

# Finding Similar 2D X-Ray Coronary Angiograms

Tanveer Syeda-Mahmood<sup>1</sup>, Fei Wang<sup>1</sup>, R. Kumar<sup>1</sup>, D. Beymer<sup>1</sup>, Y. Zhang<sup>1</sup>,  
Robert Lundstrom<sup>2</sup>, and Edward McNulty<sup>2</sup>

<sup>1</sup> IBM Almaden Research Center, San Jose, CA, USA

<sup>2</sup> Kaiser San Francisco Medical Center, San Francisco, CA, USA

**Abstract.** In clinical practice, physicians often exploit previously observed patterns in coronary angiograms from similar patients to quickly assess the state of the disease in a current patient. These assessments involve visually observed features such as the distance of a junction from the root and the tortuosity of the arteries. In this paper, we show how these visual features can be automatically extracted from coronary artery images and used for finding similar coronary angiograms from a database. Testing on a large collection has shown the method finds clinically similar coronary angiograms from patients with similar clinical history.


## 1 Introduction

X-ray Coronary angiography is a commonly used technique to assess the state of coronary artery disease (CAD). During assessment, clinicians look for characteristic visual features, taking into account the overall disease burden, the complexity of individual lesions, and placing more weight on proximal stenoses of the coronary arteries. Even though there are quantitative assessment scores such as the Syntax Score[12], they require manual input of angiographic information. Thus the clinicians still characterize the disease by 'eyeballing' on salient visual features such as lumen variation or the relative thickness of arteries (see Fig. 1a-c)[5], the distance of the junctions from the root, the number of trifurcations, etc. In this paper, our goal is to automatically extract features from coronary angiograms that mimic this process and learn a distance matrix to retrieve similar coronary angiograms for purposes of clinical decision support.

Automatically deriving such salient features from coronary artery imaging is, however, a challenging problem. It requires reliable separation of the major coronary arteries from the background. Complete delineation of arteries is difficult due to the similarities in intensity distribution in the regions surrounding the arteries. It also requires a reliable detection of all major junctions and the tubular arterial segments between junctions to allow computation of features such as arterial width and curvature changes or tortuosities.

## 2 Related Work

Much of the existing work on coronary angiogram analysis has focused on the preprocessing and segmentation of coronary angiograms. The majority of the



**Fig. 1.** Observable semantic features in coronary angiograms. (a) The dark blobs along arterial walls are calcifications, and irregular thickness variations can be observed. (b) Lumen variations along arteries. (c) Thinner/fine arteries showing diffused vessels.

methods locate vessel regions using various filters, deformable model methods, and supervised or unsupervised learning-based approaches [3,4,10,2]. Completely automatic artery tree extraction has also been attempted in 3D CT Angiography [15], but often rely on user identification of root in 2D X-ray angiography [7]. Previous work has also studied the junction detection in arteries. While the majority of the work is on 3D CT angiography data, relying on a good 3D vessel tree model and a robust 2D-3D shape alignment algorithm, [1,15], junction detection in 2D X-ray angiography has been restricted to either sensing Y or X junctions in pixel neighborhoods or using the intersection of artery centerlines.

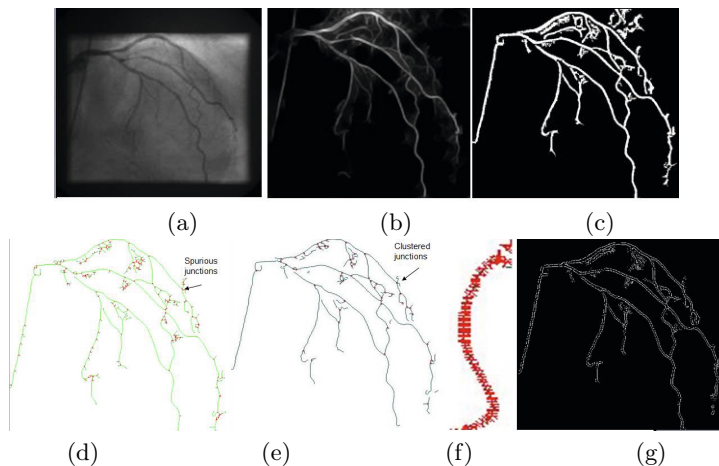
There is also work on quantitative characterization of coronary artery disease in the identification of coronary artery root. Popular angiography tools offer measurements such as luminal cross-sectional area and percentage area stenosis. Most of these tools, however, require some manual assistance including the identification of the coronary tree root.

### 3 Image Pre-processing of Coronary Angiograms

Our pre-processing uses well-known techniques put together in a new sequence to delineate coronary arteries. Since the clinical assessment focuses on major coronary arteries, accurate and complete tree reconstruction is not necessary.

#### 3.1 Highlighting Coronary Arteries

Starting from coronary angiogram video frames, we first extract a region of interest containing the arteries by exploiting the spatial and temporal variance in pixels. We then highlight the coronary artery vessel structure using a suitable ridge detection filter. While several filters could be used [4], we adopted Radon-Like Features (RLF)[6] as it does a non-isotropic sampling of neighborhoods based on edge sensing along different orientations. It has been shown to give a more complete highlighting of coronary arteries, including minor segments, while still suppressing noise. The result of RLF-filtering the coronary artery image of Fig. 2a is shown in Fig. 2c.



**Fig. 2.** Illustration of artery preprocessing. (a) original image. (b) RLF-filtered image. (c) Binary image obtained by adaptive thresholding. (d) Skeletonization and initial junction detection. (e) Junction clustering. (f) Tubular region detection for a single skeletal curve. (g) Tubular segments surrounding skeletal curves between junctions.

### 3.2 Locally Adaptive Statistical Thresholding

Next, we threshold the intensity gradients in the filtered region to separate the foreground vessel region from the background. Since a global threshold is insufficient, we model the filtered coronary image as the output of a short-space stationary process. Since coronary arteries have small thickness (less than 16 pixels), an overlapping window analysis with a small window size  $W \times W$  ( $W \leq 16$ ), is sufficient. Within each window of size  $W \times W$ , we find an optimal threshold  $T$  using similar ideas to Otsu thresholding [8], such that it separates the pixels within the regions into two classes with minimized intra-class variance  $\sigma_w^2(T) = \omega_1(T)\sigma_1^2(T) + \omega_2(T)\sigma_2^2(T)$  where  $\omega_1, \omega_2$  are the fraction of pixels belonging to the two classes. Fig. 2c shows the result of adaptive thresholding of the filtered image in Fig. 2b. Due to the narrow artery widths and the use of local Otsu thresholding, vessel fragmentation is minimized.

### 3.3 Junction Extraction from Foreground Regions

To locate the junctions, we adopt Zhang and Suen[14] to skeletonize the binary thresholded image, as it is fast, simple, and outperformed other approaches we tried. The skeletonization of Fig. 2c is shown in Fig. 2d, and it preserves the main artery centerlines. By grouping connected components on interior pixels (non-junctions) of the skeletal image, we form skeletal curves  $S_i = \{(x, y)\}$ . The set of junctions is  $J_m = (S_1, S_2, \dots, S_k)$  where the  $m$ th junction is the intersection of the incident skeletal curves  $S_1, \dots, S_k$  to give a junction of degree  $k$ . Spurious

initial junctions which are pixels with at least two incoming curves, are clustered using this grouping and the centroidal pixel becomes a robust indicator of the actual junctions, as illustrated in the change from Fig. 2d to Fig. 2e.

### 3.4 Extracting Coronary Artery Segments

To extract the tubular regions from the skeletal curves, we look for boundary pixels on either side of the skeletal curve proceeding along surface normals at each skeletal point. The tubular boundary may have points out of order at turning points, particularly where there is ambiguity in surface normals (Fig. 2f). These out-of-order points are corrected by treating the chain of endpoints on either boundary as pairs of curves to be aligned using dynamic time warping [13]. The resulting artery fragments are shown in Fig. 2g for the skeletal curves in Fig. 2d.

Each coronary artery segment  $C_i$  is represented by an ordered set of skeletal points  $\{(x, y, \delta_1, \delta_2, \theta)\}$  where  $(x, y) \in S_i$  is a skeletal pixel on the skeletal curve  $S_i$  passing through the tubular segment, and  $\delta_1$ , and  $\delta_2$  are the units along the surface normal  $\theta$  at which the tubular boundary points are detected.

## 4 Feature Extraction from Coronary Angiograms

Given the skeletal representation, we next extract clinically meaningful features. The proposed features are supported by several clinical studies including those in the SYNTAX score and JACC011 guidelines[5,12].

### 4.1 Number of Significant Junctions

This gives an indication of the bushiness of arteries as diffuse arteries tend to have a larger number of junctions. This feature  $f_1$  is simply recorded by the number of junctions  $J_m$  computed in the section above.

### 4.2 Thickness of Arteries

A blockage in the middle of the artery appears as a sudden change in the width of the artery. The average thickness of a coronary artery segment is given by  $W_i = \frac{\sum_j(\delta_{2j} - \delta_{1j})}{P}$  where the  $\sum_j$  is over the  $P$  skeletal points. The range of thickness variation within tubular regions is given as  $R_i = (\delta_{2i}^{max} - \delta_{2i}^{min})$ , and gives an indication of stenosis. The distribution of thickness of arteries is given by the feature  $f_2 = H(W_i)$  and  $f_3 = H(R_i)$ , where  $H(W_i)$  and  $H(R_i)$  are the histograms of the average thickness distribution and range respectively. The peaks in the histograms indicate the widths of dominant arteries and are useful in identifying the major coronary segments.

### 4.3 Number of Trifurcations

Trifurcations are useful to detect in cases when left main trifurcating coronary artery disease is present. Certain viewpoints (e.g. caudal) can cause trifurcation junctions to be detected, which is a good viewpoint descriptor. This feature is simply computed as  $f_4 = \{J_m, \text{degree}(J_m) > 3\}$ .

### 4.4 Tortuosity

Tortuosity is the number of curvature changes in the skeletal curves, measured using a histogram. The significant peaks in the histogram indicate the variation in tortuosity across the coronary artery segments. To estimate the curvature changes, we form a line segment approximation to the skeletal curve by recursively partitioning it at points of maximum deviation. These points of deviation are places where there is significant change in curvature. The tortuosity is normalized by taking the ratio of the curvature change points  $N_k$  over the total number of points  $N_i$  to give  $T_i = \frac{N_k}{N_i}$ . By repeating this over all curves, we form the tortuosity feature vector as a histogram over  $T_i$  as  $f_5 = H(T_i)$ .

### 4.5 Lengths of Artery Segments

The length of artery segments is important to assess early bifurcation of the left main coronary artery. Since the skeletal curves are available, this feature is easily computed by the pixel length of the skeletal curves and forming a histogram of it  $f_6 = H(|S_i|)$ , where  $|S_i|$  is the length of skeletal curve  $S_i$ .

### 4.6 Lumen Variations

To measure the lumen variations, we sample the intensity in the original image along surface normals to the skeletal curve and average it at each point along the skeletal curve. The range in intensity variations is similarly normalized by the maximum intensity in the region and its histogram is feature  $f_7 = H(I_i)$  where  $I_i$  is the normalized range of intensity variation in skeletal curve  $S_i$ .

## 5 Finding Similar Coronary Angiograms

By arranging the features into one long feature vector, each angiogram is represented by a vector  $F_c$ . Simple Euclidean distance comparisons between feature vectors is not sufficient to retrieve similar angiograms, both because of errors in vessel detection and the inherent variation in raw feature vectors.

Going past the Euclidean metric, we attempt to learn a distance metric so that vessels that are "similar" end up close to each other in feature space. Specifically, we adopted a recent work on a supervised metric learning method called Relevant Component Analysis (RCA)[11] as it has been shown to significantly improve clustering performance. RCA works by eliminating those dimensions that are

unlikely to be useful for classification using small subsets or chunklets of sample points. The chunklet covariance matrix is estimated as  $C = \frac{1}{N} \sum_i \sum_j (x_{ji} - m_j)(x_{ji} - m_j)^T$  where  $m_j$  denotes the mean of the  $j$ -th chunklet and  $x_{ji}$  the  $i$ -th element of the  $j$ -th chunklet. A whitening transformation is then associated with the covariance matrix  $C$   $W = C^{-\frac{1}{2}}$  to apply to the data points after dimensionality reduction[11].

To obtain the learning matrix, we first normalize all the feature vectors to be unit range and then annotate the features of a training set from distinct viewpoints, so that those from the same viewpoint and similar coronary anatomies are annotated with the same label. Using the resulting metric learned  $W$ , the distance between any two coronary angiograms is simply given by the Mahalanobis distance  $(X_2 - X_1)^T W (X_2 - X_1)$ . This distance is finally used to rank coronary angiogram images in a database using their respective feature vectors.

## 6 Results

From a collection of 1600 runs of X-ray angiography videos from 70 patients, we applied a keyframe detection method [13] to retain the top 10 key frames from each run, generating a ground truth test set of 600 images drawn across multiple patients, viewpoints and coronary arteries. The training set for metric learning was derived from another subset of keyframes chosen from the runs to depict distinct viewpoints such as anterior oblique projection, left anterior oblique, caudal and right coronary artery view. The training and testing was done on different sets of patients.

### 6.1 Accuracy of Junction Detection

First, trained experts manually counted the number of junctions observed in a set of about 130 left and right coronary images in 3 viewpoints. The spatial overlap with the automatically detected junctions is shown in Table 1. The 10% non-overlap in the spurious junctions is mostly from non-artery regions or from vessel overlaps and intersections (e.g. cross-overs) while all manually identified junctions were consistently detected. There is a large agreement between the manual and automatically found junctions. The lower overlap of RCA is due to the lower visibility of RCA over LCA in the X-ray images provided.

**Table 1.** Accuracy of junction detection in coronary arteries

Artery class	Number of Images	Avg. Manually detected Junctions per image	Automatically detected junctions	% Spatial Overlap
LCA	31	62.6	75.3	89.8%
LCA Caudal View	50	45.8	67.6	93.5%
RCA	50	24.3	37.5	74.3%

## 6.2 Accuracy of Coronary Artery Segment Extraction


First, trained experts marked major arteries using LabelMe[9] to trace the contours of the segments. We automatically detected tubular structures to extract the coronary segments and measured the spatial overlap between the two regions. Table 2 summarizes the average accuracy of the detection process in comparison with the ground truth for the three artery classes in our dataset. Here we define average overlap as the fraction of pixels that are in both manual and automatically extracted artery segments over the pixel area of the manual segments. Similarly, average non-overlap is defined the fraction of pixels in automatic regions that are not within manually indicated segments over the total number of white pixels in all manual and automatically detected regions. Due to the different normalization used, the two numbers need not add up to 100%. From this, we conclude that a large fraction of the artery regions are accurately detected.

**Table 2.** Accuracy of coronary artery segment detection

Artery class	Number of Images	Average Overlap	Average Nonoverlap
LCA	31	95.3%	12.4%
LCA Caudal View	50	93.8%	13.5%
RCA	50	86.2%	6.2%

## 6.3 Similarity Retrieval Performance

Using the learned distance matrix, we used query images from the 600 image test set to retrieve the top 10 most similar images from the same set. Fig. 3 shows a sample result with the query image in the top-left, and the top 5 similar images in ranked order ordered left to right, top to bottom. The retrieved images have similar topology, lumen variation, variation in artery thickness, and the same disease (left main). To evaluate precision and recall, we selected 200 images from the viewpoint set and asked trained experts to mark clinically similar images in the 600 image data set. We then measured the recall as the fraction of these images returned by the similarity ranking in the top K list while precision was



**Fig. 3.** Similarity retrieval of coronary angiograms. (a) Query angiogram image. (b)-(f) ranked order of matching coronary angiograms. (g) Precision-recall curve.

measured by the fraction of the top  $K$  matches returned that were relevant. The parameter  $K$  was varied to obtain the precision-recall curve. Fig. 3 shows the precision-recall curve using RCA-based metric learning in comparison to the Euclidean metric, indicating a large improvement in performance. In general, we found that the similarity retrieval preserved the identity of the arteries in the top 10 hits when the viewpoints were mixed in the dataset.

## 7 Conclusions

In this paper, we address for the first time, the problem of finding similar coronary angiograms using clinically meaningful features whose variation across patient population is learned using metric learning.

## References

1. Chalopin, C., Finet, G., Magnin, I.E.: Modeling the 3D coronary tree for labeling purposes. *Medical Image Analysis* 5(4), 301–315 (2001)
2. Chen, T., Funke-Lea, G., Comaniciu, D.: Robust and Fast Contrast Inflow Detection for 2D X-ray Fluoroscopy. In: Fichtinger, G., Martel, A., Peters, T. (eds.) MICCAI 2011, Part I. LNCS, vol. 6891, pp. 243–250. Springer, Heidelberg (2011)
3. Dehkordi, M.T., Sadri, S., Doosthoseini, A.: A review of coronary vessel segmentation algorithms. *Journal of Medical Signals and Sensors* 1(1) (2011)
4. Frangi, A.F., Niessen, W.J., Vincken, K.L., Viergever, M.A.: Multiscale Vessel Enhancement Filtering. In: Wells, W.M., Colchester, A.C.F., Delp, S.L. (eds.) MICCAI 1998. LNCS, vol. 1496, pp. 130–137. Springer, Heidelberg (1998)
5. Hsu, J.T.: Impact of calcification length ratio on the intervention for chronic total occlusions. *Intl. J. of Cardiology*, 135–141 (2011)
6. Kumar, R., Reina, A.V., Pfister, H.: Radon-like features and their application to connectomics. In: *IEEE Workshop MMBIA* (2010)
7. Lara, S., et al.: A semi-automatic method for segmentation of the coronary artery tree from angiography. In: *ACM Symposium on Graphics*, vol. 27(3), pp. 236–239 (2009)
8. Otsu, N.: A threshold selection method from gray-scale histogram. *IEEE Trans. Syst., Man, Cybern.* 9(1), 62–66 (1979)
9. Russell, B.C., et al.: LabelMe: a database and web-based tool for image annotation. *International Journal of Computer Vision*, 157–173 (2008)
10. Sato, Y., et al.: Three-dimensional multi-scale line filter for segmentation and visualization of curvilinear structures in medical images. In: *Medical Image Analysis*, pp. 143–168 (1998)
11. Shental, N., Hertz, T., Weinshall, D., Pavel, M.: Adjustment Learning and Relevant Component Analysis. In: Heyden, A., Sparr, G., Nielsen, M., Johansen, P. (eds.) *ECCV 2002, Part IV*. LNCS, vol. 2353, pp. 776–790. Springer, Heidelberg (2002)
12. Sianos, G., et al.: The syntax score: and angiographic tool grading the complexity of coronary artery disease. *EuroIntervention*, 219–227 (2005)
13. Syeda-Mahmood, T., et al.: Automatic selection of keyframes from angiogram videos. In: *ICPR* (2010)
14. Zhang, T.Y., Suen, C.: A fast parallel algorithm for thinning digital patterns. *Communications of ACM* 27(3), 236–239 (1984)
15. Zhao, F., Bhotika, R.: Coronary artery tree tracking with robust junction detection in 3D CT angiography. In: *ISBI*, pp. 2066–2071 (2011)