2008)
13. Metz, C., Schaap, M., van Walsum, T., Niessen, W.: Two point minimum cost path approach for CTA coronary centerline extraction. *The MIDAS Journal* (2008)
14. Friman, O., Kühnel, C., Peitgen, H.: Coronary artery centerline extraction using multiple hypothesis tracking and minimal paths. *The MIDAS Journal* (2008)