Model-Guided Extraction of Coronary Vessel Structures in 2D X-Ray Angiograms

Shih-Yu Sun, Peng Wang, Shanhui Sun, and Terrence Chen

Siemens Corporation, Corporate Technology, Princeton, NJ, USA
{shanhui.sun,terrence.chen}@siemens.com

Abstract. Analysis of vessel structures in 2D X-ray angiograms is important for pre-operative evaluation and image-guided intervention. However, automated vessel segmentation in angiograms, especially extraction of the topology such as bifurcations and vessel crossings, remains challenging mainly due to the projective nature of angiography and background clutter. In this paper, a novel framework for model-guided coronary vessel extraction in 2D angiograms is presented. In this framework, a graph is constructed using a sparse set of pixels in the angiogram. With a single user-supplied click as the starting point, the vessel tree structure in the angiogram is automatically extracted from the graph. Ambiguities in this tree structure caused by 3D-to-2D projection are then resolved using topological information from the 3D vessel model of the same patient. By incorporating this prior shape information, the proposed method is effective in extraction of vessel topology, and is robust to background clutter and uneven illumination. Through quantitative evaluation on 20 angiograms, it is shown that this model-guided approach significantly improves detection of vessel structures and bifurcations.

1 Introduction

Vessel analysis in 2D X-ray angiograms is important for pre-operative evaluation and image-guided intervention. Automation of this process, however, remains challenging. One reason is the 3D-to-2D projective nature of X-ray angiography, which results in loss of structural information in 3D geometry and thus makes it difficult to tell vessel bifurcations from vessel crossings, for instance, purely from 2D imagery. The existence of background clutter, such as organs, bones, and interventional instrument, further complicates the situation and makes accurate vessel analysis even more challenging. Some of the challenges are shown in Fig. 1.

Most work on vessel analysis in 2D imagery is focused on pixel-wise detection of vessels, usually aided by enhancement filtering. In [7], feature vectors are constructed for vessel detection, which could potentially give a discontinuous mask. On the other hand, exploratory graph-based approaches give a continuous mask marking occurrence of vessels [1,5,9]. Furthermore, [4] and [8] make assumptions on vessel shape continuity, and incorporate the scale and orientation information given by enhancement filtering for better graph-based vessel extraction. See [3] for an extensive review of the related literature as well as vessel analysis on 3D

P. Golland et al. (Eds.): MICCAI 2014, Part II, LNCS 8674, pp. 594–602, 2014.

(a) (b) (c)

Fig. 1. Challenges for vessel extraction in angiograms: (a) uneven illumination (bottom and top right), (b) catheter and vessel crossings, and (c) spine and catheter

imagery. The methods described above are effective in giving a mask indicating occurrence of vessels or a single curve connecting two user-supplied points that approximates the centerline of a vessel segment. Nevertheless, none of these methods extract topology of the vessel structures, and therefore vessel crossings could be incorrectly identified as bifurcations, for instance. Moreover, these methods tend to be sensitive to background clutter and uneven illumination, which routinely occur in 2D X-ray angiography.

To address the above issues, a graph-based framework is presented in this paper for model-guided extraction of vessel structures in 2D X-ray angiograms. In this framework, the vessel tree structure is automatically extracted with a single user-supplied click on the angiogram indicating the root. This structure is then refined by following the topology of a 3D vessel model extracted from the Computed Tomography Angiography (CTA) of the same patient. Overall, the major contributions of this work are: (1) A novel framework is presented for semi-automatic extraction of the topological vessel structure in 2D angiograms, and not just vessel occurrences. (2) A novel use of prior shape information from 3D vessel models is presented, which enhances efficacy of the proposed method in telling vessel bifurcations from crossings and its robustness to background clutter as well as uneven illumination.

2 Graph-Based Vessel Tree Extraction in 2D Angiograms

To extract vessel structures, a graph is constructed using a sparse set of image pixels as nodes. Nodes in the neighborhood of each other are connected by graph edges, whose cost functions incorporate the output of a vessel enhancement filter, edge length and edge orientation. Vessel tree extraction is then initialized by a single user-supplied point around the root in the angiogram. The paths with the lowest accumulated edge costs from this starting point to all the other nodes in the graph are found by the Dijkstra's algorithm. These paths are guaranteed to form a tree structure, which is consistent with vessel topology. Finally, branches in the tree structure are pruned based on quality of the constituent edges.

2.1 Graph Construction and Exploratory Shortest Paths

Graph nodes are defined at a sparse set of pixels sampled from an initial segmentation mask, which marks highly likely vessel occurrences. This mask is obtained by thresholding the pixel-wise "vesselness" values given by a multiscale enhancement filter [2]. This mask is skeletonized and sparsified by iteratively sampling local maxima of vesselness values within a pre-defined kernel. The user-supplied point is added as the starting node for vessel tree construction.

A bi-directional edge $e_{i,j}$ is created between the i-th and j-th nodes if they are within N pixels of each other ($N = 20$ here). The probability of $e_{i,j}$ being part of vessel is denoted by $p_{i,j}$. From all the possible paths from the starting node to a given node (denoted by Γ), a naive Bayes classifier is employed to find the optimal path Γ^*, which is the most likely to correspond to a vessel segment:

$$\Gamma^* = \underset{\Gamma}{\operatorname{argmax}} \prod_{e_{i,j} \in \Gamma} p_{i,j} = \underset{\Gamma}{\operatorname{argmin}} \sum_{e_{i,j} \in \Gamma} (-\log(p_{i,j})). \tag{1}$$

This optimization is modeled as a shortest-path problem, where the sum of edge costs $C_{i,j}$ is minimized over possible paths. $C_{i,j}$ is defined as follows:

$$C_{i,j} = C^L_{i,j} \cdot (C^V_{i,j})^\alpha \cdot (C^I_{i,j})^\beta \cdot (C^\theta_{i,j} \cdot C^{\phi_1}_{i,j} \cdot C^{\phi_2}_{i,j})^\gamma, \tag{2}$$

where the superscripts L, V and I correspond to the edge length, vesselness value, and image intensity, respectively. The cost terms with superscripts θ, ϕ_1 and ϕ_2 describe consistency between the edge orientation and the pixel-wise orientations given by the enhancement filter. The non-negative weights α, β and γ adjust influences of these cost terms and are determined empirically (a zero weight essentially eliminates the corresponding term.) The respective cost terms are defined as follows:

$$C^L_{i,j} = (d_{i,j})^2 \tag{3}$$

$$C^V_{i,j} = \frac{1}{N_{i,j}} \sum_{x \in \mathcal{X}_{i,j}} (v_x^{-1})^2 \tag{4}$$

$$C^I_{i,j} = \frac{1}{N_{i,j}} \sum_{x \in \mathcal{X}_{i,j}} (I_x - (I_{f(i)} + I_{f(j)})/2)^2 \tag{5}$$

$$C^\theta_{i,j} = \frac{1}{N_{i,j}} \sum_{x \in \mathcal{X}_{i,j}} (\theta_x - (\theta_{f(i)} + \theta_{f(j)})/2)^2 \tag{6}$$

$$C^{\phi_1}_{i,j} = (\phi_{i,j} - \theta_{f(i)})^2,\ C^{\phi_2}_{i,j} = (\phi_{i,j} - \theta_{f(j)})^2 \tag{7}$$

The term $d_{i,j}$ denotes the edge length, $\mathcal{X}_{i,j}$ denotes the set of pixels passed by the edge, and $N_{i,j}$ is the number of elements in this set. For the x-th pixel, v_x denotes the vesselness value, I_x denotes the image intensity, and θ_x denotes the orientation given by the enhancement filter. The function $f(i)$ gives the pixel index at which the i-th node is located. Finally, $\phi_{i,j}$ denotes the edge orientation.

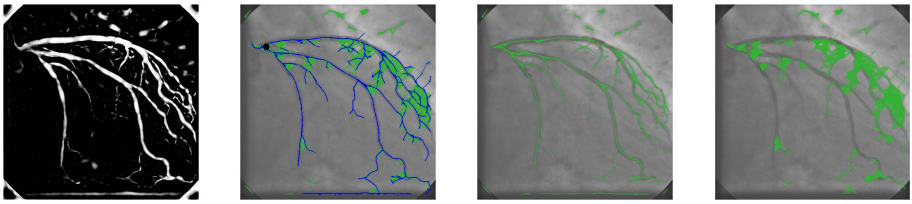

(a) (b) (c) (d)

Fig. 2. (a) The vesselness map for Fig. 1(a). (b) All the edges in the graph (green) and the extracted tree structure using the Dijkstra's algorithm (blue). The black dot around the root of the vessel tree (top-left) marks the user-supplied starting point. (c) Graph edges with a quality score of +1. (d) Graph edges with a quality score of -1.

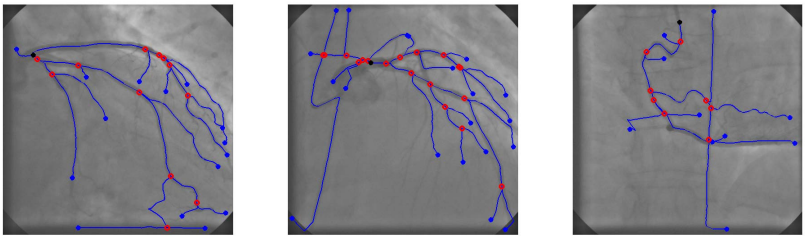

(a) (b) (c)

Fig. 3. Graph-based 2D segmentation results on the angiograms shown in Fig. 1. Identified bifurcations are marked by red circles.

The Dijkstra's shortest path algorithm is performed on this graph, starting from the user-supplied point to all the other nodes, which gives a directed tree structure rooted by the starting point. See Fig. 2(a) and (b) for the vesselness map and the shortest paths, respectively, for Fig. 1(a).

2.2 Branch Pruning

The extracted tree is pruned based on edge quality to eliminate branches that are less likely to be part of vessel. All edge costs are first clustered into two groups by the k-means method. The group of edges with lower costs are assigned a quality score of +1, and the other assigned -1. As shown in Fig. 2(c) and (d), a +1 score is a strong indicator of vessel occurrence whereas -1 is not. The quality score of a sub-tree is defined as the sum of scores of all the constituent edges.

Quality of the sub-tree rooted by each node is examined. For a non-bifurcation node (i.e. a node with one or zero child), its sub-tree is preserved only if its quality score is positive. For a bifurcation node, the sub-trees rooted by its children with non-positive scores are removed. Out of the remaining sub-trees, only the one

with the lowest mean edge cost and those with a total length of over 30 pixels are preserved (around 8.1 mm). Essentially, this pruning process allows both shortening and total elimination of branches. A balance is also struck between maximizing the quality score of the whole tree and preserving long branches. See Fig. 3 for the final 2D segmentation results on the angiograms shown in Fig. 1.

3 Tree Structure Segmentation Guided by 3D Models

As shown in Fig. 3, the 2D segmentation method suffers from background clutter and vessel crossings in the angiogram. These challenges are particularly difficult to address in single 2D angiograms due to the 3D-to-2D projective nature. Therefore, a model-based approach is presented for refinement of the 2D vessel tree.

In this approach, the 2D tree is refined by using prior shape information from a 3D vessel model of the same patient, which could be extracted pre-operatively from CTA, for instance. The 3D-to-2D projection of the 3D model is first non-rigidly registered with the 2D tree [6], an example of which is shown in Fig. 4(a). Based on this spatial alignment, correspondences are established between vessel segments from the 2D tree structure and those from the 3D tree. Finally, the 2D vessel segments are restructured following topology of the 3D vessel tree, which gives a 2D tree that accurately describes the vessel structure in the angiogram.

3.1 Correspondences between 2D and 3D Vessel Segments

The 2D vessel tree is first disconnected at bifurcations, giving a set of 2D segments. A segment is further broken into multiple shorter segments at edges with a quality score of -1 (Section 2.2), if any. All these 2D segments form a set $\mathcal{S}_{2D}$, an example of which is shown in Fig. 4(b). The registered projection of the 3D model is also disconnected at bifurcations, giving another set of segments $\mathcal{S}_{3D}$.

Segment correspondences between $\mathcal{S}_{2D}$ and $\mathcal{S}_{3D}$ are established based on spatial closeness. For each segment in $\mathcal{S}_{2D}$, the closest segment in $\mathcal{S}_{3D}$ is found based on mean pixel distances. All the 2D segments corresponding to the i-th 3D segment in $\mathcal{S}_{3D}$ forms a set $\mathcal{S}_{2D}^{(i)}$, and $\{\mathcal{S}_{2D}^{(1)}, \mathcal{S}_{2D}^{(2)}, \ldots\ldots \mathcal{S}_{2D}^{(K)}\}$ are therefore disjoint subsets of $\mathcal{S}_{2D}$, where K is the number of segments in $\mathcal{S}_{3D}$. To avoid false vessel segments in the final tree, a 2D segment is removed if less than 80% of its pixels are within 30 pixels (around 8.1 mm) from the closest 3D segment.

3.2 Restructuring 2D Vessel Segments Following 3D Topology

Given the segment correspondences, segments in $\mathcal{S}_{2D}$ are structured following connectivity of segments in $\mathcal{S}_{3D}$, which describes topology of the 3D tree. This structuring is first performed within each set $\mathcal{S}_{2D}^{(i)}$ and then between the sets $\{\mathcal{S}_{2D}^{(1)}, \mathcal{S}_{2D}^{(2)}, \ldots\ldots \mathcal{S}_{2D}^{(K)}\}$. Segments in each set $\mathcal{S}_{2D}^{(i)}$ are first connected locally using the Dijkstra's algorithm (Section 2.1), giving a connected group of 2D segments for each 3D segment. These groups $\{\mathcal{S}_{2D}^{(1)}, \mathcal{S}_{2D}^{(2)}, \ldots\ldots \mathcal{S}_{2D}^{(K)}\}$ are then connected locally by the Dijkstra's algorithm following connectivity of the 3D segments.

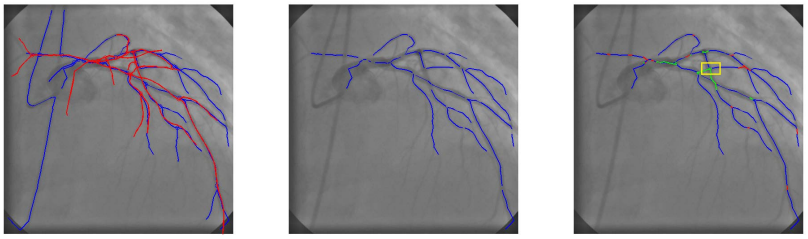

(a) (b) (c)

Fig. 4. (a) Projection of the 3D vessel model from CTA (red) registered with the 2D segmentation (blue). (b) Segments from the 2D tree structure. (c) Local connections joining the segments in (b). Connections corresponding to the same segment in $\mathcal{S}_{3D}$ are shown in red, while those corresponding to different segments are shown in green. A zoom-in image of the region highlighted in yellow is shown in Fig. 5(b).

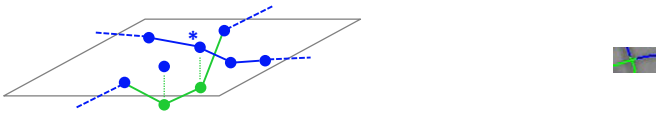

(a) (b)

Fig. 5. (a) Node splitting to avoid incorrect node connection. To establish the local connection marked in green, the nodes along the path are split (in green) to avoid incorrect vessel connection at the node marked by an asterisk. The original and new nodes are essentially independent nodes at the same image coordinates. (b) An example from Fig. 4(c) where creation of a false bifurcation is avoided by node splitting, as can be seen from the resulting tree structure shown in Fig. 6(b).

To avoid creation of false bifurcations during restructuring, the nodes connected by the newly created local paths are split, except the source and destination nodes (illustrated in Fig. 5). Multiple local paths are therefore allowed to cross each other without creating bifurcation nodes. Fig. 4(c) shows an example of local connections. As can be seen from the results in Fig. 6, this model-guided approach is effective in telling vessel crossings from bifurcations (Fig. 6(a) and (b)), as well as in coping with background clutter (Fig. 6(b) and (c)).

4 Experimental Results and Discussion

Vessel structures were extracted from 20 512×512 angiograms with and without guidance from 3D models, which were generated by manual segmentation. Results were compared with ground truth obtained by manually segmenting the 2D angiograms by an expert, where there were 130 bifurcations and 66581 vessel points in total. Parameters were selected based on data sets not included in the test set and remained consistent. The imager pixel spacing is 0.27×0.27 mm.

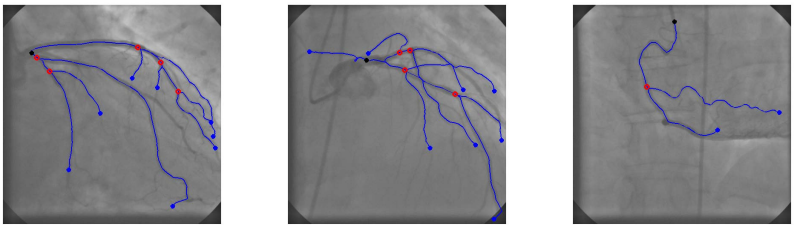

(a) (b) (c)

Fig. 6. Extracted tree structures from the angiograms shown in Fig. 1 with model-guided refinement. Identified bifurcations are marked by red circles.

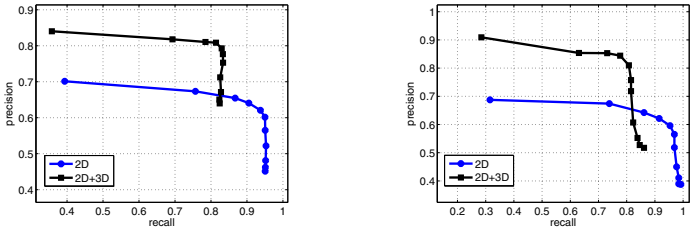

(a) (b)

Fig. 7. Precision-recall curves for (a) vessel points and (b) bifurcation points. Results from both 2D segmentation and model-guided segmentation (2D+3D) are shown.

The precision-recall curves for detection of vessel and bifurcation points are shown in Fig. 7, where operating points were changed by varying the cost threshold for assigning edge quality scores in branch pruning (Section 2.2). For a given threshold, the precision is the fraction of detected points that are correct, and the recall is the fraction of ground-truth points that are correctly detected. Here a detected vessel point is considered correct if there is a ground-truth point at the same pixel location. A detected bifurcation is correct if there is a corresponding ground-truth bifurcation in its 30-pixel neighborhood (around 8.1 mm).

As shown in Fig. 7, with 3D model guidance, the precision improves significantly for both vessel and bifurcation detection (by around 0.15 at a recall of 0.8, for instance). This improvement is expected since the model provides additional information about the vessel structure, which makes segmentation more robust to noise, background clutter, uneven illumination, and vessel crossings. For instance, the false positives due to catheters and spines in Fig. 3(b)-(c) are removed since these false vessel segments do not correspond to any 3D segments in restructuring. Similar improvement can be seen in Fig. 3(a) and Fig. 6(a), where uneven illumination at the bottom could cause false positives. Additionally, as can be observed by comparing Fig. 3(b) and Fig. 6(b), false-positive bifurcations due to vessel crossings are removed thanks to the information on connectivity of

vessel segments from the 3D model. Reduction of these false positives based on the 3D model directly contributes to the improvement in precision.

There is a drop in the highest recall rate that can be achieved after restructuring the 2D segmentation results with 3D model guidance, as also can be observed from the results in Fig. 6. This phenomenon is due to the fact that the refined structure is essentially a subset of the 2D segmentation that is consistent with the 3D model, which is created from manual segmentation and includes only clinically important vessel structures. Therefore, this drop in recall reflects mainly the omittance of less clinically important vessel branches and could be easily addressed during the phase of segmentation on CTA.

5 Conclusion

A framework is presented for vessel structure extraction in 2D X-ray angiograms with a single user-supplied click. Through experiments on 20 angiograms, this method was quantitatively evaluated and the performance is shown to be significantly improved by incorporating prior shape information from 3D vessel models. Future work includes joint optimization of vessel segmentation and registration, as well as development of automated 3D model extraction.

Acknowledgement. Shih-Yu Sun (from MIT, Cambridge, MA, USA) and Peng Wang (from Google, New York City, NY, USA) contributed to this work during their employment (as an intern and a full-time employee, respectively) at Siemens Corporation, Corporate Technology, Princeton, NJ, USA.

References

1. Estrada, R., Tomasi, C., Cabrera, M.T., Wallace, D.K., Freedman, S.F., Farsiu, S.: Exploratory Dijkstra Forest Based Automatic Vessel Segmentation: Applications in Video Indirect Ophthalmoscopy (VIO). Biomedical Optics Express 3(2), 327–339 (2012)
2. Frangi, A.F., Niessen, W.J., Vincken, K.L., Viergever, M.A.: Multiscale Vessel Enhancement Filtering. In: Wells, W.M., Colchester, A., Delp, S. (eds.) MICCAI 1998. LNCS, vol. 1496, pp. 130–137. Springer, Heidelberg (1998)
3. Kirbas, C., Quek, F.: A Review of Vessel Extraction Techniques and Algorithms. ACM Computing Surveys (CSUR) 36(2), 81–121 (2004)
4. Pechaud, M., Keriven, R., Peyre, G.: Extraction of Tubular Structures Over an Orientation Domain. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR 2009), pp. 336–342. IEEE Press (2009)
5. Poon, K., Hamarneh, G., Abugharbieh, R.: Live-Vessel: Extending Livewire for Simultaneous Extraction of Optimal Medial and Boundary Paths in Vascular Images. In: Ayache, N., Ourselin, S., Maeder, A. (eds.) MICCAI 2007, Part II. LNCS, vol. 4792, pp. 444–451. Springer, Heidelberg (2007)

6. Rivest-Henault, D., Sundar, H., Cheriet, M.: Nonrigid 2D/3D Registration of Coronary Artery Models with Live Fluoroscopy for Guidance of Cardiac Interventions. IEEE Transactions on Medical Imaging 31(8), 1557–1572 (2012)
7. Sofka, M., Stewart, C.V.: Retinal Vessel Centerline Extraction Using Multiscale Matched Filters, Confidence and Edge Measures. IEEE Transactions on Medical Imaging 25(12), 1531–1546 (2006)
8. Wink, O., Niessen, W.J., Viergever, M.A.: Multiscale Vessel Tracking. IEEE Transactions on Medical Imaging 23(1), 130–133 (2004)
9. Yim, P.J., Choyke, P.L., Summers, R.M.: Gray-Scale Skeletonization of Small Vessels in Magnetic Resonance Angiography. IEEE Transactions on Medical Imaging 19(6), 568–576 (2000)