Segmentation and Removal of Pulmonary Arteries, Veins and Left Atrial Appendage for Visualizing Coronary and Bypass Arteries

Hua Zhong[1], Yefeng Zheng[1], Gareth Funka-Lea [1],
and Fernando Vega-Higuera[2]

[1]Siemens Corporate Research
755 College Road East, Princeton, NJ 08540

[2]Siemens Healthcare
91301 Forchheim, Germany

{hua-zhong, yefeng.zheng, gareth.funka-lea, fernando.vega-higuera}@siemens.com

Abstract

In this paper we present an automatic heart segmentation system for helping the diagnosis of the coronary artery diseases (CAD). The goal is to visualize the heart from a cardiac CT image with pulmonary veins, pulmonary arteries and left atrial appendage removed so that doctors can clearly see major coronary artery trees, aorta and bypass arteries if exist. The system combines model-based detection framwork with data-driven post-refinements to create voxel-based heart mask for the visualization. The marginal space learning [6] algorithm is used to detect mesh or landmark models of different heart anatomies in the CT image. Guided by such detected models, local data-driven refinements are added to produce precise boundaries of the heart mask. The system is fully automatic and can process a 3D cardiac CT volume within 5 seconds.

1. Introduction

Coronary Artery Disease (CAD) or Coronary Heart Disease (CHD) is the leading cause of death in the world. It is caused by accumulation of plaque in coronary arteries (CA). Eventually such plaque blocks or reduces blood flow to heart muscles. Deprived of oxygen, myocardium is damaged and other heart diseases develop. The early symptom for CAD/CHD is usually chest pain which can be easily mistaken for other less serious diseases until the patients experience a heart attack.

Computed tomography (CT) is often used for diagnosis and treatment planing for CAD/CHD. In cardiac CT images, not only the heart but also surrounding anatomical structures are imaged and can block the direct view of the heart in a 3D visualization. In the past, algorithms were developed to isolate the heart from surrounding structures like lungs, spine and sternum [9] [2]. However, the heart structures like the Pulmonary Arteries (PA), the Pulmonary Veins (PV) and the Left Atrial Appendage (LAA) still partially occlude the coronary arteries (mainly the left part). It is desirable to remove these three structures from a cardiac CT image visualization so that physicians can easily see the coronary arteries. Figure 1 (a) shows a raw CT scan image. Ribs, sternum, and other structures totally block any view of the heart. It is impossible for physicians to see any CA for diagnosis. Figure 1 (b) is the previous heart isolation result [9]. It can automatically isolate the heart from surrounding structures by detecting the pericardial mesh of the heart. With this result, physicians can easily see many detailed heart structures. However the LAA, the PA and the PV still block the left coronary artery (LCA). Figure 1 (c) is the result of our algorithm described in this paper. The LAA, the PA and the PV are removed and the LCA can be clearly seen without any occlusion. Such a 3D view can greatly help physicians to perform diagnosis of CAD/CHD.

Since these structures are very close to the cornary artery tree and they are all connected to heart chambers we want to keep, pure data-driven algorithms such as region growing cannot segment them cleanly without leaking into nearby chambers or the aorta. While model-based segmentation algorithms, such as the marginal space learning [6], can reliably detect the anatomies based on mesh models. However, there are some limitations. First it works well for anatomies with relatively less variations, like the four chambers, but not highly variable structures like the LAA, and the PA. They can hardly be modeled by a single-part mesh model. Second, note that standard local refinements based on a statistical shape model [1] and mesh smoothing algorithms [3] [4] are used by the algorithm to generate a smoothed mesh result. However such smoothed meshes, when converted to voxel masks, may not cover all the voxels of the detected

978-1-4673-6412-5/12/$31.00 ©2012 IEEE

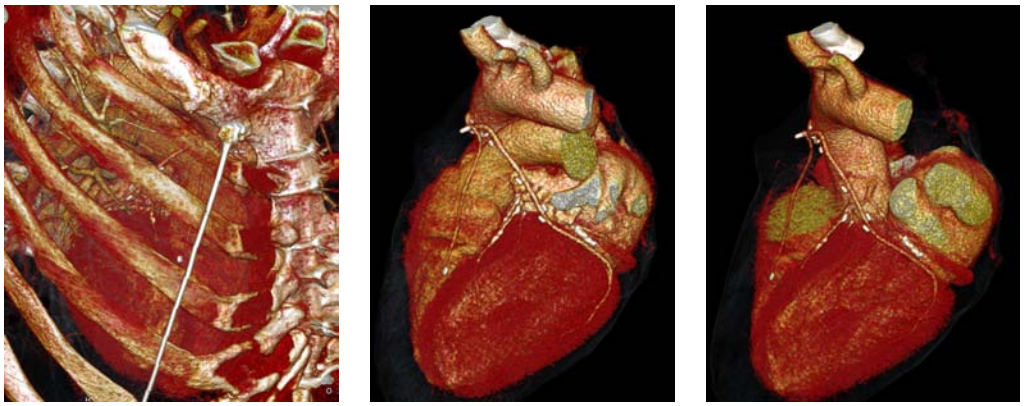

(a) (b) (c)

Figure 1. Heart isolation visualization. (a) The original CT scan. Note that bones blocked any direct view of the heart. (b) The result of the previous algorithm which only isolates the whole heart. Still, the pulmonary artery (PA), the pulmonary veins (PV) and the left atrial appendage (LAA) occlude the left coronary artery (LCA). (c) The result of our system. The PA, the PV and the LAA are removed automatically and the LCA is easily seen. Further more, the plaques that block the LCA can be easily identified. Note that in this case, there are two bypass arteries which are reliably kept intact by the algorithm.

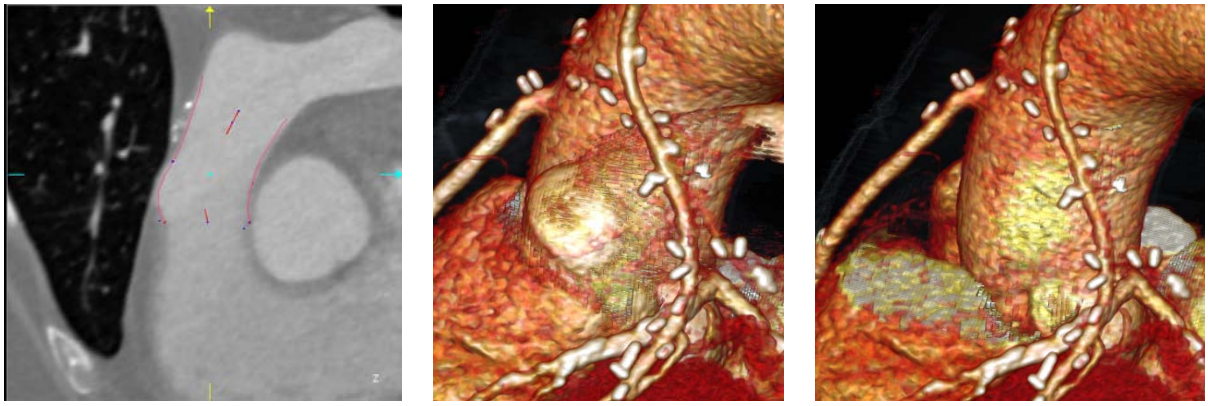

(a) (b) (c)

Figure 2. Directly applying the model-based machine-learning algorithm from [6] on PA root usually result in a thin layer of PA left in the image (b), even though the mesh looks accurate and smooth in (a). That's because the mesh model's resolution cannot capture the voxel-level details of the shape. (c) While our algorithm can create a clean mask of PA for removal.

anatomy and generate visible artifacts (Figure 2).

To overcome the mentioned problems, we combine a local region growing algorithm with the global shape model to solve the PA, the PV and the LAA segmentation problems. The idea is to use a machine learning algorithm to learn a global shape model, either mesh based or fiducial control point based, to locate the approximate location and orientation of the object. Then we use constrained local intensity based region growing algorithms to refine the shape and generate a detailed mask. In order to avoid any removal of the CA or the aorta which are desired to be kept, we also use a model-based algorithm to create "protection" zones for them where no removal is allowed to such protected objects. The result is a fully-automatic, efficient and clean removal of the PA, the PV, and the LAA for 3D visualization of the CA.

2. Methods

Before the segmentation of the PA, the PV and the LAA, we first segment the whole heart using algorithms described in [9]. The result is a pericardial mesh containing four heart chambers, part of the aorta, the PA, the PV and the LAA as well as the CA. Then the aorta root's location is detected using algorithms described in [7]. We also detected the left atrium using [5]. These structures usually can be reliably detected since they have relatively less shape variations and clean boundaries. Furthermore, the segmentation results of these structures can then be used to constrain the segmentation of the PA, the PV and the LAA's locations.

For the PA, the PV and the LAA, we use slightly different segmentation algorithms for each. However, the frameworks of all these algorithms are similar: global shape-

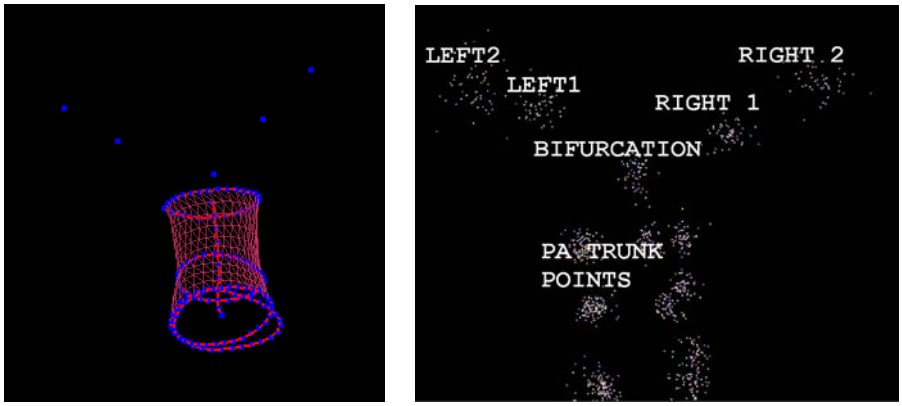

(a) (b)

Figure 3. PA model: (a) the mesh and five fiducial point model. (b) the statistical shape model for detecting the fiducial points (bifurcation, left 1 and 2, right 1 and 2). Based on 120 manually labeled data, we select nine points from the PA trunk mesh and combine them together with the five PA fiducial points to create a statistical shape model.

model detected by the marginal space learning algorithm [6] and local refinement based on the statistics of intensities.

2.1. Globe Shape Segmentation

2.1.1 The Pulmonary Artery Model

The PA trunk root: the portion of the PA from the right ventricle (RV) to the bifurcation, is modeled as a tubular mesh. From the bifurcation, it is difficult to approximate the shape with a tube. In this case, we use five fiducial control points: one at the bifurcation, two at the left PA branch and two at the right PA branch as shown in Figure 3 (a). We first describe how the PA trunk mesh is detected.

For the PA trunk mesh, we use the segmentation algorithm of [6]. The shape model, the bounding box detector and the boundary detector are trained with 320 manually annotated volume data. The detector returns the probability of the image I given PA trunk bounding box at the combination c of location, orientation and scale:

$$P(I|c) = P(I|x, y, z, \alpha, \beta, \gamma, s) \tag{1}$$

where (x, y, z) is the location of the bounding box, (α, β, γ) is the oritentation of the bounding box and (s) is the scale. The final boudning box is the one that maximize the probability of:

$$P(c|I) = P(I|c)P(c)/P(I) \tag{2}$$

P(I) is defined as a constant. The prior probability of $P(c)$ can be learned from annotated data. We define it as the range of the location, the orientation and the scale of all PA trunk bounding boxes we have seen in annotated training data. If a bounding box is within the range, its $P(c) = 1$, otherwise its $P(c) = 0$. However, with different scanning protocols, the arbitrary location and orientation of the PA trunk in an image can vary a lot. Thus the range is very large and it allows the bounding box to be almost anywhere with any orientations. Then during the detection, some false positives happened. For example, the aorta root is detected as the PA trunk since they look similar to each other.

To reduce such errors, we use the detected AO root and LA bounding boxes to constrain the detection of the PA trunk bounding box. We modify the equation 2 to:

$$P(c_{PA}, c_{AO}, c_{LA}|I) = P(c_{PA}|c_{AO}, c_{LA}, I))P(c_{AO}, c_{LA}|I) \tag{3}$$

Since the AO and the LA are already detected, $P(c_{AO}, c_{LA}|I)$ is constant. And the problem becomes maximizing the probability of:

$$P(c_{PA}|c_{AO}, c_{LA}, I) \tag{4}$$

with Bayes' theorem, it is equal to maximizing the likelihood of:

$$P(I|c_{PA}, c_{AO}, c_{LA})P(c_{PA}, c_{AO}, c_{LA}) \tag{5}$$

Again, the AO and the PA are already detected and fixed, the problem can be further simplified as maximizing the probability of:

$$P(I|c_{PA}, c_{AO}, c_{LA})P(c_{PA}|c_{AO}, c_{LA}) \tag{6}$$

The likelihood of $P(I|c_{PA}, c_{AO}, c_{LA})$ is calculated by the bounding box detectors for the three anatomies. The conditional prior of $P(c_{PA}|c_{AO}, c_{LA})$ is defined as the constraint by the detected aorta (AO) and LA bounding boxes. It is learned from training database with annotated AO, LA and PA trunk bounding boxes. Different from the prior we mentioned above, this constraint is defined as the relative location, orientation and scale range of the PA to the LA and the AO. This relative constraint gives us a much tighter

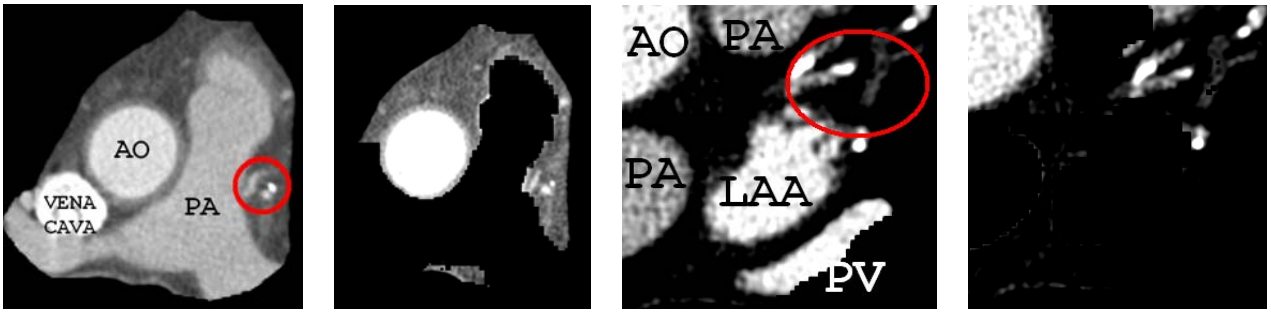

(a) (b) (c) (d)

Figure 4. Voxel-based refinement for the PA, the PV, and the LAA. (a) Before removal, the bypass arteries are highlighted by the red circle. (b) the PA and the vena cava are removed by region growing while bypass right adjacent to PA is kept intact. (c) before removal, we can see the small isolated chambers of the LAA very close to the coronary arteries highlighted by the red circle. (d) the LAA, the PA and the PV are removed cleanly while the CA is intact.

range and thus eliminated a lot of false positives during the detection. The result then is much more robust.

After the bounding box is detected, a statistical shape model is fit within it. Five fiducial points on the left and right PA branches are trained with 120 manually annotated volumes. The detection of the fiducial points is a mixture of a statistical shape model and individual fiducial point detectors using the marginal space learning (MSL) algorithm [6]. The reason for this mixture is that in many images, the PA fiducial points' locations are not inside, or are very close to the image borders. Thus, the MSL-based bounding box detector may fail in these cases since it relies on image features which are not available outside the image. However, the statistical shape model can handle this out-of-boundary situations well. In our method, we build a statistical shape model [1] containing nine PA trunk points selected from the PA trunk mesh and the five PA fiducial points: bifurcation, left 1 and 2, right 1 and 2 as shown in Figure 3 (b). When the PA trunk is detected, we extract the nine PA trunk points from the detected mesh. Then we use the statistical shape model to estimate the optimal location of the five PA fiducial points given the nine PA trunk points' locations. The statistical shape model can estimate the location of a fiducial point even it is outside the volume. We select only nine PA trunk points instead of all the mesh points because we want the statistical shape model to capture variations for both the PA trunk and the left and right PA branches in a balanced way. If all the PA trunk mesh points are included, the statistical shape model will be dominated by the shape variations of the PA trunk, and makes the estimation of the left and right PA less accurate. We found that with nine PA trunk mesh points the algorithm works very well for our purpose. Next we use the learned MSL detectors to refine each of the five estimated PA fiducial points. The MSL detectors will only search a small neighborhood around the current estimated locations thus it is reliable and fast. If MSL detectors failed because a fiducial point is close to or out of the image border, the statistical shape model result will be used as final detection result. Otherwise the MSL detector's result will be used.

2.1.2 The Pulmonary Vein Model

The PV's shape varies too much to be represented by a single mesh model. Instead, we use two fiducial points defined on the detected left atrium (LA) mesh model to locate the root of the left and the right pulmonary veins. In practice they are defined as two specified vertices on the LA mesh. The detailed mask for PV's is handled by a region growing method described in the next section.

2.1.3 The Left Atrial Appendage Model

We model the LAA using the same mesh model as a heart chamber. The mesh is designed to capture the outer boundary of LAA. However the LAA's shape varies much more than any heart chambers both for its topology and size. This mesh model usually cannot capture the exact boundary of a LAA. Instead, we only use this model to locate the LAA's bounding box so that the exact boundary can be segmented using the intensity based refinement described in the next section.

2.2. Local Voxel-based Refinement

As we have stated before, the global shape model usually cannot generate the exact voxel mask for the PA, PV and LAA. A local refinement is necessary for our heart isolation application. For the PA, PV and LAA, we use different refinement strategies. However the goals are the same: to find clear boundaries without cutting into any of the CA or bypass.

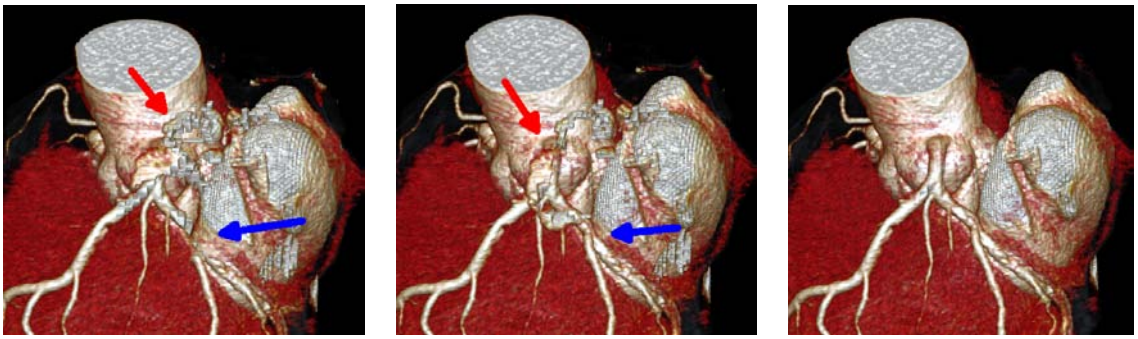

(a) (b) (c)

Figure 5. LAA removal: (a) With LAA mesh model only: some LCA is cut (blue arrow) while some LAA is not removed (red arrow). (b) First pass of connected component analysis: only the largest connected region in the bounding box of LAA mesh is removed. LCA is intact (blue arrow) however still some small isolated regions of LAA remain (red arrow). (c) Second pass of connected component analysis: run in the whole image and keep the largest piece while any small isolated pieces within the LAA bounding box are removed. The result is a clean removal of all LAA voxels.

2.2.1 The Pulmonary Artery

For the PA, the global shape model contains two parts: the PA trunk mesh and the five fiducial points. For the mesh, we first close its openings and then mask out any voxels inside the mesh. As shown in Figure 2, usually a thin layer of the PA trunk still remains due to the mesh smoothing. We then use the region growing algorithm to dilate the mask outward for 2-3 millimeters depending on the image resolution. The region growing algorithm's threshold is determined by the mean and standard deviation of the voxels which are already in the mask. With such statistics, the region growing can work for images with or without contrast agent in the PA to successfully remove the thin layer left by the PA trunk mesh. For the left PA, right PA and the PA bifurcation regions, we start region growing from each of the five fiducial points. In this step, the range for region growing is limited to 15mm since the PA fiducial points are defined as less than 15mm apart from each other. The region growing from the fiducial points thus can cover all the voxels of the PA bifurcation and the two branches. However, it tends to leak into surrounding objects as it only relies on the local information, especially to nearby bypass arteries or LCA. To prevent such "leaks" and protect important structures, we use a model-based method which will be described later.

2.2.2 The Pulmonary Vein

For the PV, we apply the same region growing algorithm from the two root fiducial points of left and right PV as for the PA fiducial points. The intensity threshold is based on the statistics of voxel intensities within the detected LA mesh model. To prevent leakage into nearby structures, we limit the growing range to 25mm.

2.2.3 The Left Atrial Appendage

The LAA is more complex. First, the LAA mesh model only gives an approximate boundary: it may not cover the whole LAA and it may include some LCA or other structures. This is due to the high variation of the LAA's shape. Second, there usually are many small chambers in the LAA which make the LAA not look like a single connected region in the image. To deal with these challenges, we design an algorithm composed with model-based mesh detection and connected component analysis (CCA). The algorithm consists of 3 steps:

1. The LAA mesh is detected. It gives us an initial estimation of the LAA's location and shape. Then we create a bounding box slightly larger than the mesh to make sure we cover the whole LAA regions as the LAA mesh may be smaller than the exact LAA region.

2. The first CCA pass is run *within the bounding box* and the largest connected region found is removed. We assume it is the largest chamber of the LAA. However smaller isolated chambers still remain and they are difficult to be separated from LCA pieces within the bounding box.

3. The second CCA pass is run on the *whole image*. The LCA pieces in the LAA bounding box in this pass should be connected to the whole LCA tree and eventually to the aorta and LV. Thus they should form the largest connected region. While the remaining LAA pieces will form smaller isolated regions. We then remove any such small regions *only within* the LAA bounding box.

This three-step algorithm is illustrated in Figure 5 and as it shows, the algorithm can successfully generate a clean mask of the LAA.

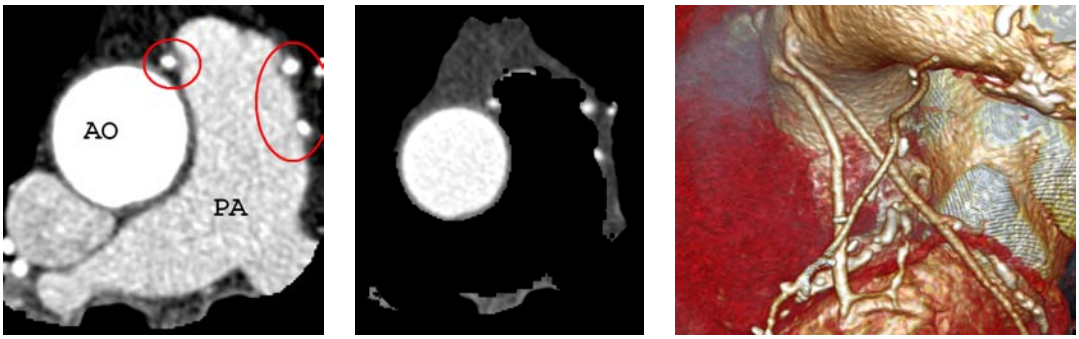

(a) (b) (c)

Figure 6. Protection of vessels while removing the PA. (a) One case where bypass is deeply embedded in the PA as shown with the red circles. (b) Region growing constrained by vessel classification can reliably remove any voxels belong to the PA while keep the bypass arteries' voxels untouched. Also the aorta is protected by segmentation. (c) 3D visualization of the case, the bypass arteries are intact and clearly visible for physicians.

2.3. Chamber and Vessel Protection

Sometimes pieces of the important structures such as the AO, the LA or the CA are removed by the leakage of the region growing algorithms because of similar voxel intensities of them to the PA, the PV or the LAA. To prevent this, we introduce several measures to protect these structures. First we use the segmentation results of the aorta and the LA to mask them as "not possible to grow". The region growing algorithms for the PV and PA and the connected component analysis for the LAA then will ignore any such regions. It is more difficult to protect the CA and the bypass arteries since they are small and usually very close to the PA, the PV and the LAA. Furthermore, we do not have a clean mask of the CA tree as we have for the AO and the LA. Here we use a machine-learning based vessel protection algorithm.

As described in [8], the idea is to train a voxel classifier based on image context to tell the probability of the voxel being in a vessel. This algorithm is capable of quickly classifying a voxel to be vessel or not by applying a threshold to the returned vessel probability. We found that a threshold equal to 75% works well for our purpose. However, there would be a lot of waste of computation power if we classified all voxels in an image. Instead, we confine the classification to only those voxels around the PA trunk, the LAA and the PV where cutting of the CA or bypass arteries by the region-growing or CCA algorithms could happen.

For arteries around the PA trunk, any voxels within 3mm to the PA trunk mesh will be classified for vessel. For regions around the LAA and the PV, usually only the LCA maybe cut. In order to efficiently identify the LCA region around LAA and PV, we build a similar fiducial point model as the PA trunk: it contains the left coronary ostium point, the point where left main (LM) coronary artery bifurcated into LAD and LCX, 20 control points along LCX and 20 selected points from the LA mesh. Then we train a statistical shape model for these 42 points based on a manually labeled training database. During the detection, the 20 LA points from the detected LA mesh, the detected left coronary ostium and the bifurcation point (using [7]) are used to estimate the positions of the 20 LCX points based on the learned statistical model. Then we run vessel classification around the region of this estimated LCA control points.

In our tests, the vessel classification in the regions described above takes only 0.02 seconds with a 4-core Xeon 2.53GHz CPU. After the vessel classification, we can create a vessel protection mask. This vessel protection method can preserve LCA, RCA and bypass very well in our application, as shown in Figure 6.

3. Test Results

The goal of the algorithm is to remove most of the LAA, PA and PV so that the coronary arteries and bypass can be clearly seen in 3D visualization. The removal should not touch any coronary arteries or bypass arteries. The algorithm is tested on a database containing around 120 cardiac CT images and most of them are bypass cases. Then the result is visually examined by experienced testers and a score of 1-5 is given for each case:

1. Major CA cut or bypass cut, important structures removed, considered as failed

2. Large piece of the PA, the PV or the LAA may not be removed but not blocking the CA, some minor shave or cut on the CA or the bypass arteries, considered as acceptable

3. The PA, the PV and the LAA are largely removed with only small pieces left, no cut on CA or bypass, considered as good

4. Only very little of the PA, the PV or the LAA voxels left, no cut on CA or bypass, considered as very good

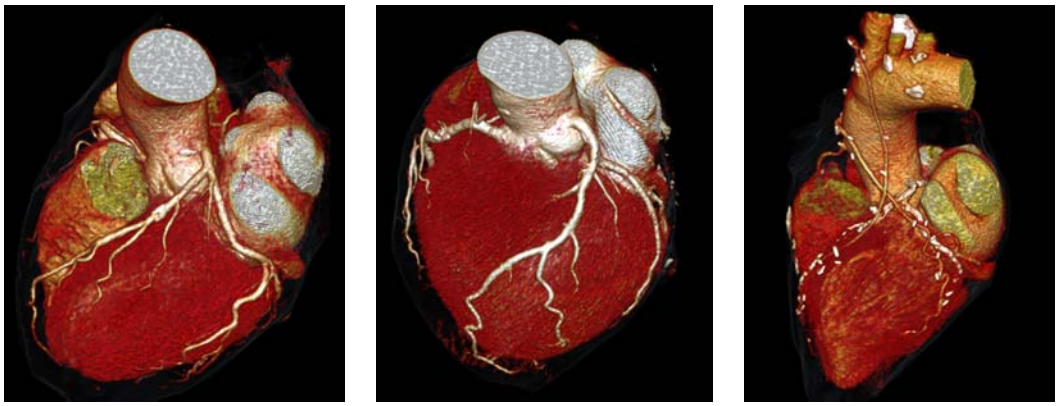

(a) (b) (c)

Figure 7. Some results of our algorithm. (a) and (b) are normal cases and (c) is a bypass cases.

Table 1. Testing score of our algorithm over test dataset. Score 1 is failed, 2 is acceptable, 3 is good, 4 is very good and 5 is perfect. Our algorithm achieved an average score of 3.73.

Score	1	2	3	4	5
Percentage	0.00%	13.33%	13.33%	60.00%	13.33%

5. Clean mask of the heart with the PA, the PV and the LAA totally removed, no CA or bypass cut. Perfect

A score of 3 is thought to be useful, a score of 4 is very good and 5 is perfect. Anything less than 2 is not useful and thought to be failed. Our algorithm's average score is 3.73 and there is no failed case. The distribution of scores is shown in Table 1. Some examples of our result images are shown in Figure 7. We tested the speed on 80 cardiac CT scans. The size of the scans are from $512 \times 512 \times 419$ to $512 \times 512 \times 667$. Resolution of the scans is around $0.4mm \times 0.4mm \times 0.4mm$. The longest processing time is less than 5 seconds with a 2.53GHz Xeon E5630 CPU.

4. Conclusion

In this paper we presented an algorithm that can reliably remove the PA, the PV and the LAA for 3D visualization of coronary arteries for physicians. The system combines global shape model based on machine learning algorithms with local intensity based region growing to segment the structures in voxels. We also provide important structure protection mechanisms to make sure no coronary or bypass arteries are cut during the removal. The system is fully automatic. The test results demonstrate that this algorithm can achieve the goal well, and thus is useful for CAD/CHD diagnosis and treatment planning.

References

[1] T. F. Cootes, C. J. Taylor, D. H. Cooper, and J. Graham. Active shape models—Their training and application. *Computer Vision and Image Understanding*, 61(1):38–59, 1995.

[2] G. Funka-Lea, Y. Boykov, C. Florin, M.-P. Jolly, R. Moreau-Gobard, R. Ramaraj, and D. Rinck. Automatic heart isolation for CT coronary visualization using graph-cuts. In *Proc. IEEE Int'l Sym. Biomedical Imaging*, pages 614–617, 2006.

[3] G. Taubin. Curve and surface smoothing without shrinkage. In *Proc. Int'l Conf. Computer Vision*, pages 852–857, 1995.

[4] G. Taubin. Optimal surface smoothing as filter design. In *Proc. European Conf. Computer Vision*, pages 283–292, 1996.

[5] Y. Zheng, A. Barbu, B. Georgescu, M. Scheuering, and D. Comaniciu. Fast automatic heart chamber segmentation from 3D CT data using marginal space learning and steerable features. In *Proc. Int'l Conf. Computer Vision*, 2007.

[6] 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):1668–1681, 2008.

[7] 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*, pages 1–8, 2010.

[8] Y. Zheng, M. Loziczonek, B. Georgescu, S. K. Zhou, F. Vega-Higuera, and D. Comaniciu. Machine learning based vesselness measurement for coronary artery segmentation in cardiac CT volumes. In *Proc. of SPIE Medical Imaging*, pages 1–12, 2011.

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