

Similarity Retrieval of Angiogram Images BASED on a Flexible Shape Model

Tanveer Syeda-Mahmood, Colin B. Compas, David Beymer, and Ritwik Kumar

IBM Almaden Research Center, San Jose, CA, USA

Abstract. In this paper we address the problem of finding similar coronary angiograms from a database of angiograms using a new constrained nonrigid shape model for the description of coronary arteries. The model captures the non-rigid variations in the artery shapes while still preserving the overall perceptual spatial layout based on the articulation constraints between arteries. Shape matching involves testing for class membership using the constraints specified in the model. The shape similarity method is demonstrated in a similarity retrieval application on a large database of angiogram images.

1 Introduction

Cardiologists use 2D X-ray coronary angiography routinely to image the coronary arteries that supply blood and oxygen to the heart muscle. They grade the severity of coronary artery disease by analyzing the coronary tree layouts, detecting arterial lesions in key coronary artery segment locations, and assessing their severity. The predominant way of using angiograms currently is in manual diagnosis at the point of care, in a single sample-guided manner. That is, only the data from the given patient is used along with a physician's a priori knowledge to make decisions. If the imaging data from the patient could be used to retrieve similar cases, it can lead to enhanced diagnosis where physicians can validate their current hypothesis, and check for any overlooked possibilities or alternate interpretations. The use of imaging as a modality to do population analysis has now become popular in cardiac decision support as evidenced by recent work using a variety of modalities such as Doppler imaging [10]. A key subproblem to be solved in such decision support is the development of a disease-specific similarity metric to judge the similarity between images of similar patients. Our paper focuses on characterizing the shapes of coronary arteries depicted in individual coronary angiogram video frames of a prior-labeled image database to find similarities.

Characterizing shapes of coronary arteries is challenging. There is considerable variation in the appearance of the arteries depending on the viewpoint (rotation and skew of the imaging setup) at which the sequence was taken, or based on the individual differences in the coronary anatomy of patients or due to effects of a disease. Figure 1a shows the variation in the appearance of the same coronary artery across persons for similar viewpoints. Figure 1b shows the difference in

the appearance of the same coronary artery (left coronary) of the same person under changes in viewpoint across runs of an angiogram study. Finally, the matching should be robust to the effects of imaging conditions such as brightness variations, background clutter and zoom that can also be seen in Figure 1. In this paper we present a novel shape model for description and recognition




Fig. 1. (a) Illustration of variation in appearance of the left coronary artery across patients under similar viewpoint. (b) Illustration of the variation in appearance of the left coronary artery for the same patient under changes in viewpoint.

of coronary arteries that captures the non-rigid variations in their shapes under the variations described above. The model captures the variations in artery shapes while still preserving the overall perceptual spatial layout based on the articulation constraints between arteries. We also demonstrate its application for similarity retrieval on a large database of angiogram images.

2 Related Work

Although there are no commercial tools that can interpret coronary arteries in a completely unattended fashion [1], there is a large body of work on interpretation of 3D and 2D angiograms including filter based methods (derivative-based filter [3], matching filter [7], etc.), deformable models, learning-based approaches (Neural Network [6], and Support Vector Machines [6]). The predominant way of capturing shape layout of arteries has been through a graph formalism using built-in knowledge of coronary artery tree [5]. Precise geometric modeling of the non-rigid shapes of these articulated structures however, is lacking. There is large body of work also in the field of shape descriptors. Besides rigid body recognition methods, there are a considerable number of models that allow nonrigid shape matching for regional shapes such as variants of 2D strings[12], region adjacency graphs, shape context [2], super pixels, and more recently, variants of pyramid match kernels[4]. While recognition could be attempted by such methods, they are not as suitable for capturing the correspondence of artery segments in the primarily long and thin network of tubular structures found in coronary artery images.

Thus although there are a number of shape models available for describing 2D and 3D shapes, automated recognition and retrieval from low visibility X-ray coronary angiographic imagery has been difficult and to our knowledge, has only been attempted recently using features such as tortuosity, lumen variations, etc. without establishing correspondence between the coronary tree layouts.

3 Artery Pre-processing

Adopting the approach described in [9], we first pre-process the image using a conventional edge detector (e.g. Canny edge detector). Since large intensity gradients are more likely to correspond to edges, a thresholding is applied at two levels to yield both weak (lower threshold) and strong (higher threshold) edges. We extract curves from these edge maps using 8- connected neighbors and tracing curves using depth-first search between junctions. We then detect tube-like vessel structures looking for evidence of closely-spaced almost parallel curves among the retained curves as described in [9]. Figure 4b shows the result of vessel extraction on a raw image shown in Figure 4a after these operations. Finally, corner features are extracted from the curves to serve as basic shape features for later curve matching during shape correspondence. The result of such corner extraction is overlaid on Figure 4b.

While the effect of background is minimized through this process, curves from the background will still be present and will be accounted under matching where we address spurious and missing features in the shape model as described below.

4 Modeling Shape Layout of Coronary Arteries

Since the viewpoint considerably changes the appearance of arteries, we factor its effect out by using the azimuth and elevation angle information provided in the DICOM annotation of the sequence. These angles can be clustered to form viewpoint classes using straightforward K-means clustering. Each artery image can then be classified into a viewpoint class and the following description shows how correspondence can be found between two artery configurations of the same viewpoint class (and artery). To motivate the need for articulated non-rigid shape modeling for arteries, consider the left coronary artery system consisting of three major arteries, left main (LMCA), left anterior-descending (LAD), and left circumflex (LCX) along with their many branches as shown in Figure 2a. Different patients may show variations in the number of branches but the three major arteries are mostly visible as shown in Figure 2b for a different person with the same disease. To reliably recognize such structures, the shape model and the associated matching metric must take into account both global and local deformations. Figure 2c illustrates this difficulty for the coronary artery images of Figure 2a and b. Even by using a good correlation measure, superposing the vessel structures from these images shows evidence of not only residual global rotation and translation but also local translational shifts along

a common direction which is roughly orthogonal to the dominant orientation of the shape conveyed by the network of arteries.

Our goal is to capture this characteristic of artery networks in the shape model. The key observation we make is that the coronary artery system is a specific class of articulated non-rigid shapes in which the movement of arteries is constrained to be a relative translation that preserves the ordering of the arteries. Specifically, we model a coronary artery system S_M as a collection of curves $\{C_1, C_2, \dots, C_m\}$. Members of the same shape class obey the following shape constraints:

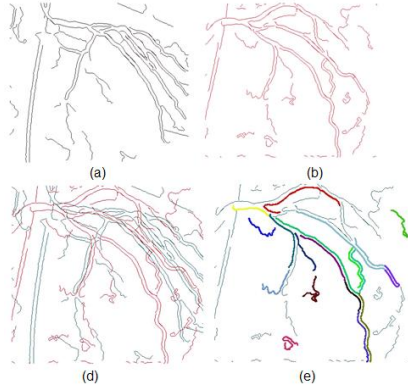


Fig. 2. Illustration of articulated non-rigid shape matching. (a) Candidate image from the database. (b) A query image showing the same artery but from a different person with the same disease. (c) Registering with global translation shows nonrigid deformations. (d) Result of nonrigid matching. The set of curves in the query that found a valid match are shown in different colors.

Direction of Residual Translation. All corresponding curves belonging to two shapes in the same class are related by a global translation along a common direction θ , i.e.

$$\mathcal{Y}_{ijy} = \mathcal{Y}_{ijx} \tan \theta \quad \forall i, j \quad (1)$$

where $\mathcal{Y}_{ijx} = C'_i x - C_j x$ and $\mathcal{Y}_{ijy} = C'_i y - C_j y$ and $C'_i x, C_j x$ are the corresponding points along the two corresponding curves C'_i and C_j .

Extent of Residual Translation Constraint. The extent of translation of all corresponding regions is bounded by δ so that

$$\|\mathcal{Y}_{ijx}\| \leq \|\delta \cos \theta\| \text{ and } \|\mathcal{Y}_{ijy}\| \leq \|\delta \sin \theta\| \quad (2)$$

for $\theta = 90$, the bound on the extent of translation becomes $\|\mathcal{Y}_{ijy}\| \leq \delta$. The above bound includes those cases where the curves do not translate allowing for non-rigid shifts in position along the common direction of translation.

Ordering of Curves Constraint. The ordering of corresponding curves along the direction of residual translation remains the same across members of the shape

class. The ordering of a set of curves $\{C_1, C_2, \dots, C_M\}$ is defined as a permutation of the indices $\{i_1, i_2, \dots, i_M\}$ such that

$$\begin{aligned} C_{i_k} \leq C_{i_{k+1}} \text{ if } \exists p(t_1) \in C_{i_k} & \quad (3) \\ q(t_2) \in C_{i_{k+1}} & \\ t_1 \leq t_2 & \\ \alpha(p) \leq \alpha(q). & \end{aligned}$$

where $i_k = j$ for some $1 \leq j \leq M$ and $\alpha(P)$ is the α -coordinate representation of a point P along the line of orientation θ so that $P = \alpha P_1 + (1 - \alpha)P_2$ and P_1, P_2 are two reference points along the line of orientation θ .

Under the above constraints, individual curves of the candidate instances of the model are allowed to undergo non-rigid shape warping in their parametric representations. Thus any two members $S_M = \{C_1(t), C_2(t), \dots, C_m(t)\}$ and $S_N = \{C'_1(t), C'_2(t), \dots, C'_n(t)\}$ are related under the constrained nonrigid shape model as

$$C'_i(t) = AC_j(\Phi_{ij}(t)) + \Upsilon_{ij} \quad (4)$$

where $\Phi_{ij} = b_{ij}t + \Gamma_{ij}(t)$.

The above shape model allows non-rigidity in the parametric representation of the curves through Γ_{ij} which is a function of t . Here A captures the global rotation and scaling of the overall shape (e.g. due to a change in viewpoint). Since Γ is also a function of the pair of curves C'_i, C_j thus allowing individual corresponding curves on the shape to stretch and compress to different extents. In addition, using a translation Υ_{ij} defined in Equation 1, we ensure that the corresponding curves are also allowed nonuniform translations as long as such translation does not disturb the spatial ordering of curves in Equation 3.

4.1 Shape Matching Using the Nonrigid Shape Model

Using the constrained nonrigid shape model, shape matching reduces to testing membership in the shape class under the model. That is, it involves determining if a sufficiently large set of corresponding curves can be found that satisfy the class constraints. Consider first the simpler case when there is no global rotation or scale. Given a model shape S_M from the class, and a candidate member shape S_I , we find potential matches to candidate curves in S_I within the model shape S_M that match to it under non-rigid shape warping while still preserving the curve ordering and obeying other model constraints.

Handling Pair-Wise Non-Rigid shape Distortions. The goal of allowing non-rigid distortion between corresponding curves for coronary vessel images is to allow for variations in length and thickness across patients, as well tolerance to local shape distortions due to stenosis. This means we must allow for gaps and insertions while still finding the longest possible pairs of curves.

To achieve this, we adapt the well-known technique of dynamic time warping used for time series matching to find matching candidate curve fragments

between a model shape and a candidate shape derived from a coronary artery map. Details of this method are described in [8].

Checking Ordering Constraints. Using a suitable threshold the best pairwise matches of curves can be retained. Since curves of shapes S_M and S_I are pre-ordered with respect to the direction of residual translation (eigenvectors), we obtain an ordered sequence of curves for each as $S_M = \{C_{r1}, C_{r2}, ..C_{rM}\}$, and $S_I = \{C'_{r1}, C'_{r2}, ...C'_{rN}\}$. A subset of size T in S_I was found as matching curves in S_M using dynamic shape warping earlier. For each curve $C'_i \in S_I$, the candidate match set can be denoted by $R_i = \{C_{i1}, C_{i2}, ..C_{ik}\}$ where $C_{ij} = C_{rl}$ for some $1 \leq l \leq M$. The candidate sets can be combined with the curve ordering information in S_M by representing the match as a set sequence $(R_1, ...R_T)$.




Fig. 3. Illustration of similarity retrieval of angiogram images. In each case, the first image is the query and the matches in ranked order are indicated in the same row. (a) Case of left coronary system. (b) Case of right coronary system with right dominance. (d) case of right coronary artery (RCA) with possibly left dominance (PDA is not easily visible).

Using the set sequence, the largest set of corresponding curves that preserve the spatial ordering with respect to the direction of residual translation can then be found as the longest common subsequence (LCS) of ordered model curves in S_M that are also a member sequence of the set sequence above. It can be shown that as in the case of LCS of two ordinary sequences [11], the problem of finding an LCS of a pair of sequence and set sequence also has optimal substructure property and can thus be derived using a similar algorithm as the dynamic shape warping mentioned above.

5 Finding Similar Coronary Angiograms

The overall algorithm for similarity retrieval from angiogram images consists of several offline processing steps and an online query-driven search. In the offline

processing stage, all frames of incoming angiogram videos first segmented into viewpoint classes as described earlier. An optional keyframe selection process can be applied to prune the large number of very similar successive images that are often present in the angiogram sequences as described in [9] so that the viewpoint clusters could be formed from keyframes alone. Each viewpoint cluster images are then processed to highlight vessel regions and their boundary curves as described in Section 3.

Given a new patient study, similar processing is applied to all video frames of an incoming angiogram study of a patient (query). Each image is then first classified into its corresponding viewpoint cluster. For each query image, the vessels are similarly extracted, the dominant orientation determined and curve ordering sequence is similarly obtained. The match to candidate images of its viewpoint cluster is then done serially using the constrained non-rigid shape matching described in the previous section. Candidate matching curves are determined using dynamic shape warping as described in Section 4.1. The longest common set subsequence using the constraints of the nonrigid model are then applied to the candidate matches to recover a curve correspondence of the query curves. The fraction of the perimeter covered by the matched curves is used to rank the images for retrieval.

6 Results

We assembled a large database of angiogram images from actual angiogram video sequences recorded at a catheter lab at a medical center in our area. This dataset depicts patients with various forms of coronary artery disease. A typical study consists of 20-25 runs, with each run consisting of 15-120 images depending on the complexity of the arteries being shown and their pathology. In our database we have a collection of over 70 patients covering a total of 1628 runs and 80,000 images. For the purpose of the study, we automatically classified the collection into 10 viewpoint classes using K-means clustering of viewpoint parameters in the DICOM files. For enabling precision and recall studies, the images within viewpoint classes were given a ground-truth label using a training process where physician experts chose one of 8 artery class labels (and 3 outlier class labels) including those that depict LMCA (left-main), LCX (circumflex), LAD (anterior-descending), RCA (right coronary), RMA, and PDA (Posterior-descending) for each depicted keyframe image. In cases where more than one artery was visible, the image was multiply classified. Three outlier classes were also added including a background class (BG) and incomplete depiction of artery class (INC) and an ambiguous class (AMG) where even experts had difficulty identifying the artery depicted due to noise, viewpoint or anatomic variation.

Retrieval Performance. We first illustrate results of similarity retrieval for example images from the various classes. Figure 3 shows the top ranked images for three queries corresponding to left main (LMCA), RCA with PDA and RCA. As can be seen, in each case the top images depict another instance of the same

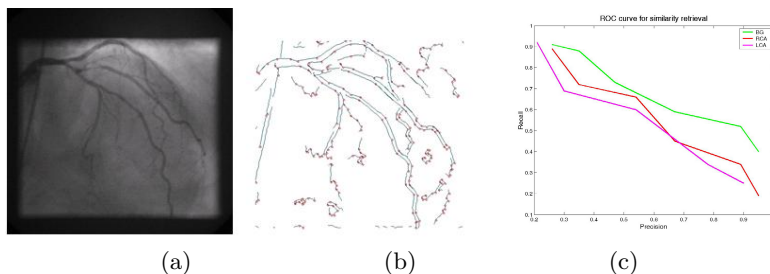


Fig. 4. (a) Illustration of vessel extraction - original image. (b) extracted vessel features. (c) precision-recall for similarity retrieval of angiogram images. Vertical axis is Recall, horizontal axis is precision.

artery under changes in global pose, and non-rigid artery variations. The occasional mismatches (4th match in row 3) are also reasonable depicting a PDA dominant RCA. The performance of our similarity search algorithm is illustrated by the ROC curve in Figure 4c. In general RCA variations are much less thus explaining the higher recall (red curve). Due to the vesselness measure distinguishing between background and vessel-depicting images, the performance is also good for rejecting outliers (green curve).

7 Conclusions

In this paper we have presented a new shape matching framework for similarity retrieval of angiogram video sequence images. The model accounts for both global and local deformations of arteries under changes in viewpoint, diseases and patient-specific variations. Experiments show promising results on artery recognition as an initial filter for decision support and diagnosis.

References

1. Medis medical imaging systems, Inc., <http://www.medis.nl/index.htm>
2. Belongie, S., Malik, J., Puzicha, J.: Shape matching and object recognition using shape contexts. *IEEE Trans. PAMI* 24, 509–522 (2002)
3. 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)
4. Grauman, K., Darrell, T.: The pyramid match kernel: Efficient learning with sets of features. *Journal of Machine Learning Research*, 725–760 (2007)
5. Haris, K., Efstratiadis, S., Maglaveras, N., Pappas, C., Gourassas, J., Louridas, G.: Model-based morphological segmentation and labeling of coronary angiograms. *IEEE-TMI* 18(10), 1003–1015 (1999)
6. Perfetti, R., Ricci, E., Casali, D., Costantini, G.: A cnn based algorithm for retinal vessel segmentation. In: *ICC 2008: Proceedings of the 12th WSEAS International Conference on Circuits*, pp. 152–157 (2008)

7. Sato, Y., Nakajima, S., Shiraga, N., Atsumi, H., Yoshida, S., Koller, T., Gerig, G., Kikinis, R.: Three-dimensional multi-scale line filter for segmentation and visualization of curvilinear structures in medical images. *IEEE Medical Image Analysis*, 143–168 (1998)
8. Syeda-Mahmood, T., Beymer, D., Wang, F.: Shape-based matching of ECG recordings. *IEEE EMBC, 2012–2018* (2007)
9. Syeda-Mahmood, T., Beymer, D., Wang, F., Mahmood, A., Lundstrom, R., Shafee, N., Holve, T.: Automatic selection of keyframes from angiogram videos. In: *ICPR*, pp. 4008–4011 (2010)
10. Syeda-Mahmood, T., Turaga, P., Beymer, D., Wang, F., Amir, A., Greenspan, H., Pohl, K.: Shape-based similarity retrieval of doppler images for clinical decision support. In: *IEEE CVPR*, pp. 855–862 (2010)
11. Cormen, T., et al.: *Introduction to algorithms*, p. 1985. MIT Press (1985)
12. Yeh, M., Cheng, K.T.: A string matching approach for visual retrieval and classification. In: *ACM-MIR* (2008)