Sequential Monte Carlo Tracking for Marginal Artery Segmentation on CT Angiography by Multiple Cue Fusion

Shijun Wang, Brandon Peplinski, Le Lu, Weidong Zhang, Jianfei Liu,
Zhuoshi Wei, and Ronald M. Summers

Imaging Biomarkers and Computer-Aided Diagnosis Laboratory,
Radiology and Imaging Sciences, National Institutes of Health Clinical Center,
Bethesda, MD, 20892-1182, U.S.
rms@nih.gov

Abstract. In this work we formulate vessel segmentation on contrast-enhanced CT angiogram images as a Bayesian tracking problem. To obtain posterior probability estimation of vessel location, we employ sequential Monte Carlo tracking and propose a new vessel segmentation method by fusing multiple cues extracted from CT images. These cues include intensity, vesselness, organ detection, and bridge information for poorly enhanced segments from global path minimization. By fusing local and global information for vessel tracking, we achieved high accuracy and robustness, with significantly improved precision compared to a traditional segmentation method (p=0.0002). Our method was applied to the segmentation of the marginal artery of the colon, a small bore vessel of potential importance for colon segmentation and CT colonography. Experimental results indicate the effectiveness of the proposed method.

Keywords: Sequential Monte Carlo tracking, multiple cues, particle filtering, marginal artery, CT angiography.

1 Introduction

The marginal artery is a small blood vessel located within the abdominal mesentery which travels parallel to the colon, and communicates between the inferior and superior mesenteric arteries (IMA and SMA). Segmentation of the marginal artery can improve supine-prone colonic polyp registration and help connect collapsed colonic segments in CT colonography (CTC). The purpose of this pilot study is to automatically detect the marginal artery on high-resolution abdominal CT angiograms (CTA) using a sequential Monte Carlo (SMC) tracking method.

Vessel enhancement filtering, region-growing, active contours, centerline extraction, and stochastic framework are five major approaches to 3D vessel segmentation [1]. Among these methods, SMC tracking, or particle filtering (PF), has been widely used for its accuracy, robustness, and computational feasibility. Florin et al. proposed a PF-based approach for segmentation of coronary arteries [2]. In their model, state

variables include position, orientation, shape, and vessel appearance. Later, Schaap et al. presented a Bayesian tracking framework for tubular structures such as vessels [3]. The key contribution of their work is a novel observation model designed for tube-like objects. Lacoste et al. employed Markov marked point processes for segmentation of coronary arteries on 2D angiograms [4]. More recently, Friman proposed a multiple hypothesis template tracking scheme for small 3D vessel structures [5].

SMC has been used in computer vision to handle athlete or vehicle tracking in video sequences. The utilization of more target and background information will typically improve accuracy and robustness for a given noise level. In recent years, incorporating multiple cues in the Bayesian tracking framework has been a major research direction. Wu and Huang proposed a factorized graphical model to integrate multiple cues for Bayesian tracking [6]. Brasnett et al. proposed visual cues including color, edge, and texture for object tracking in video sequences [7]. The work of Moreno-Noguer et al. [8] focused on integrating dependent multiple cues for robust tracking.

In this work, we propose a new Bayesian vessel segmentation method by fusing multiple cues extracted from CT images to automatically detect the marginal artery on high-resolution abdominal CTA. The remainder of this paper is organized as follows: in Sec. 2 we introduce our SMC multiple cue tracking framework; in Sec. 3 we show experimental results on a 7 patient CTA dataset. We conclude our findings in Sec. 4 with a short discussion.

2   Sequential Monte Carlo Tracking by Multiple Cue Fusion

2.1   Bayesian Tracking Framework

First we will introduce the SMC tracking framework and notation. Observations $\{\mathbf{y}_t ; t \in N\}$, $\mathbf{y}_t \in R^{m_y}$ are typically captured in a sequential order. Each observation has an associated hidden variable $\{\mathbf{x}_t ; t \in N\}$, $\mathbf{x}_t \in R^{m_x}$ which generally corresponds to the location of the target and speed at time point t. For each t, the observation $\mathbf{y}_t$ is only conditionally dependent on x_t, i.e. $p(\mathbf{y}_t | \mathbf{y}_{1:t-1}, \mathbf{x}_{1:t}) = p(\mathbf{y}_t | \mathbf{x}_t)$, where $\mathbf{y}_{1:t-1}$ represents all observations from time point 1 to time point t-1 and $\mathbf{x}_t$ represents all hidden variables from time point 1 to time point t. We also assume that the time sequence $\mathbf{x}_t$, t=1,2,…T has a Markov property of order one: $p(\mathbf{x}_t | \mathbf{x}_{1:t-1}) = p(\mathbf{x}_t | \mathbf{x}_{t-1})$. The dynamics of the Markov chain can be described by the following two steps:

1) Prediction step:

$$p(\mathbf{x}_t | \mathbf{y}_{1:t-1}) = \int p(\mathbf{x}_t | x_{t-1}) p(\mathbf{x}_{t-1} | \mathbf{y}_{1:t-1}) d\mathbf{x}_{t-1} \tag{1}$$

2) Update step:

$$p(\mathbf{x}_t | \mathbf{y}_{1:t}) = \frac{p(\mathbf{y}_t | \mathbf{x}_t) p(\mathbf{x}_t | \mathbf{y}_{1:t-1})}{p(\mathbf{y}_{1:t})} \tag{2}$$

In our implementation, the state variable **x** was composed by $\mathbf{x} = \left(x, x', y, y', z, z' \right)$, corresponding to the current location and moving speed of the vessel during the dynamic tracking process.

2.2   Mixture Dynamic Model Combined with Vector Field

Considering that the majority of vessel segments are smooth in 3D space and exhibit a tube structure, we chose a constant velocity model to capture translational motion:

$$\mathbf{x}_{t+1} = \mathbf{F}\mathbf{x}_t + \mathbf{d}_t, \ \ \mathbf{d}_t \sim N\left(0, \Sigma_d\right), \tag{3}$$

where matrix **F** controls the speed at which the target (vessel segmentation) can proceed during the tracking process. $\mathbf{d}_t$ follows a zero-mean Gaussian distribution whose covariance matrix was determined empirically based on the training set.
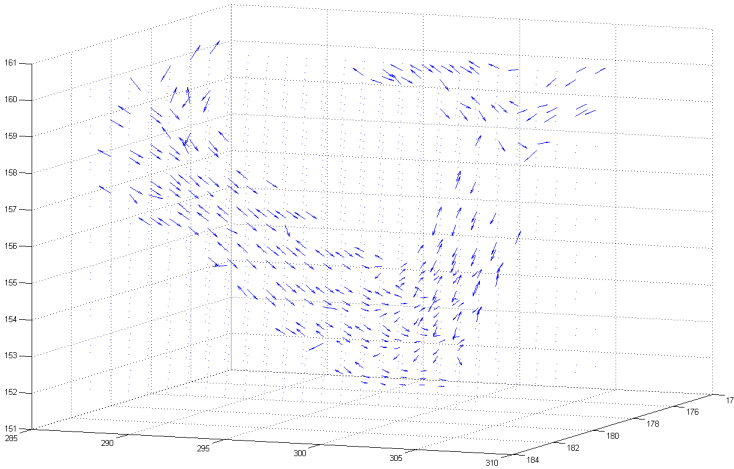

Fig. 1. 3D vector field plot from Hessian eigenvector decomposition. Eigenvectors with the lowest magnitude eigenvalue correspond to the direction of least curvature and thus point in the direction of vessel flow. Vectors were used to create accurate PF prediction steps.

Some vessel segments change direction abruptly and cannot be captured by the translational motion model, especially at bifurcation points. To track this movement we employed a vector field model for motion prediction. A vector field was produced by eigenvector decomposition of the Hessian matrix. The eigenvector associated with the lowest magnitude eigenvalue indicates the direction of least curvature, corresponding to the direction of vessel flow. The algorithm switches from translational motion to vector field motion model in the presence of a strong vector signal. Such utilization of Hessian eigenvectors for motion prediction is novel to medical image analysis. Fig. 1 shows the vector field on a short segment of the artery.

2.3   Likelihood Models Combined with Multiple Cues

In previous particle filtering vessel segmentation work on CT   [2, 3], intensity is used as the dominant information. Upon inspection of CT images for vessel segmentation, radiologists not only check intensity information, but also utilize anatomical information such as organ location, regional vesselness, and fat and muscle tissue. Thus human vision combines multiple cues during vessel tracking. Inspired by radiologists, we propose a new likelihood model for vessel segmentation by fusing multiple cues. Fig. 2 shows four tracking domains used to produce cues on an axial slice CT image.


Fig. 2. Domains used to generate vessel tracking cues. A) Vesselness response from Hessian analysis, with a mask generated by thresholding and post-processing. B) A minimum spanning tree algorithm was applied to intensity and vesselness features to connect low contrast segments of the artery. C) A MIP mask was used to amplify vessel signal in low contrast segments. D) Global, unthresholded vesselness response to Hessian analysis.

Intensity Cue. As with traditional vessel segmentation methods, intensity is the most important information for vessel tracking on CT. For a particle $\mathbf{x}_t^i$ at time point t, $i = 1, ..., N$, where N is the total number of particles, we extracted a spherical search region and summed the intensities within the sphere as our intensity cue for tracking. The single voxel particle intensity was used as an additional cue.

Vesselness Cue. Because the majority of vessel segments exhibit tube structure, a vesselness cue is essential to differentiate true vessels from noisy, bright, blob-like areas. We employed Li's multiscale vessel enhancement filtering [9] to provide this vesselness cue. Spatial scale standard deviations from 0.5 to 2 voxels with 0.25 voxel incremental steps were used for multi-scale analysis. Three vesselness cues were utilized: single voxel particle vesselness, vesselness sum within the spherical search region, and a binary vessel mask produced by thresholding vesselness response and applying ray casting and connected component analysis post-processing.

Organ Cue by Ray Casting. Nearby organs are a major source of false positives, including the bowel, liver, and kidneys. Tracking paths can be attracted to organ boundaries having line or curve character. To avoid these tracks, a ray casting technique was applied at each particle. Rays are casted in 26 spatial directions, and halt at either low intensity or a maximum distance, both determined heuristically.

Maximum Intensity Projection Cue. Maximum intensity projection (MIP) provides a method to amplify intensity signal in a selected direction. This is an informative cue for noisy data and thin, peripheral vessel segments with poor contrast enhancement. MIP was applied in pre-processing to the volumetric data based on several pre-selected directions. We then project the 2D detections back to 3D space.

Missing Vessel Cue. Due to non-uniform blood flow or vessel constriction, some vessel segments may have particularly low enhancement. These segments are not well distinguished by intensity and vesselness cues alone, which necessitates global context information to track these difficult areas. We employed a minimum spanning tree to connect segments with very high vesselness response, and generated a missing vessel cue mask prior to tracking. The use of a minimum spanning tree as a tracking cue is also novel to the medical image analysis field.

Fusing of Multiple Cues. Eight tracking cues are fused as a likelihood product:

$$L\left(\mathbf{y}_t \middle| \mathbf{x}_t\right) = L_{I_s}\left(\mathbf{y}_t \middle| \mathbf{x}_t\right) \times L_V\left(\mathbf{y}_t \middle| \mathbf{x}_t\right) \times L_O\left(\mathbf{y}_t \middle| \mathbf{x}_t\right) \times L_M\left(\mathbf{y}_t \middle| \mathbf{x}_t\right) \times L_B\left(\mathbf{y}_t \middle| \mathbf{x}_t\right), \quad (4)$$

where the five likelihood functions correspond to intensity, vesselness, organ, MIP and bridge cues. In the fusion process, each cue is treated independently and uniformly regarding its weight. Cues were taken to be independent, which is a common assumption used in computer vision for a Naïve Bayes methodology [7, 8]. For each cue, a kernel density estimator (KDE) was leveraged to learn the target distribution based on a 10 patient training set. During tracking, cue observations for each particle were weighted probabilistically using the respective KDE's to update the vessel location. Fig. 3 shows the KDE for each cue. The intensity and vesselness cues above contain 2 and 3 independent sub-cues, respectively.

2.4   Automatic Bifurcation Detection

The marginal artery is composed of several loops that bifurcate frequently, presenting a challenge to a local tracking method. To solve this, we implemented a robust automatic bifurcation detection system using a spherical shell search region. At each step, the shell was checked for high intensity voxels in the enhanced vessel range. A single vessel entering and leaving the shell produced two high intensity patches, but at a bifurcation the shell identified three patches indicating three paths leaving the sphere. In this case, multiple parallel paths were initialized to complete the vessel tree.

3   Dataset and Experimental Results

Our dataset contained 17 patients with high-resolution contrast-enhanced CT angiograms, 10 for KDE training and 7 for validation. Data acquisition and analysis were
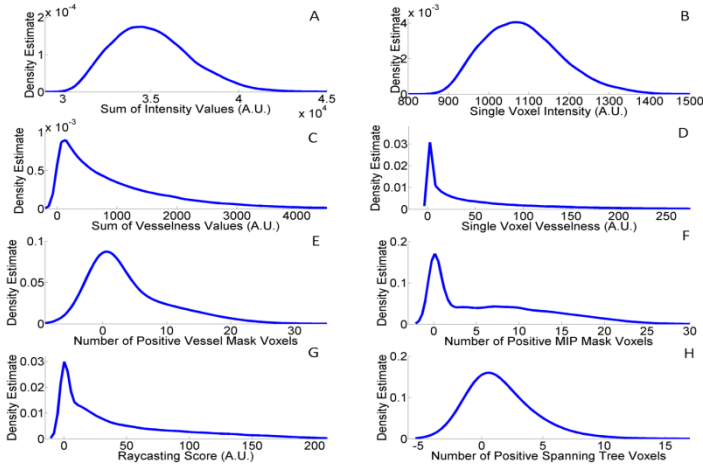
Fig. 3. Trained KDE for each cue. A) Intensity sum in spherical search region. B) Single voxel intensity. C) Hessian vesselness sum in spherical search region. D) Single voxel Hessian vesselness response. E) Number of positive voxels from binary thresholded vessel mask with post-processing. F) Number of positive voxels from MIP mask. G) Single voxel ray casting score to identify organ. H) Number of positive voxels from binary minimum spanning tree mask.

conducted under an Institutional Review Board (IRB) approved protocol. CT scans were acquired following oral administration of 3 bottles Volumen and intravenous administration of 130 ml Isovue-300 with 5 ml/sec injection rate and 30 second delay. The scanning parameters were section collimation 1.0-mm, reconstruction interval 0.5 mm, 512x512 matrix and in-plane pixel dimensions of 0.82 mm to 0.94 mm depending on the participant's body size. A major inclusion criterion for the testing set was high levels of visceral fat content for good spatial separation of the artery.

The proposed method was evaluated on the two largest and typically best enhanced segments of the marginal artery, which run parallel to the transverse and descending colon. A manual seed point was designated at the bifurcation point between these two segments, and the algorithm was allowed to track in the three initial vessel directions. The tracking algorithm required a runtime of approximately one hour per patient. Fig. 4 shows the segmentation result on these branches for one patient. Fig. 5 shows the recall and precision rates for each testing patient. Compared to the traditional baseline Hessian analysis method for vessel segmentation [9], our SMC multiple cue fusion algorithm achieved an average recall of 88.5% while improving the average segmentation precision to 32.2%. Baseline average recall and precision rates were 91.4% and 7.9%, respectively. Recall was defined as the fraction of ground truth voxels detected by the algorithm, and precision was defined as the fraction of detected voxels that were true detections. Paired student t-test comparison between our algorithm and the baseline method showed a p value of 0.639 for recall, and a precision p value of 0.0002, indicating significance in the precision improvement.

Fig. 4. 3D segmentation of the marginal artery with pelvis and spine for reference. The portion shown communicates between the SMA and IMA. Ground truth is labeled in green, and SMC detection is labeled in red. Detection shown has recall of 94.9% and precision of 58.3%.


Fig. 5. Experimental results comparison between the SMC multiple cue fusion method and baseline Hessian vessel analysis. The SMC cue fusion average recall for the 7 testing patients was 88.5%, compared to the baseline average recall of 91.4%. Average precision for SMC cue fusion was 32.2% compared to the baseline average precision of 7.9%.

4   Conclusion and Discussion

We have proposed a novel Bayesian tracking framework using SMC and multiple cue fusion to automatically track and segment the marginal artery of the colon on contrast-enhanced CTA. Such an algorithm was novel to medical image analysis, and advantageous compared to other vessel tracking methods by incorporating more

information for robustness. Utilizing local and global information, we achieved high recall and an increase of precision by a factor of 4 compared to the baseline method. It is important to note that the vast majority of false positive detection occurred on other segments of abdominal vasculature due to frequent anastomosis and our robust bifurcation detector. Thus, an extended study evaluating the algorithm on the complete marginal artery or abdominal vessel tree would likely increase precision results.

References

1. Lesage, D., Angelini, E.D., Bloch, I., Funka-Lea, G.: A review of 3D vessel lumen segmentation techniques: Models, features and extraction schemes. Medical Image Analysis 13, 819–845 (2009)
2. Florin, C., Paragios, N., Williams, J.: Particle filters, a quasi-monte carlo solution for segmentation of coronaries. In: Duncan, J.S., Gerig, G. (eds.) MICCAI 2005. LNCS, vol. 3749, pp. 246–253. Springer, Heidelberg (2005)
3. Schaap, M., Manniesing, R., Smal, I., van Walsum, T., van der Lugt, A., Niessen, W.J.: Bayesian tracking of tubular structures and its application to carotid arteries in CTA. In: Ayache, N., Ourselin, S., Maeder, A. (eds.) MICCAI 2007, Part II. LNCS, vol. 4792, pp. 562–570. Springer, Heidelberg (2007)
4. Lacoste, C., Finet, G., Magnin, I.E.: Coronary tree extraction from X-ray angiograms using marked point processes. In: IEEE International Symposium on Biomedical Imaging: Nano to Macro (2006)
5. Friman, O., Hindennach, M., Kuhnel, C., Peitgen, H.O.: Multiple hypothesis template tracking of small 3D vessel structures. Medical Image Analysis 14, 160–171 (2010)
6. Wu, Y., Huang, T.S.: Robust visual tracking by integrating multiple cues based on co-inference learning. International Journal of Computer Vision 58, 55–71 (2004)
7. Brasnett, P., Mihaylova, L., Bull, D., Canagarajah, N.: SequentialMonte CarloTracking by Fusing Multiple Cues in Video Sequences. Image and Vision Computing 25, 1217–1227 (2007)
8. Moreno-Noguer, F., Sanfeliu, A., Samaras, D.: Dependent multiple cue inte-gration for robust tracking. IEEE Transactions on Pattern Analysis and Machine Intelligence 30, 670–685 (2008)
9. Li, Q., Sone, S., Doi, K.: Selective enhancement filters for nodules, vessels, and airway walls in two- and three-dimensional CT scans. Medical Physics 30, 2040–2051 (2003)