Vessel-guided airway segmentation based on voxel classification

Pechin Lo[1], Jon Sporring[1], Haseem Ashraf[2], Jesper Johannes Holst Pedersen[3], and Marleen de Bruijne[1,4]

[1] Image Group, Department of Computer Science, University of Copenhagen, Denmark, pechin@diku.dk,
[2] Department of Respiratory Medicine, Gentofte University Hospital, Denmark,
[3] Department of Thoracic Surgery, Gentofte University Hospital, Denmark,
[4] Biomedical Imaging Group Rotterdam, Departments of Radiology & Medical Informatics, Erasmus MC, Rotterdam, The Netherlands.

Abstract. This paper presents a method for improving airway tree segmentation using vessel orientation information. We use the fact that an airway branch is always accompanied by an artery, with both structures having similar orientations. This work is based on a voxel classification airway segmentation method proposed previously. The probability of a voxel belonging to the airway, from the voxel classification method, is augmented with an orientation similarity measure as a criterion for region growing. The orientation similarity measure of a voxel indicates how similar is the orientation of the surroundings of a voxel, estimated based on a tube model, is to that of a neighboring vessel. The proposed method is tested on 20 CT images from different subjects selected randomly from a lung cancer screening study. Results from our experiments showed that length of the airway branches segmented using the proposed method are significantly longer ($p = 0.0125$) as compared to only using probability from the voxel classification method.

1 Introduction

It has been shown in various studies that analysis of airways in CT, mainly the measurement of airway wall thickness, plays a significant role in the analysis of various lung diseases [1]. Airway tree segmentation plays a critical role in these studies, offering a starting point for conducting measurements on the airways. Nevertheless, current available airway segmentation methods are still far from perfect, limiting the measurements obtainable from these airway analysis studies to the larger airways that are easier to segment.

Most airway segmentation methods are based on the region growing algorithm [2–6]. The main difficulty in using the region growing algorithm lies in the fact that there often exist regions that have low contrast between airways and their surroundings, due to noise or pathologies such as emphysema. These regions often cause the region growing algorithm to leak into surrounding lung tissues. Currently there are two approaches to address this problem: explosion control and use of local image descriptors.

The idea of explosion control is to stop the segmentation in the low contrast regions where otherwise leakage would occur, while segmentation continues in other regions. Strategies for explosion control generally involve heuristic rules based on geometrical properties of the regions labelled. Some examples of these geometrical properties are: volume of the regions segmented [2], radius of propagation front [3], cross section area [4] and topology of thinned structure [5].

The second approach makes use of local image information to better differentiate between airways and their surroundings, for instance using pattern recognition techniques [7, 6] or local tube detection [8].

Previously, we have proposed a method for airway segmentation based on voxel classification and region growing [6]. In this paper, we propose to incorporate airway and vessel orientation information to further improve the voxel classification based method. This is done by using an orientation similarity measure that is computed from the orientation of a candidate airway voxel and the orientation of a neighboring vessel. The orientation similarity measure is then used as an additional criterion in the region growing.

The motivation for our work lies in the fact that every airway branch is accompanied by an artery. Sonka et al [9] decribed an approach that uses vessels to improve airway segmentation. The differences however is that our method uses a segmented vessel tree and the orientation computed from it, while the work described in [9] uses the proximity of the airway to the vessel, which is assumed to be a bright object nearby.

2 Vessel-guided airway segmentation

We start by first describing the construction of a voxel classification based airway appearance model, which is proposed in a previous work. We then proceed to explain the way the vessels are extracted and the computations of their orientations. After that, we present the way we compute the orientation of an airway candidate voxel, and how this is used with the orientation from a neighboring vessel to form an orientation similarity measure. Finally the segmentation framework is described, where the airway appearance model and the orientation similarity measure are used to form a decision function for a region growing algorithm.

2.1 Airway Appearance Model

An appearance model based on voxel classification is used. We based this appearance model on [6], where a k^{th} nearest neighborhood (KNN) classifier is used for differentiating between voxels from airway and non-airway classes. A brief review on how the model is constructed and used is presented here for the convenience of the reader. Refer to [6] for details.

Ideally, a gold standard provided by hand-tracing by a human expert should be used to construct or train the appearance model. However, such a ground truth of the airway trees is not feasible to obtain due to the extreme amount of

manual labour involved [5]. Therefore, a surrogate ground truth is used instead, which is imperfect but easier to obtain. We will refer to this surrogate ground truth simply as 'ground truth' in the following text.

The ground truth is obtained using a simple intensity based interactive region growing algorithm, where the user is required to provide a seed point and an intensity threshold. The highest threshold possible without any observable leakage is selected for each of the images individually. A second segmentation is produced using a slightly higher threshold. Due to the higher threshold, this 'exploded segmentation' usually has more airway branches and significantly more leakage. The exploded segmentation is used to exclude potential airways voxels that were missed by the ground truth from the non-airway class in training the appearance model.

The airway class then consists of all voxels labelled in the ground truth, but excluding the trachea and main bronchus. The non-aiway class consists of voxels surrounding the airways that are within the lung fields and are not marked by the exploded segmentation.

To ensure approximately independent training samples, only a small percentage (5%) of the voxels belonging to the airway class, selected randomly, are used as training samples. The same number of training samples are also selected randomly from the non-airway class. To prevent the samples belonging to airway class from having a bias towards the larger airway branches, the random sampling is done such that more samples will be drawn from the the smaller branches.

An initial set of local image descriptors is first used to compute the features of each training sample. This set of features consists of partial derivatives up to and including the second order, eigenvalues of the Hessian matrix, determinant and gaussian curvature of the Hessian, as well as combinations of eigenvalues that measure tube, plate and blob ($\sqrt{\lambda_1^2+\lambda_2^2+\lambda_3^2}$, $|\lambda_2/\lambda_1|$, $|\lambda_3/\lambda_1|$, $(|\lambda_1|-|\lambda_2|)/(|\lambda_1|+|\lambda_2|)$, $|\lambda_3|/\sqrt{|\lambda_1\lambda_2|}$). The partial derivatives of the image are computed at multiple scales by convolving the image with the partial derivative of a Gaussian kernel, based on scale space theory [10].

Sequential forward floating feature selection [11] is used to select an optimal subset of features from the initial set, which maximizes the area under the receiver operating characteristic (ROC) curve of the KNN classifier. In the feature selection process, training samples are randomly separated into two equal parts, a training set and a validation set. The final KNN classifier used is constructed using the optimal features of all the training samples.

Given a set of optimal features $\boldsymbol{x}$ computed at a particular position in the image, the posterior probability of $\boldsymbol{x}$ belonging to the airway class is defined as

$$P(A|\boldsymbol{x}) = \frac{K_A(\boldsymbol{x})}{K} \tag{1}$$

where A is the airway class, $K_A(\boldsymbol{x})$ is the number of nearest neighbors around $\boldsymbol{x}$ belonging to the airway class, obtained from a total of K nearest neighbors.

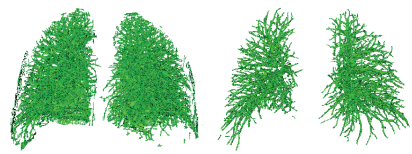

Fig. 1. Vessel segmentation with thresholding only on intensities (left) and on both intensities and tubeness measures (right).

2.2 Vessel Orientation

Extraction of vessel orientation involves three steps: vessel tree segmentation, centerline extraction and orientation computation.

The segmentation of the vessel tree in CT image starts by first segmenting the lung fields based on thresholding and morphological smoothing [12]. We then threshold the image with an intensity of t_v, followed by a second threshold process based on the eigenvalues of the Hessian matrix to remove other high intensity structures such as airway walls and fissures.

We obtain two measurements based on the eigenvalues [13], given as

$$m_1 = \frac{|\lambda_1| - |\lambda_2|}{|\lambda_1| + |\lambda_2|} \quad \text{and} \quad m_2 = \frac{|\lambda_1| - |\lambda_3|}{|\lambda_1| + |\lambda_3|}$$

where $|\lambda_1| \geq |\lambda_2| \geq |\lambda_3|$ are the eigenvalues of the Hessian matrix. Note that $0 \leq m_1, m_2 \leq 1$. Within a solid bright tube structure, λ_1 and λ_2 correspond to the principal curvatures along the direction from inside to the outside of the tube, and λ_3 corresponds to the principal curvature along the direction of the tube. Therefore, the eigenvalues within a tube structure will have a relationship of $|\lambda_1| \approx |\lambda_2| \gg |\lambda_3|$, thus resulting in $m_1 \approx 0$ and $m_2 \approx 1$.

The second threshold retains voxels with $m_1 \leq t_{m_1}$ and $m_2 \geq t_{m_2}$. Additional connected component analysis is used to remove small isolated objects (of less than 20 voxels in size) to obtain the final segmentation. Fig. 1 shows the result of vessel segmentation with and without thresholding via m_1 and m_2.

The centerlines are extracted from the segmented vessel tree using a 3D thinning algorithm [14]. Subsequently, the vessel orientation at the centerline voxels is measured as the eigenvector corresponding to λ_3. The reason for using orientation estimated through the Hessian eigen analysis is because it is less sensitive to noise and inaccuracies in the vessel segmentation, as compared to the orientation that would be obtained from the skeleton itself.

2.3 Orientation Similarity Measure

The orientation of the airways is extracted using Hessian eigen analysis in the airway probability image. Let $\boldsymbol{a} = (a_x, a_y, a_z)$ be the orientation of an airway

candidate voxel and $\boldsymbol{v} = (v_x, v_y, v_z)$ be the orientation of the voxel belonging to the centerline of the vessel that is nearest to it, we defined the orientation similarity measure as

$$s = \frac{| < \boldsymbol{a}, \boldsymbol{v} > |}{\| \boldsymbol{a} \| \| \boldsymbol{v} \|}$$

where s will be near to 1 only when the orientation of $\boldsymbol{a}$ is similar to $\boldsymbol{v}$.

2.4 Segmentation Framework

The airway tree is extracted using a region growing algorithm on the airway probability as described in [6]. The orientation similarity measure is used along with the airway probability (1) from the KNN classifier, when deciding whether a voxel belongs to an airway or not.

In region growing on the airway probability, it was often observed that a whole subtree of the airway is not segmented due to a small amount of voxels with low airway probability 'blocking' the way. This is especially pronounced for the smaller high generation branches, where 1 or 2 low probability voxels are sufficient enough to block the entire subtree after them.

We propose to solve this problem by reducing the threshold of the airway probability in cases of high orientation similarity. To accomplish this, we introduced 3 thresholds: an upper airway probability threshold T_u, a lower airway probability threshold T_l and an orientation similarity measure threshold T_s. The decision function for airway is then defined as

$$D(p(x,y,z), s(x,y,z)) = \begin{cases} 1, & p(x,y,z) \geq T_u, \\ 1, & T_u > p(x,y,z) \geq T_l \text{ and } s(x,y,z) \geq T_s, \\ 0, & \text{otherwise.} \end{cases} \quad (2)$$

where $p(x,y,z)$ is the airway probability and $s(x,y,z)$ is the orientation similarity measure of the voxel located at (x,y,z). The voxel is labelled as an airway if $D(p(x,y,z), s(x,y,z)) = 1$. Suitable values for these thresholds can be found for instance using cross validation.

2.5 Optimal Threshold Selection

Our method requires the selection of 3 threshold values, T_u, T_l and T_s. Due to the conservative nature of our ground truth, threshold selection based on measurements such as accuracy or segment overlap would result in an over conservative segmentation. Instead, we will aim to maximize the total length of branches segmented, while minimizing the chances of explosion.

A modified fast marching algorithm based on [3] is used to detect possible explosion and measure the branch length. This algorithm works by constantly monitoring the propagating front of the fast marching algorithm in a particular airway branch. The fast marching algorithm is initialized at each branch at bifurcations, which is detected when there is a discontinuity in the front.

Different from [3], which uses the minimum radius of the branches from previous generation for explosion detection, our approach uses the radius of the current branch after the first N step as reference instead. Explosion is detected when the ratio between the radius of the current front and the reference radius exceeded β. The number of branches at bifurcations is also monitored, where explosion is said to occur when the number of branches at a bifurcation exceededed γ.

The centroids of all the fronts at each time step of a particular branch are stored. The branch length can then be computed by summing up all the distances between the centroids from neighboring time steps. The total branch length is computed by summing up the length of all branches, excluding the trachea.

3 Results on 20 Low-dose CT Images

Experiments were conducted on 20 low-dose CT images from 20 different subjects enrolled in the Danish Lung Cancer Screening Trial (DLCST), with a voxel size of 0.78125×0.78125×1mm (except for one image that has a voxel size of 0.75×0.75×1mm). The 20 subjects were selected randomly from the screening study. A two-fold cross validation experiment was conducted, where the 20 subjects were randomly separated into two groups: A and B. Group A was then used as training set for constructing the classifier that was to be tested on group B and vice-versa.

3.1 Parameter Settings

A fast implementation of KNN based on approximate nearest neighbor (ANN) searching [15] is used as the classifier for the appearance model. The error *eps* is set to zero to turn off the approximation part of the ANN searching algorithm. A K of 21 was used for the KNN classifier of the appearance model. The features are calculated at 7 scales exponentially distributed within a range from 0.5mm to 3.5mm.

For vessel tree segmentation, t_v was set to -600HU, σ of 1mm was used for the computation of the Hessian matrix, with both t_{m_1} and t_{m_2} set to 0.5. In our experiments, the orientations of the vessels and airways were computed at a scale of 2mm for the orientation similarity measure. There are two reasons for this higher scale, as compared to the 1mm used in vessel tree segmentation: to compensate for the noise in the probability image and to make sure that the scale is large enough such that orientation in the centerline of the vessels can be estimated reliably.

For threshold selection, the N is set to 2, β to 3 and γ to 5. The airway probability thresholds T_u and T_l were varied over 21 different values (with 0 excluded), which was equivalent to the K used for the KNN classifier. The threshold T_s was varied over 21 different values ranging from 0 to 1. A total of 4011 different combinations of thresholds were tested.

The thresholds were optimized within the two-fold cross validation experiment, where only images from the training set were involved. Airway probability for each image was produced by a KNN classifier that was constructed in a leave-one-out manner. The threshold combination selected was the one that had the highest total branch length without any explosion detected for all cases.

3.2 Results

Maximum (Max), minimum (Min) and average (Avg) total branch length (measured in cm), along with true positive rate (TPR = TP/(TP+FN)) and false discovery rate (FDR = FP/(FP+TP)) of the segmentation results are presented in Table 1. Results from the best setting, obtained using the threshold selection process described previously, with airway probability alone and on image intensity alone are also included for comparison purpose. Computation for TPR and FDR was done with respect to the ground truth described in Sect. 2.1. Note that FDR does not only indicates false positives (leakage), but also newly discovered actual airway branches that were missed in the conservative ground truth.

Results from region growing on the image intensity were significantly worse than the ones that use the airway probability, be it with or without the orientation similarity criterion. Also during the optimization process, the criterion for the number of cases with explosion detected needed to be increased to 1. This was because there were a few images where leakage occurs no matter what intensity threshold was used. In the test results in Table 1, one of the test image was excluded due to explosion, as shown in Fig. 2(a). Fig. 2(b) shows the best result, while Fig. 2(c) and Fig. 2(d) shows representative results of region growing with intensity.

Both segmentation using the airway probability are better than the ground truth, with more new branches found than missed. Results with orientation similarity measure are in general more complete, with more and longer branches, as indicated in the results in Table 1. A paired t-test performed on the total branch length calculated from the segmentation results showed that the increase was significant ($p = 0.0125$). Examples are shown in Fig. 2(e) and Fig. 2(f). Fig. 2(g) and Fig. 2(h) show 2 examples where orientation similarity measure shows obvious improvements visually than the ones without it, with the former comparable and the later clearly worse than the surrogate ground truth. Examples where orientation similarity measure performed slightly worse, with either more obvious leakage or less branches, as compared to using airway probability alone are also shown in Fig. 2(i) and Fig. 2(j).

4 Discussion and Conclusions

An airway tree segmentation method that augments an airway appearance model with vessel orientation information is presented. The use of the airway probability image makes it possible to determine the orientation of an airway candidate

Table 1. Results with airway probability and orientation similarity measure (P+S), with only airway probability (P) and with image intensity (I). The values in brackets are those with a case excluded due to explosion from region growing on intensity.

	Max(cm)	Min(cm)	Avg(cm)	TPR(%)	FDR(%)
P+S	381(381)	103(103)	213(209)	97.58(97.62)	19.40(19.17)
P	299(299)	90(90)	186(186)	93.25(93.78)	14.01(14.05)
I	-(366)	-(94)	-(136)	-(86.82)	-(0.37)

voxel using Hessian matrix eigen analysis. The airway orientation of the candidate voxel is then compared with the orientation from a vessel nearest to it to form an orientation similarity measure. This orientation similarity measure is used to lower the threshold for airway probability during the region growing process, resulting in a more complete segmentation with longer airway branches. Results from our experiments showed that augmenting the airway appearance model with our orientation similarity measure gives better segmentations than with only the airway appearance model.

The explosion detection based on [3] that was used in our experiments worked well in general, but was not without problems. There were cases where it was either too sensitive or failed to detect leakage. Due to this reason, the thresholds obtained and used in our experiments were not really optimal. Tuning the thresholds manually or employing another explosion criterion may still improve the results.

It should be noted that we employed explosion detection only in the optimization process. Airway tree segmentations were subsequently generated using standard region growing on airway probability and orientation similarity. Employing a smarter region growing algorithm would likely improves the results as well.

We have showed a way to incorporate vessel orientation information into voxel classification based segmentation methods. However, the idea of using vessel orientation can also be useful when applied in combination with other methods, such as intensity region growing. In this work, we focused on improving the detection of small airway branches using orientation information extracted at a single, small scale. A natural extension of this work would be to use automatic scale selection for computing airway and vessel orientations, which we believe would further improve the accuracy and sensitivity of the orientation similarity measure, thus further improve the segmentation results.

Acknowledgments. This work is partly funded by the Danish Council for Strategic Research (NABIIT), the Netherlands Organisation for Scientific Research (NWO), and AstraZeneca, Lund, Sweden.

References

1. Berger, P., Perot, V., Desbarats, P., de Lara, J.M.T., Marthan, R., Laurent, F.: Airway wall thickness in cigarette smokers: quantitative thin-section CT assess-

ment. Radiology **235**(3) (Jun 2005) 1055–1064
2. Kiraly, A.P., Higgins, W.E., Hoffman, E.A., McLennan, G., Reinhardt, J.M.: 3D human airway segmentation for virtual bronchoscopy. In: SPIE Medical Imaging 2002: Physiology and Function from Multidimensional Images. Volume 4683. (April 2002) 16–29
3. Schlathölter, T., Lorenz, C., Carlsen, I.C., Renisch, S., Deschamps, T.: Simultaneous segmentation and tree reconstruction of the airways for virtual bronchoscopy. Volume 4684., SPIE (2002) 103–113
4. Kitasaka, T., Mori, K., Suenaga, Y., Hasegawa, J., Toriwaki, J.: A method for segmenting bronchial trees from 3D chest X-ray CT images. In: MICCAI (2). (2003) 603–610
5. Tschirren, J., Hoffman, E., McLennan, G., Sonka, M.: Intrathoracic airway trees: segmentation and airway morphology analysis from low-dose CT scans. Medical Imaging, IEEE Transactions on **24**(12) (Dec. 2005) 1529–1539
6. Lo, P., de Bruijne, M.: Voxel classification based airway tree segmentation. In: Medical Imaging 2008: Image Processing. Volume 6914., SPIE (2008) 69141K
7. Ochs, R.A., Goldin, J.G., Abtin, F., Kim, H.J., Brown, K., Batra, P., Roback, D., McNitt-Gray, M.F., Brown, M.S.: Automated classification of lung bronchovascular anatomy in CT using AdaBoost. Medical Image Analysis **11**(3) (June 2007) 315–324
8. Graham, M.W., Gibbs, J.D., Higgins, W.E.: Robust system for human airway-tree segmentation. In: Medical Imaging 2008: Image Processing. Volume 6914., SPIE (2008) 69141J
9. Sonka, M., Park, W., Hoffman, E.: Rule-based detection of intrathoracic airway trees. IEEE Transactions on Medical Imaging **15**(3) (1996) 314–326
10. Weickert, J., Ishikawa, S., Imiya, A.: On the history of Gaussian scale-space axiomatics. In Sporring, J., Nielsen, M., Florack, L., Johansen, P., eds.: Gaussian Scale-Space Theory. Kluwer Academic Publishers, Dordrecht, The Netherlands (1997) 45–59
11. Pudil, P., Novovičová, J., Kittler, J.: Floating search methods in feature selection. Pattern Recogn. Lett. **15**(11) (1994) 1119–1125
12. Hu, S., Hoffman, E., Reinhardt, J.: Automatic lung segmentation for accurate quantitation of volumetric X-ray CT images. Medical Imaging, IEEE Transactions on **20**(6) (June 2001) 490–498
13. Bülow, T., Wiemker, R., Blaffert, T., Lorenz, C., Renisch, S.: Automatic extraction of the pulmonary artery tree from multi-slice CT data. Volume 5746., SPIE (2005) 730–740
14. Wang, T., Basu, A.: A note on 'A fully parallel 3D thinning algorithm and its applications'. Pattern Recognition Letters **28**(4) (March 2007) 501–506
15. Arya, S., Mount, D.M., Netanyahu, N.S., Silverman, R., Wu, A.Y.: An optimal algorithm for approximate nearest neighbor searching fixed dimensions. J. ACM **45**(6) (1998) 891–923

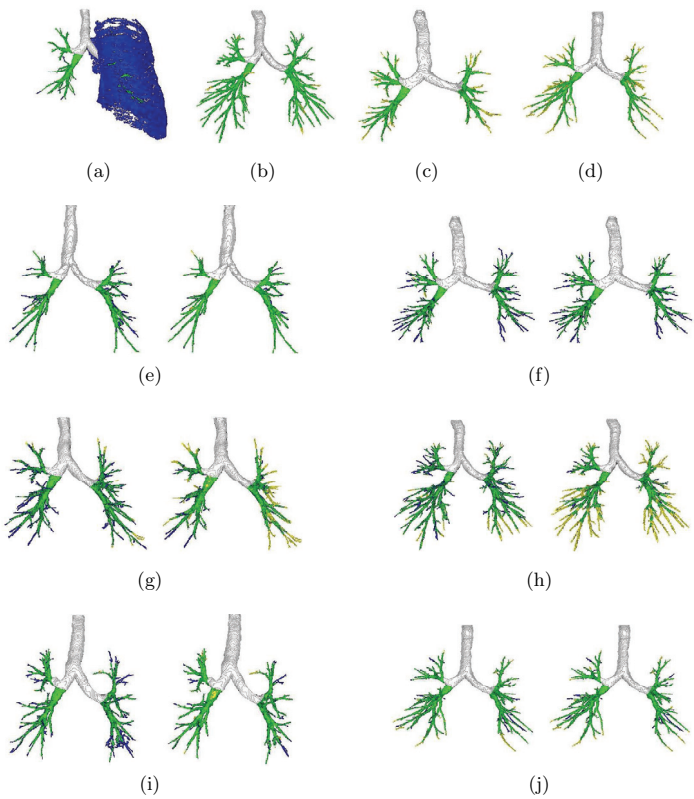

Fig. 2. Surface renderings of the segmentation results. Results from region growing with intensity are given in (a)-(d). Comparison of segmentation results using airway probability with (left) and without (right) orientation similarity measure are given in (e)-(j). Results from the following pairs of figures are from the same test image: (a) and (i), (b) and (j), (c) and (f), (d) and (h). The pre-segmented trachea and main bronchus are shown in white, true positives are show in green, false positives are shown in blue and false negatives in yellow, all with respect to the surrogate ground truth. (Refer to the electronic version for colours.)