Machine Learning Based Vesselness Measurement for Coronary Artery Segmentation in Cardiac CT Volumes

Yefeng Zheng[1], Maciej Loziczonek[1], Bogdan Georgescu[1], S. Kevin Zhou[1]
Fernando Vega-Higuera[2], and Dorin Comaniciu[1]

[1]Image Analytics and Informatics GTF, Siemens Corporate Research, Princeton, NJ, USA
[2]Computed Tomography, Siemens Healthcare, Forchheim, Germany

Abstract

Various coronary artery segmentation methods have been proposed and most of them are based on shortest path computation given one or two end points on the artery. The major variation of the shortest path based approaches is in the different vesselness measurements used for the path cost. An empirically designed measurement (e.g., the widely used Hessian vesselness measurement) is by no means optimal in the use of image context information. In this paper, a machine learning based vesselness measurement is proposed by exploiting the rich domain specific knowledge embedded in an expert-annotated dataset. For each voxel, we extract a set of geometric and image features. The probabilistic boosting tree (PBT) is then used to train a classifier, which assigns a high score for voxels inside the artery and a low score to those outside. The detection score can be treated as a vesselness measurement in the computation of the shortest path. To speed up the system, we perform classification only for voxels around the heart surface, which is achieved by automatically segmenting the whole heart from the 3D volume in a preprocessing step. Experiments demonstrate that the proposed learning based vesselness measurement outperforms the conventional Hessian vesselness in both speed and accuracy.

1 Introduction

Cardiovascular diseases are the leading cause of death in western countries and coronary artery disease (CAD) is the most popular among them. Automatic extraction of the coronary artery tree will facilitate physicians in diagnosing CAD. However, coronary artery segmentation is a difficult problem because of the irregularity of coronaries, imaging artifacts, and variations of contrast, etc. Various coronary artery segmentation methods have been proposed. Recently, MICCAI conference held a competition for coronary artery tracking (CAT) [1]. Most of the top performed approaches are based on shortest path computation given one or two end points on the artery. Since the shortest path computation is standard, the major variation is in the different vesselness measurements used for the path cost. The classical method by Frangi *et al*. [2] assumes that a vessel has an approximately tubular structure. The eigenvalues of the Hessian matrix at any point give us a measure of the tubularity around that region. A similar approach has also been proposed by Sato *et al*. [3]. Hessian vesselness measures were used in a few submissions for the MICCAI CAT competition [4–6]. Detailed analysis of the 2D cross section of a vessel provides another family of vesselness measurements. For example, the medialness measurement [4, 7, 8] uses the circularity assumption of the 2D cross section and edge responses obtained from multi-scale filters. Friman *et al*. [9] presumes a model for ideal intensity distribution in a vessel cross section. Intensity should be the greatest at the vessel centerline and exponentially decrease toward the artery wall. Their vesselness measures how well actual intensity values fit the model. Alternatively, Zambal *et al*. [10] used two concentric cylinders for vessel tracking. Intensity is sampled along both the exterior and the interior cylinders. The vesselness response depends on how well the exterior intensity histogram can be separated from the interior. The inherent disadvantage of the existing vesselness measurements is that they introduce strong assumptions about the structure of the coronary. It is clear, however, that coronary arteries are very erratic. To make matters worse, some CT volumes may lack sufficient contrast or contain artifacts. Therefore, an empirically designed measurement is by no means optimal in the use of image context information.

In this paper, a machine learning based vesselness measurement is proposed by exploiting the rich domain specific knowledge embedded in an expert-annotated dataset. We use machine learning techniques to train a classifier that assigns a high score for voxels inside the artery and

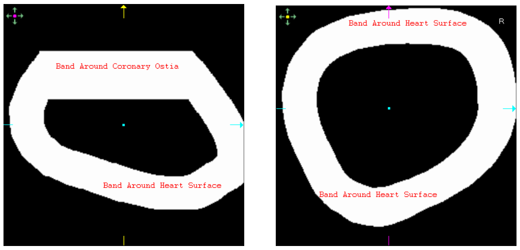

Figure 1. Coronary artery mask (which is used to constrain the detection of coronary artery voxels) is generated from the coronary ostia and the whole heart surface. Two orthogonal views are shown here.

a low score to those outside. The classification score can be treated as a vesselness measurement in the computation of the shortest path. Two kinds of features (geometric and image features) are extracted to train the classifier. For example, as a geometric feature, the position of a voxel in the heart-oriented coordinate system (which is defined by three landmarks, namely, the aortic valve center, the mitral valve center, and the left ventricle endocardium apex) provides a priori knowledge about how likely this voxel lies inside a coronary artery. Normally, a coronary artery is a bright thin structure when contrast agent is applied. Therefore, we also extract a set of image features using the flexible framework of steerable features [11]. To speed up the system, we perform classification only for voxels around the detected coronary ostia and the heart surface, which is achieved by automatically segmenting the whole heart from the 3D volume in a preprocessing step [12]. By constraining the classification, we can also reduce the false positives (a non-artery voxel getting a high vesselness score). Experiments demonstrate the accuracy and efficiency of the proposed approach, compared to the widely used Hessian vesselness [2]. At the same detection rate, it can reduce the false positive rate by a half. It only requires approximately 3.5 seconds to process a large volume (e.g., $512 \times 512 \times 200$ voxels), which is significantly faster than the Hessian vesselness.

2 Learning-Based Vesselness Measurement

In this section, we present our learning-based vesselness measurement. The main goal is to train a classifier that assigns a high score to voxels inside the coronary and a low score to those outside.

2.1 Coronary Artery Mask

Normally, a cardiac CT volume has a high resolution (less than 0.5 mm inside a slice). It often contains approximately $512 \times 512 \times 200$ voxels. As a consequence, it is very time consuming to test every voxel using the classifier. Furthermore, there is a strong constraint on the distribution of the coronary arteries. Starting from the coronary ostia, the arteries emerge quickly to the heart surface. By constraining the voxel classification to a region with high occurrence probability of the coronary arteries, we can significantly reduce the computation time and effectively reduce the false positive rate.

In the preprocessing step of our system, we automatically segment the whole heart from the 3D volume [12] and detect the two coronary ostia [13], too. Our preprocessing step is very fast and takes about 1.5 seconds to finish both tasks. Given the heart surface mesh, we expand it by 5 mm and shrink it by 10 mm to generate a coronary mask. The arteries go deep into the heart to connect to the coronary ostia, they are not close to the heart surface around that region. We add a band of ± 10 mm around the coronary ostia to the mask. The band widths are empirically tuned on a few datasets to make sure all coronary arteries are within the mask. Since we only care about the arteries around the ventricles, we can cut the part above the detected coronary ostia (above of the ostia band). An example of the coronary mask is shown in Fig. 1.

2.2 Vesselness Classifier Training

We collected 54 expert-annotated volumes. Among them, 40 were randomly selected for training and the remaining 14 were reserved for testing. The method takes advantage of the precision of the annotations in the data set. For every volume, the whole coronary tree is annotated. Each annotation of a coronary branch consists of densely placed cross sections of the vessel. The coronary lumen on each cross section is delineated with a polygon. All voxels inside the labeled coronary arteries are regarded as positive training samples. The negative samples are those voxels within the coronary artery mask, but with a distance of at least 5 mm away from the artery. We randomly sample five million negative samples for training.

For each sample, we extract a set of features for the classification purpose. Two groups of features (namely, geometric and image features) are extracted. The geometric relation of the coronary arteries and the heart chambers are well constrained, which is helpful for automatic detection. However, the pose of the heart inside a volume varies a lot. We need to project the sample position into the heart-oriented coordinate system. Our heart coordinate system is defined by three landmarks (namely, the aortic valve center, the mitral valve center, and the left ventricle endocardium apex). As shown in Fig. 2, the z axis is defined as the direction pointing from the aortic valve center to the left ventricle apex. The x axis is defined as the vector perpendicular to the z axis and lies inside the plane formed by three landmarks. The y axis is the cross product of the z and x axes. In the preprocessing step, we need to segment the whole heart surface from the CT volume, which is composed with

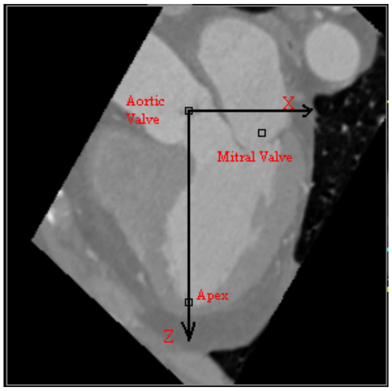

Figure 2. Heart coordinate system defined by the aortic valve center, the mitral valve center, and the left ventricle apex.

two steps [12]. In the first step, we estimate the pose of the heart. After that, a mean shape is aligned with the estimated pose, following by non-rigid boundary refinement. We use the above heart coordinate system to define the pose of the heart. Therefore, the estimate of the heart coordinate system is a by-product of heart segmentation, without demanding any extra computation efforts. Suppose the position of voxel in the heart coordinate system is (X, Y, Z) and the angles between vector (X, Y, Z) and the three axes of the heart coordinate system are α, β, and θ, respectively. The following seven geometric features are extracted: X, Y, Z, $\sqrt{X \times X + Y \times Y + Z \times Z}$, α, β, and θ.

We also extract a set of image features. Steerable features are an efficient framework, which is capable of capturing complicated image patterns [11]. Around a given position, we sample $5 \times 5 \times 5$ points using the regular sampling pattern of the steerable features. The distance between neighboring sampling point is set to 3 mm. On each sampling point, we extract 24 local features based on the image intensity and gradient. To be specific, the following features are extracted. Suppose a sampling point (x, y, z) has intensity I and gradient $g = (g_x, g_y, g_z)$. The three axes of the volume coordinate system are n_x, n_y, and n_z. The angle between the gradient g and the z axis is $\alpha = arccos(n_z.g)$, where $n_z.g$ means the inner product between two vectors n_z and g. The following 24 features are extracted: I, $\sqrt{I}$, $\sqrt[3]{I}$, I^2, I^3, $\log I$, $\|g\|$, $\sqrt{\|g\|}$, $\sqrt[3]{\|g\|}$, $\|g\|^2$, $\|g\|^3$, $\log \|g\|$, α, $\sqrt{\alpha}$, $\sqrt[3]{\alpha}$, α^2, α^3, $\log \alpha$, g_x, g_y, g_z, $n_x.g$, $n_y.g$, $n_z.g$. In total, we have 24 local features for each sampling point. The first six features are based on intensity and the remaining 18 features are transformations of gradients. Feature transformation, a technique often used in pattern classification, is a process through which a new set of features is created. We use it to enhance the feature set by adding a few transformations of an individual feature. In total, we get 125×24 image features.

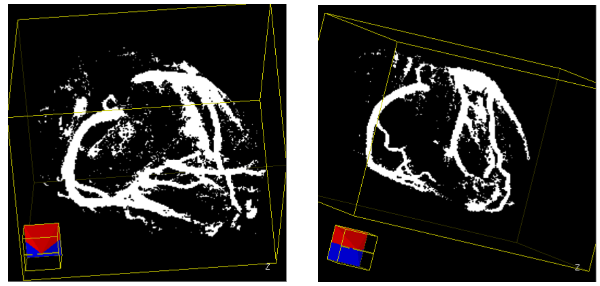

Figure 3. Two examples of the detected coronary arteries.

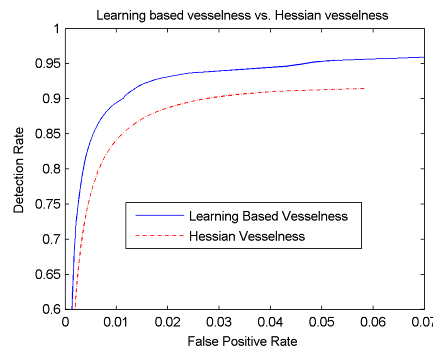

Figure 4. Detection rate vs. false positive rate of the learning based vesselness and Hessian vesselness.

Probabilistic boosting tree (PBT) [14] is used to train a classifier to distinguish voxels inside the coronary arteries from those outside. The PBT classifier is equivalent to a classification tree with an AdaBoost classifier at each node. Cross validation has shown that a tree of level four achieves good balance between speed and classification accuracy. In order to speed up the classification, we want to use as few weak classifiers as possible at the tree root node to reject a significant number of negative samples. Therefore, easy negative samples can be rejected quickly without the need to compute too many features. The training performance goal of the root node is set to achieve a detection rate no less than 99% and a false positive rate no more than 50%. We find an AdaBoost classifier with five features is enough to achieve this accuracy level. Since the classification problem is becoming harder and harder along the tree structure, we gradually increase the number of weak classifiers toward the tree leaves. To be specific, the following parameters are used to train the PBT: 5, 10, 20, and 30 for levels 1 to 4, respectively.

2.3 Vesselness Evaluation

Vesselness evaluation phase consists of two steps. First, the the whole heart is segmented from the volume [12] and the coronary ostia are detected [13]. A coronary artery mask is generated, as described in Section 2.1. Then, voxels inside the mask are evaluated by the PBT classifier. The score output by the PBT classifier (in the range of [0, 1]) is treated as a vesselness measurement. Fig. 3 shows two examples of the detected coronary arteries (using 0.7 as a threshold).

In Fig. 4, we quantitatively compare the proposed learning based vesselness and Hessian vesselness on the test set (14 volumes). The curves show the detection rate vs. false positive rate. Please note that the coronary mask is also applied to the Hessian vesselness to exclude non-artery tissues. Otherwise, the false positive rate of the Hessian vesselness is significantly higher since the pulmonary arteries in the lung are enhanced by the filter too. The Hessian vesselness cannot achieve a high detection rate (no higher than 88%). At the same detection rate, our learning based vesselness can reduce the false positive rate by a half. Average computation time for the voxel classification procedure is about 3.5 seconds on a computer with 3.2 GHz CPU and 3 GB memory. This is also significantly faster than the Hessian vesselness.

3 Conclusion

In this paper, we proposed a learning based vesseleness measurement for coronary artery segmentation. Instead of using an empirically designed vesseleness measurement, we exploit the rich domain-specific knowledge embedded in an expert-annotated dataset. For each voxel, we extract a set of geometric and image features. The probabilistic boosting tree (PBT) is used to train a classifier, which assigns a high score for voxels inside the artery and a low score to those outside. To improve the efficiency of the proposed method, we constrain the detection to voxels inside a coronary artery mask, which is generated based on the whole heart surface and the detected coronary ostia. The system can be applied as pre-filtering tool and a Dijkstra-like approach can be used for extraction the whole coronary tree out of the detected voxels.

References

[1] C. Metz, M. Schaap, T. van Walsum, A. van der Giessen, A. Weustink, N. Mollet, G. Krestin, and W. Niessen, "3D segmentation in the clinic: A grand challenge II — coronary artery tracking," in *MICCAI Workshop on 3D Segmentation in the Clinic: A Grand Challenge*, 2008.

[2] A. Frangi, W. Niessen, K. Vincken, and M. Viergever, "Multiscale vessel enhancement filtering," in *Proc. Int'l Conf. Medical Image Computing and Computer Assisted Intervention*, 1998.

[3] Y. Sato, S. Nakajima, N. Shiraga, H. Atsumi, S. Yoshida, T. Koller, G. Gerig, and R. Kikinis, "Three-dimensional multi-scale line filter for segmentation and visualization of curvilinear structures in medical images," *Medical Image Analysis* **2**, pp. 143–168, 1998.

[4] K. Krissian, H. Bogunovic, J. Pozo, M. Villa-Uriol, and A. Frangi, "Minimally interactive knowledge-based coronary tracking in CTA using a minimal cost path," *The Insight Journal* , 2008.

[5] C. Metz, M. Schaap, T. van Walsum, and W. Niessen, "Two point minimum cost path approach for CTA coronary centerline extraction," *The Insight Journal* , 2008.

[6] Y. Zhang, K. Chen, and S. Wong, "3D interactive centerline extraction," *The Insight Journal* , 2008.

[7] K. Krissian, G. Malandain, N. Ayache, R. Vaillant, and Y. Trousset, "Model based detection of tubular structures in 3D images," *Computer Vision and Image Understanding* **80**(2), pp. 130–171, 2000.

[8] H. Tek, M. A. Gulsun, S. Laguitton, L. Grady, D. Lesage, and G. Funka-Lea, "Automatic coronary tree modeling," *The Insight Journal* , 2008.

[9] O. Friman, C. Khnel, and H. Peitgen, "Coronary artery centerline extraction using multiple hypothesis tracking and minimal paths," *The Insight Journal* , 2008.

[10] S. Zambal, J. Hladuvka, A. Kanitsar, and K. Bhler, "Shape and appearance models for automatic coronary artery tracking," *The Insight Journal* , 2008.

[11] Y. Zheng, A. Barbu, B. Georgescu, M. Scheuering, and D. Comaniciu, "Four-chamber heart modeling and automatic segmentation for 3D cardiac CT volumes using marginal space learning and steerable features," *IEEE Trans. Medical Imaging* **27**(11), pp. 1668–1681, 2008.

[12] Y. Zheng, F. Vega-Higuera, S. K. Zhou, and D. Comaniciu, "Fast and automatic heart isolation in 3D CT volumes: Optimal shape initialization," in *Int'l Workshop on Machine Learning in Medical Imaging (In conjunction with MICCAI)*, 2010.

[13] Y. Zheng, M. John, R. Liao, J. Boese, U. Kirschstein, B. Georgescu, S. K. Zhou, J. Kempfert, T. Walther, G. Brockmann, and D. Comaniciu, "Automatic aorta segmentation and valve landmark detection in C-arm CT: Application to aortic valve implantation," in *Proc. Int'l Conf. Medical Image Computing and Computer Assisted Intervention*, 2010.

[14] Z. Tu, "Probabilistic boosting-tree: Learning discriminative models for classification, recognition, and clustering," in *Proc. Int'l Conf. Computer Vision*, pp. 1589–1596, 2005.