

Vascular Landmark Detection in 3D CT Data

David Liu^a, S. Kevin Zhou^a, Dominik Bernhardt^b, Dorin Comaniciu^a

^aSiemens Corporate Research, Princeton, NJ, U.S.A.

^bSiemens Healthcare, Siemensstr. 1, 91301 Forchheim, Germany

ABSTRACT

This work presents novel methods to accurately placing landmarks inside the vessel lumen. This task is an important prerequisite to automatic centerline tracing. Methods have been proposed in the past to determine the location of organ landmarks, and yet several challenges remain for vascular landmarks. First, placing landmarks inside the lumen could be challenging for narrow vessels. Second, contrast-enhanced arteries could be tightly surrounded by bones with similar intensity profiles, making detection difficult compared to arteries surrounded only by darker tissues. Third, landmarks not located at bifurcations could be ill-defined as they have high uncertainty in position.

We first present a method to detect landmarks that are located at vessel bifurcations. Such landmarks have well-defined positions, and we detect them using machine learning techniques. We then present a method to detect vascular landmarks not located at bifurcations. First, a segment detector is created to detect a vessel segment. Annotating multiple points along a vessel segment is easier than annotating a single landmark position, as there is no well-defined position along a vessel. This resolves the ambiguity issue mentioned above. Second, spatial features are computed from the segment detector's response map, and a regression model is created which takes as input the local spatial features surrounding a voxel, and outputs a confidence score of how likely this voxel is inside the lumen. We evaluate the system on a set of 94 3D CT datasets.

Keywords: Landmark Detection, Centerline Tracing, Vessel Landmark, Vertebralis Artery, Carotid Artery, Iliac Artery

1. DESCRIPTION OF PURPOSE

For the assessment of coronary artery disease (CAD) and peripheral artery disease (PAD), automatic extraction of vessel centerlines is a crucial technology. In the most common approach, two seed points are manually placed in the vessel and the centerline is automatically computed between these points. A comprehensive overview of centerline extraction methods is given by Kirbas and Quek.¹ With automatic vessel landmark detection, the seed points can be placed automatically, thus making the centerline extraction workflow fully automatic.

In this paper we present two methods for accurately placing landmarks inside the vessel lumen. The first method is designed for landmarks defined at vessel bifurcations, for example, a landmark at the common-internal-external carotid artery bifurcation. The second method is designed for landmarks where no bifurcation exists nearby, for example, the vertebral artery at the C1 vertebra level. Detecting such landmarks is an important prerequisite to automatic centerline tracing methods that require human or machines to indicate the start and/or end points of a vessel to be traced.

Vascular landmark detection poses several challenges. First, placing landmarks inside the lumen could be challenging for narrow vessels. Second, contrast-enhanced arteries could be tightly surrounded by bones with similar intensity profiles, making detection difficult compared to arteries surrounded only by darker tissues. Third, landmarks not located at bifurcations could be ill-defined as they have high uncertainty in position.

Further author information: (Send correspondence to David Liu.)
E-mail: david-liu@siemens.com, Telephone: 1 609 734 5000

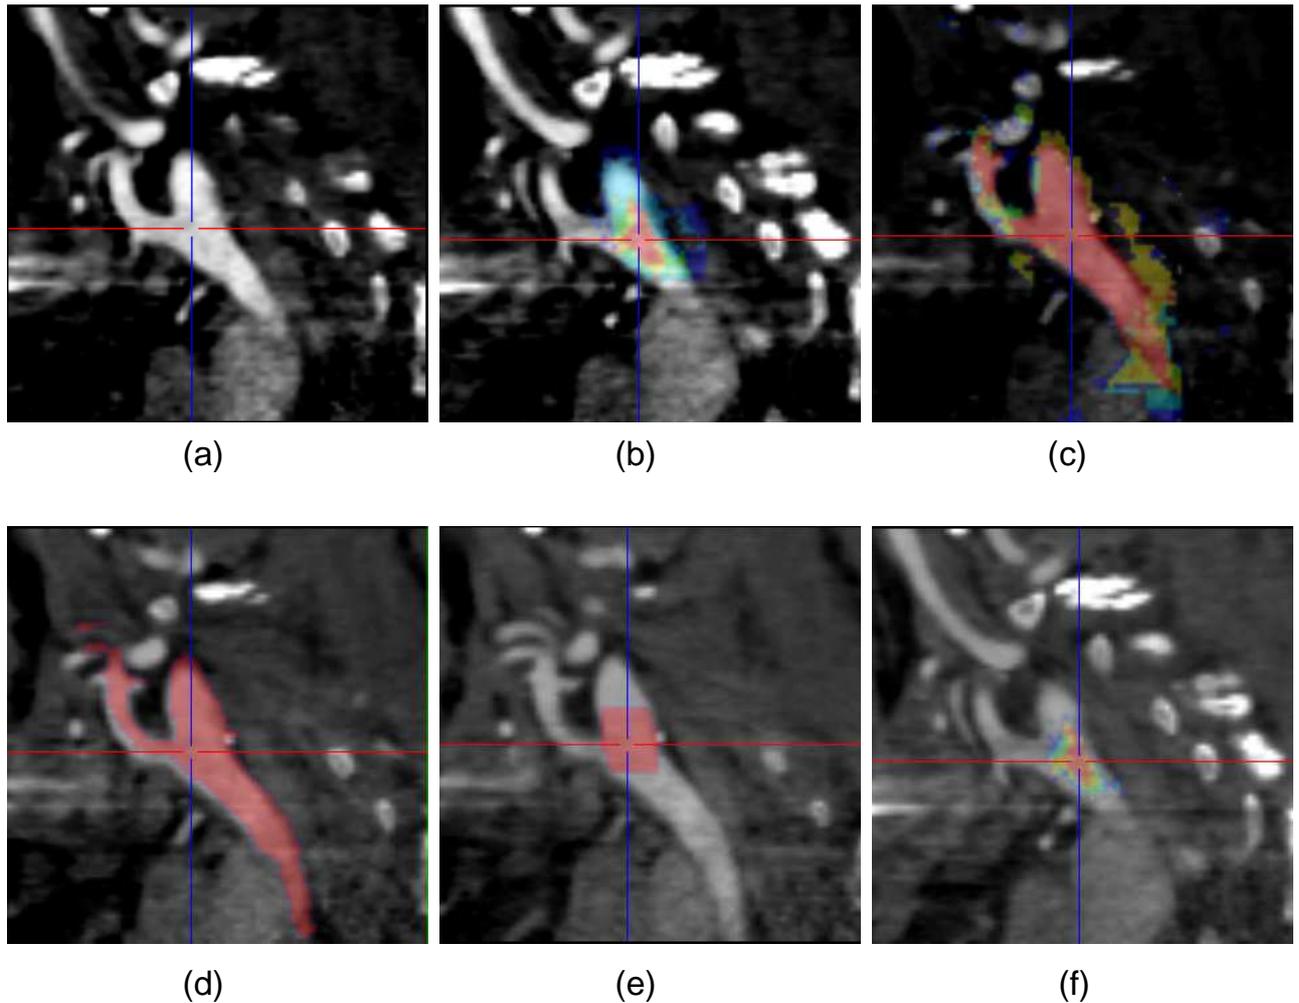


Figure 1. Detecting the carotid artery bifurcation. (a) Cropped input image. (b) Output of bifurcation detector. (c) Output of vesselness detector. (d) Output of graphcuts segmentation. (e) Output of morphological processing. (f) Weighted sum of all probability maps.

2. METHOD FOR DETECTING LANDMARKS AT BIFURCATIONS

Vessel landmark positions are in general ill-defined unless such landmarks are (1) adjacent to some non-vascular structure, such as a bone or an organ, or (2) defined at a vessel bifurcation at a certain body region. The bifurcation of the common carotid artery into the internal and external carotid arteries exemplifies such landmarks. Another example is the bifurcation of the descending aorta into the left and right iliac arteries. In the following, we use the carotid artery as the running example. Let us define the landmark of interest as L , and its position as P_L .

2.1 Multiscale Bifurcation Detector

We train a bifurcation detector with features extracted at multiple resolutions (16, 8, 4, 2mm) to quickly localize the carotid landmark to within about 50mm. The bifurcation detector uses standard machine learning techniques, AdaBoost² and Probabilistic Boosting Tree,³ together with Haar-like features.² Running the bifurcation detector on each voxel in a volume yields a response map R_{bif} of the same size as the volume. An example is shown in Figure 1(b). The response map will be used for evidence aggregation, as described in a later section.

2.2 Vesselness Detector

We annotate voxels inside the common carotid artery from 150 mm below the bifurcation up to the bifurcation. This annotation provides the training data for building a vesselness detector. Running the vesselness detector on a volume yields a response map R_{ves} which has strong response on the artery segment that is within the region that one annotates, lower response above the bifurcation where annotation is not provided, and much lower response outside the vessel. An example is shown in Figure 1(c).

The output of the vesselness detector does not directly indicate the position of the bifurcation, but it is useful because a) the appearance variation of the carotid artery varies more significantly above the bifurcation, hence we are forcing the vesselness detector to focus on the simpler task, that is, to detect the part of the artery with less variations; b) once the part of the artery below the bifurcation is detected, namely, the common carotid artery, we use it to estimate the grayscale intensity or Hounsfield unit of the artery. The grayscale intensity is utilized in the next section for segmentation.

It might appear to the reader that the vesselness detector is redundant when the bifurcation detector already exists. In practice, since the bifurcation detector does not always return accurate response maps, the additional evidence obtained from the vesselness detector could provide a more accurate solution. This is especially true in pathological cases where the bifurcation appears differently from normal cases, such as in carotid stenosis.

2.3 Segmentation

Intuitively, if the intensity of the carotid artery is different from its surrounding, one could segment the carotid artery and then analyze the bifurcation position. Our algorithm selects voxels that have a grayscale intensity close to the estimated grayscale intensity of the common carotid artery (as done in the previous section), and uses those voxels as seeds to initialize a binary segmentation. The binary segmentation intends to assign a binary ON to voxels that belong to the carotid artery, and a binary OFF to the other voxels. We use the graphcuts algorithm⁴ to produce the segmentation R_{seg} . The graphcuts algorithm solves a Markov Random Field that includes a smoothness term, hence producing smoother and more accurate segmentations than simple thresholding. An example is shown in Figure 1(d).

2.4 Skeletonization

On the segmentation map R_{seg} , a morphological thinning operation generates a skeleton that forks at the carotid bifurcation. Ideally, the estimated position of the carotid bifurcation overlaps with the position of the skeleton bifurcation. To fuse this estimated position with the maps obtained from the previous steps, a map R_{ske} is generated by assigning a kernel function (such as a Gaussian function) that centers at the estimated position. An example is shown in Figure 1(e).

The response maps obtained from all the previous steps are weighted and summed up to produce a final score map R as follows:

$$R = w_{bif}R_{bif} + w_{ves}R_{ves} + w_{seg}R_{seg} + w_{ske}R_{ske} \quad (1)$$

The score map typically consists of multiple modes. An example is shown in Figure 1(f). The mode with the highest score is determined as P_L , namely the landmark of interest.

3. METHOD FOR DETECTING LANDMARKS NOT AT BIFURCATIONS

Vessel landmarks not at bifurcations can also be defined when they are adjacent to well-defined bone structures. The vertebral artery has no bifurcations at the C1 to C6 vertebrae levels. In order to detect a landmark within the lumen at, say, the C1 vertebra level, a method is proposed as follows.

3.1 Segment of Interest

A vessel segment of interest S is defined as a portion of the vessel where L is intended to be found. In the following we use the segment at the C1 vertebra level as an example. This segment is illustrated in Figure 2.

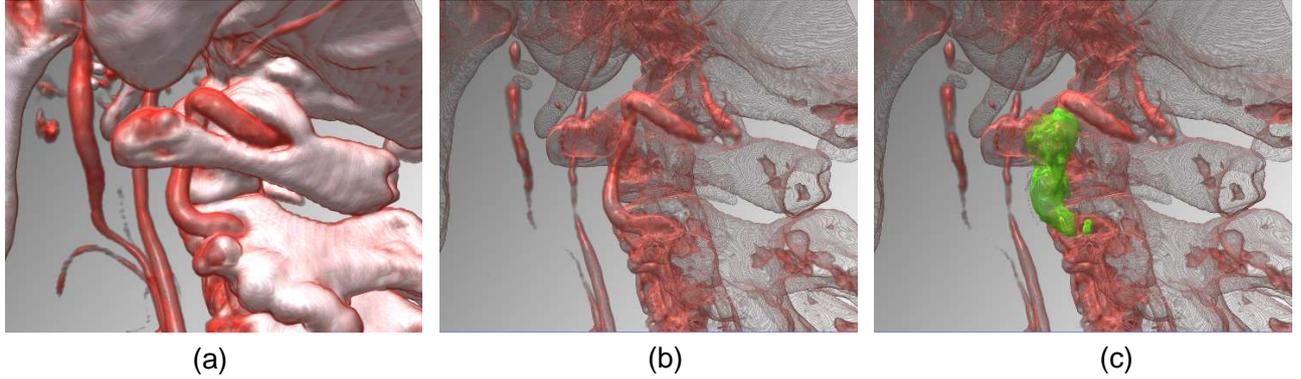


Figure 2. Vertebral artery at C1 vertebra level passes through holes of the vertebra. The same data is shown in (a) and (b) with different transfer functions for visualization. (c) 3D Response map of segment detector shown in green color.

3.2 Data Preparation

We split the training data into two sets. The first set of volumes (Train-Set-1) is used for training a segment detector, and the second set (Train-Set-2) is used for learning a regression model. The purpose of maintaining separate training datasets is to avoid overfitting.

3.3 Segment Detector

We build a detector for detecting the segment of interest. We use standard machine learning techniques, AdaBoost and Probabilistic Boosting Tree, together with Haar-like features in the same way we trained the bifurcation detector in the previous section.

3.4 Segment Response Map

After a detector is trained using Train-Set-1, we run the detector on each volume in Train-Set-2 and store the response at each voxel position as a 3D response map (image) R . Examples are shown in Figure 2(c). We see that the response is stronger at the segment of interest, and weaker at other positions. The response could contain multiple peaks. This is better observed in 2D slices of R , as shown in Figure 3. In the following, we aim at determining a landmark inside the lumen from R .

3.5 Regression Model

It is not trivial to determine P_L when R is multi-modal (has multiple peaks). This suggests that P_L should be placed at a 2D slice where the response is strong and unimodal (has a single peak). Simply smoothing the response map and counting the number of peaks is not robust due to noise. Also, there often exists multiple slices with unimodal responses, and determining the single best one from them is not reliable enough, according to our experiments. Therefore we use a statistical regression model to find P_L . This method relies on a set of annotations. A human annotator examines the response maps in 2D slices as shown in Figure 3. If the response is unimodal on a 2D slice and if the peak falls inside the lumen, then this slice is labeled as positive (Figure 3(c) lower row); otherwise as negative (Figure 3(c) upper row). These labels, together with some features extracted from the map as will be discussed in Section 3.5.1, are used in a regression model. We adopt multiple linear regression using least squares to fit the ground truth labels and features of the response maps.

$$y = f(\mathbf{X}, \beta) = \beta_1 X_1 + \dots + \beta_B X_B \quad (2)$$

where B is the number of features. During testing, the voxel position that yields the highest value from the regression function is determined as P_L . We also experimented with logistic regression and ridge regression and noticed that linear regression gave satisfactory results. Features X_i are explained next.

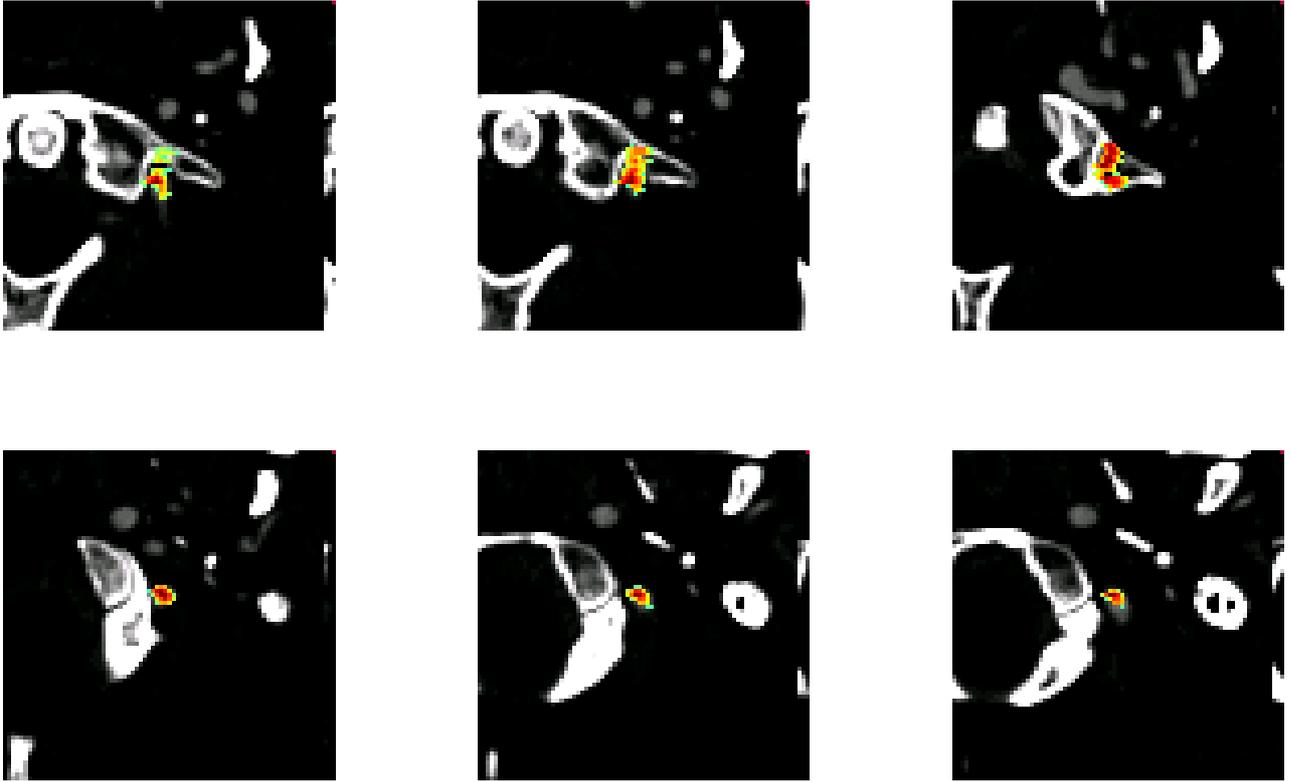


Figure 3. Response maps viewed in 2D slices. The responses are multi-modal in the first row, and unimodal in the second row. A regression model determines the landmark position based on features extracted from the response maps.

3.5.1 Map Features

Features are computed around local maximum on the response maps along 2D slices, as shown in Figure 4(a). Such features record the shape of the response. The space is divided by a log-polar mapping. Each bin records the mean and maximum response falling within that spatial bin. Intuitively, if the response map has high response at certain position, the corresponding bin has high feature values. To reduce overfitting when training data is limited, we use the simplified mapping in Figure 4(b) where only radial distance is considered. Each radial bin stores the mean and maximum values of the responses falling in that bin.

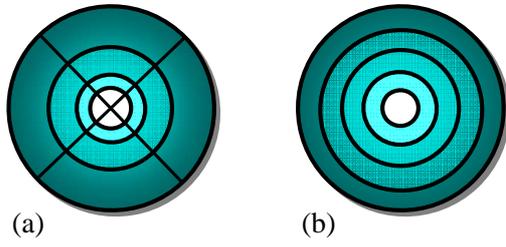
4. RESULTS

To evaluate the system, a set of 94 datasets has been tested. Figure 4(c) shows the detection accuracy of landmarks at the iliac artery bifurcation, aortic valve, aortic arch, internal carotid at cervical-petrous, and renal artery bifurcation. An error is declared when the detected position is outside the lumen. Using the detector in Section 2.1 alone already achieves satisfactory accuracy.

For the internal-external carotid bifurcation, however, using the bifurcation detector alone yields many errors. This is because the carotid artery is much thinner than the arteries mentioned above. As shown in Figure 4(d), our proposed method in Section 2 clearly outperforms using the bifurcation detector alone (shown as Baseline).

For the vertebral artery, Figure 4(e) compares the proposed method in Section 3 with a baseline method in which a single landmark point is annotated per volume and a standard boosting classifier is trained to detect the point. The proposed method clearly outperforms the standard classification technique.

As an application of our landmark detection system, Figure 4(f) shows the result of an automatic centerline tracing result that performs shortest path planning given the start and end position. In this example, the



	FPR	FNR
Iliac	0%	0%
Valve	1.1%	4.8%
Arch	1.1%	4.5%
Cervical-Petrous	5.3%	0%
Renal	1.1%	4.3%

(c)

Baseline		Proposed	
FPR	FNR	FPR	FNR
2.1%	27.4%	1.1%	12.9%

(d)

Baseline		Proposed	
FPR	FNR	FPR	FNR
30.7%	4.2%	2.1%	4.2%

(e)



(f)

Figure 4. (a)(b) Feature extraction from the response map. Each bin records the mean and maximum values from the response maps. The features are used in a regression model. (c)(d)(e) Detection results of different arteries, see text for details. (f) Automatic centerline traced from detected landmarks.

5. CONCLUSIONS

Our proposed methods can place a landmark accurately inside the artery lumen. We demonstrated the feasibility on several important arteries of different nature, including the iliac artery, which is thick and has a well-defined bifurcation; the carotid artery, which is narrow and has a well-defined bifurcation; and the vertebral artery, which is narrow and has no well-defined position.

REFERENCES

- [1] Kirbas, C. and Quek, F., “A review of vessel extraction techniques and algorithms,” *ACM Computing Surveys* **36**, 81121 (2004).
- [2] Viola, P. and Jones, M., “Robust real-time face detection,” *Intl. Journal of Computer Vision* **57**, 137–154 (2004).
- [3] Tu, Z., “Probabilistic boosting-tree: Learning discriminative models for classification, recognition, and clustering,” *IEEE Intl. Conf. Computer Vision* , 1589–1596 (2005).
- [4] Boykov, Y. and Jolly, M.-P., “Interactive graph cuts for optimal boundary and region segmentation of objects in n-d images,” *IEEE Intl. Conf. Computer Vision* **1**, 105 – 112 (2001).