Medical Image Analysis 14 (2010) 527–538

Contents lists available at ScienceDirect

Medical Image Analysis

journal homepage: www.elsevier.com/locate/media

Vessel-guided airway tree segmentation: A voxel classification approach

Pechin Lo [a,*], Jon Sporring [a], Haseem Ashraf [b], Jesper J.H. Pedersen [c], Marleen de Bruijne [a,d]

[a] Image Group, Department of Computer Science, University of Copenhagen, Denmark
[b] Department of Radiology, Gentofte University Hospital, Denmark
[c] Department of Cardio Thoracic Surgery, Rigshospitalet – Copenhagen University Hospital, Denmark
[d] Biomedical Imaging Group Rotterdam, Departments of Radiology & Medical Informatics, Erasmus MC – University Medical Center Rotterdam, The Netherlands

ARTICLE INFO

Article history:
Received 13 February 2009
Received in revised form 16 March 2010
Accepted 17 March 2010
Available online 27 March 2010

Keywords:
Airway segmentation
Lung computed tomography
Appearance model
Blood vessel
Classification

ABSTRACT

This paper presents a method for airway tree segmentation that uses a combination of a trained airway appearance model, vessel and airway orientation information, and region growing. We propose a voxel classification approach for the appearance model, which uses a classifier that is trained to differentiate between airway and non-airway voxels. This is in contrast to previous works that use either intensity alone or hand crafted models of airway appearance. We show that the appearance model can be trained with a set of easily acquired, incomplete, airway tree segmentations. A vessel orientation similarity measure is introduced, which indicates how similar the orientation of an airway candidate is to the orientation of the neighboring vessel. We use this vessel orientation similarity measure to overcome regions in the airway tree that have a low response from the appearance model. The proposed method is evaluated on 250 low dose computed tomography images from a lung cancer screening trial. Our experiments showed that applying the region growing algorithm on the airway appearance model produces more complete airway segmentations, leading to on average 20% longer trees, and 50% less leakage. When combining the airway appearance model with vessel orientation similarity, the improvement is even more significant ($p < 0.01$) than only using the airway appearance model, with on average 7% increase in the total length of branches extracted correctly.

1. Introduction

Chronic obstructive pulmonary disease (COPD) is among the leading causes of death and disability in the world, tending to be even more widespread in the future (Murray and Lopez, 1996; Rabe et al., 2007). Measurement of airway lumen dimension and wall thickness, as can be obtained from computed tomography (CT) images, play a significant role in the analysis and understanding of COPD in various studies (Nakano et al., 2000; Berger et al., 2005; Coxson and Rogers, 2005), where measurements from the smaller and higher generation airways are especially important (Lee et al., 2008). The success of such studies relies heavily on the availability of accurate and automated methods for airway tree segmentation. However, most airway tree segmentation methods are still limited to the larger and more visible airways, therefore there is an immediate need for a better airway tree segmentation method. Moreover, the lungs are anatomically divided into subregions based on the structure of the airway tree. This makes airway tree segmentation a useful starting point for tasks such as the segmentation of lobes (Kuhnigk et al., 2003; Ukil et al., 2006; Zhou et al., 2006) and pulmonary segments (Mori et al., 2008). Segmented airway trees can also be used as landmarks for guiding registration processes (see for instance Li et al. (2008)), resulting in a more accurate and natural transformation for applications such as disease progression monitoring (Gorbunova et al., 2008).

The general approach to segmentation of the airway tree involves variants of region growing applied on the image intensity (Mori et al., 1996; Sonka et al., 1996; Kiraly et al., 2002; Schlathölter et al., 2002; Aykac et al., 2003; Kitasaka et al., 2003; Singh et al., 2004; Tschirren et al., 2005). The assumption is that in CT images, the airway lumen is dark and surrounded by brighter structures, i.e. airway walls followed by lung parenchyma. The main problem with this approach is that there often exist small regions, where part of the airway wall is not visible and the airway lumen has intensities similar to the surrounding lung tissue, due to noise or pathologies such as emphysema. This results in 'leakage' in the region growing process in which surrounding lung regions are wrongly labeled as part of the airway tree.

One direct way to reduce leakage is to introduce prior knowledge of the geometry of the airway tree into the region growing algorithm. Mori et al. (1996) proposed to monitor the change in volume of the labeled region to detect leakage, and use the highest threshold without any leakage detected to segment the entire tree.

* Corresponding author.
E-mail addresses: pechin@diku.dk (P. Lo), sporring@diku.dk (J. Sporring), haseem@dadlnet.dk (H. Ashraf), jesper.holst.pedersen@rh.regionh.dk (J.J.H. Pedersen), marleen@diku.dk (M. de Bruijne).

1361-8415/$ - see front matter © 2010 Elsevier B.V. All rights reserved.
doi:10.1016/j.media.2010.03.004

Later works focused on the idea of stopping the segmentation locally where leakage occurs, while allowing the segmentation process to continue in other regions (Schlathölter et al., 2002; Kitasaka et al., 2003; Tschirren et al., 2005; van Ginneken et al., 2008). Detection of leakage via the radius of propagating fronts from the fast marching algorithm was introduced in Schlathölter et al. (2002). van Ginneken et al. (2008) replaced the fast marching with a sphere-constrained region growing and investigated a multi-threshold scheme. Both Kitasaka et al. (2003) and Tschirren et al. (2005) proposed to use geometrical properties within a volume of interest derived from previously detected airways, where the former used the area of cross-sections, and the latter used the topology of the thinned structure for leakage detection. Singh et al. (2004) proposed information gain as a region growing criteria to avoid leakage, where the local topology of the labeled region around a candidate voxel is implicitly included into the gain.

Leakage may also be avoided by improved differentiation between airway lumen and surrounding lung tissue. The current paper falls in this category. The key idea is to base the decision in the region growing not only on the intensity of a single voxel, but also on the information from the surroundings. Sonka et al. (1996) used the proximity of airways and vessels as one of the criteria for region growing. Various morphological operators designed specifically for detecting airways have also been investigated (Pisupati et al., 1996; Kiraly et al., 2002; Aykac et al., 2003). Fetita et al. (2004) coupled morphological operators with an energy based reconstruction method that takes into account the appearance of a bronchial tree. A combination of fuzzy logic rules and 2D template matching for segmenting airways was presented in Mayer et al. (2004). Ochs et al. (2007) presented a pattern recognition technique to classify various structures in lung CT images, including airways, on the basis of training points that were hand picked by experts. Graham et al. (2008) presented an airway segmentation algorithm that detects tube parts throughout the image, where the final segmentation is obtained by combining the detected tubes using a graph search technique.

The contribution of this paper is twofold. Firstly, we introduce an airway appearance model that automatically learns the characteristic appearance of the airways and the surrounding tissues from a set of segmented example images. The core of this appearance model is a classifier that is trained to differentiate between airway and non-airway voxels using a set of local image descriptors. We show that good results can be achieved without the need for high quality and complete airway tree segmentations as training data. The idea to use voxel classification for airway segmentation is similar to Ochs et al. (2007). However our appearance model differs in the choice of classifier and in the way the training samples are extracted, where we use random samples extracted from easily obtainable low quality airway tree segmentations instead of hand picked training points as described in Ochs et al. (2007).

Secondly, we propose to incorporate a segmented vessel tree to further improve the performance of the voxel classification-based appearance model. The fact that airways are accompanied by arteries is well known, and has previously been used in both airway and vessel segmentation (Sonka et al., 1996; Pisupati et al., 1996; Bülow et al., 2005). Vessels are especially useful for airway segmentation in CT because of its better visibility, as shown in Sonka et al. (1996). In our work, we extend this idea to using the orientations between vessels and airways, similar to Bülow et al. (2005) who used it for artery–vein separation. This exploits the fact that every airway branch is accompanied by an artery, and that both structures have similar orientation, which to our knowledge has not been applied to the segmentation of airway trees. The final segmentation is obtained with a 3D region growing algorithm based on a decision function that combines both the airway appearance model and the relationship between airways and arteries.

Early versions of this work were presented in Lo and de Bruijne (2008) and Lo et al. (2008). The current paper offers an extended evaluation and introduces a multi-scale approach to vessel orientation similarity. The article is organized as follows: We start by explaining the training of the airway appearance model in Section 2. Section 3 presents the various steps involved in computing the vessel orientation similarity measure. The segmentation framework that combines both the airway appearance model and the vessel orientation similarity is presented in Section 4. Section 5 presents the results of a set of experiments on 250 low-dose CT images. Discussion of the results from our experiments, comparison with other related works, and suggestions for possible improvements are presented in Section 6. Finally a conclusion is presented in Section 7.

2. Classification based airway appearance model

We propose an airway appearance model that is based on voxel classification. A potential drawback of such a classification-based appearance model is that it requires segmented training data, which may be difficult to obtain. In this work, an easily obtainable manual segmentation is used for training instead, which is incomplete but leakage free. The construction of this manual segmentation as well as steps taken to compensate for its incompleteness during the training process are explained in the following. This is followed by a description of the extraction of training samples, which focus on the smaller airways. Finally the choice of the classifier and its training process are presented.

2.1. Incomplete manual segmentation as a basis for training

Ideally, a gold standard obtained from hand-tracing by one or multiple human experts should be used for the training of a classification-based appearance model. However, such a gold standard is in general not available for airway trees, due to the extreme amount of manual labor involved. Fortunately, incomplete but leakage free airway tree segmentation can easily be obtained interactively, which we will show is sufficient to train our appearance model.

An intensity based region growing algorithm was used to obtain the manual segmentations needed, where a seed point within the trachea as well as an intensity threshold were determined manually. The highest threshold possible without causing any leakage was selected for each of the training images individually. This typically results in an over conservative segmentation that has many missing branches. However, the ‘background’ region directly surrounding such a conservative segmentation will always contain airway voxels. To exclude likely airway voxels trained as background, a second ‘leaked segmentation’ was obtained using a threshold slightly higher than the one used for the manual segmentation, which results in more and often longer airway branches, but also with some leakage. The voxels that were marked in the leaked segmentation denote uncertain regions that may be either airway or background, and were excluded from the training process. Fig. 1 shows an example of the manual and leaked segmentation.

Training was performed using two classes: the airway class and the non-airway class. The airway class consists of all voxels that were labeled in the manual segmentations, excluding the trachea, left and right main bronchi. The non-airway class was limited to voxels within the lung fields and close to the airways. The lung fields were extracted using thresholding and morphological smoothing, similar to Hu et al. (2001). The region within the lung fields that is close to the airways were obtained by dilating the manual segmentation with a sphere of radius R_{dilate}. The non-air-

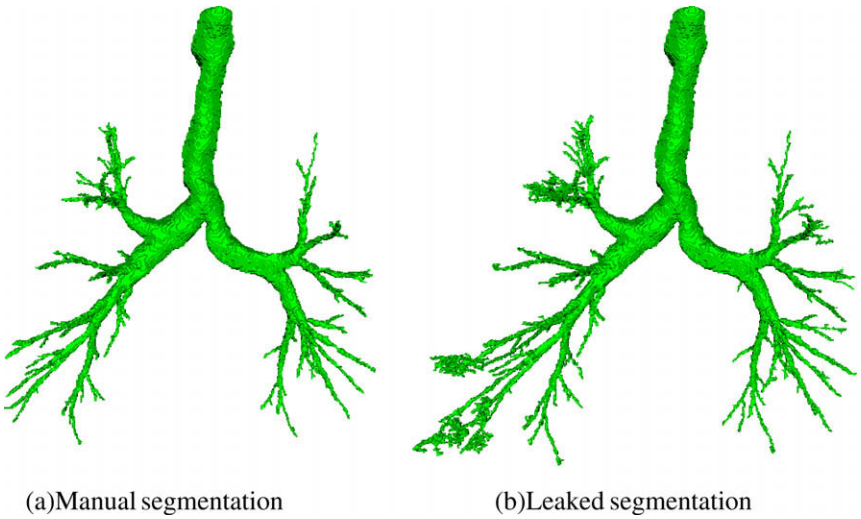

Fig. 1. Surface renderings of the interactively obtained segmentations used for training.

way class then consists of all voxels within this dilated region not marked by the leaked segmentation.

Only a fraction f_s of the voxels belonging to the airway class was used as training samples, excluding the trachea and the left and right bronchi. This results in a total of $N_s = f_s V$ training samples, where V is the total number of airway class voxels in the manual segmentation with trachea, and both left and right main bronchi excluded. The same number of training samples were also extracted from the non-airway class. In order to prevent the large number of voxels in the larger airways from dominating the appearance model, we sample evenly along the distance from the main bronchi. The distance from a voxel to the main bronchi is defined as the shortest distance measured within the segmented tree, which was obtained by applying the fast marching algorithm (Tsitsiklis, 1995; Malladi and Sethian, 1996) on the manual segmentation with the main bronchi as seeds.

The sampling process was performed by first grouping the voxels based on their distance from the main bronchi in bins of width W, and then randomly sampling a total of $N_b = N_s W / D_{max}$ training samples from each bin, where D_{max} is the maximum distance between a voxel in the manual segmentation to the main bronchi. In order to prevent a bin from being sampled too densely, only a fraction f_b ($\gg f_s$) of all voxels belonging to the bin were used as samples for extraction. Finally, samples were extracted by starting from the bin furthest away from the main bronchi. If the required number of samples from a bin was larger than the number of samples available in the bin itself, the remaining samples were extracted from the next available bin of shorter distance.

2.2. Airway probability

The training samples of Section 2.1 were used to train a classifier to differentiate between voxels belonging to the airway and non-airway class. Any classifier that outputs a posterior probability or other soft classification can be used here, but in this work we choose the k nearest neighbor (KNN) classifier (Cover and Hart, 1967; Duda et al., 2001). An initial feature set of local image descriptors was computed from the training samples, which consisted of spatial derivatives up to and including the second order, eigenvalues of the Hessian matrix (λ_1, λ_2 and λ_3, where $|\lambda_1| \geqslant |\lambda_2| \geqslant |\lambda_3|$), determinant and trace of the Hessian matrix, Frobenius norm of the Hessian matrix, and combinations of Hessian eigenvalues that measure tube, plate and blobness ($|\lambda_2/\lambda_1|$, $|\lambda_3/\lambda_1|$, $(|\lambda_1| - |\lambda_2|)/(|\lambda_1| + |\lambda_2|)$, $|\lambda_3|/\sqrt{|\lambda_1\lambda_2|}$). The partial derivatives of the image were computed at multiple scales by convolving the image with the partial derivatives of the Gaussian kernel (Weickert et al., 1997), and the features were standardized to zero mean and unit variance.

Sequential floating forward feature selection (Pudil et al., 1994) was used to find an optimal set of image descriptors that maximizes the area under the receiver operating characteristic (ROC) curve of the classifier. To this end, the training samples were randomly partitioned into two parts to compute the ROC curve: one third for training of the classifier and two thirds for validation. We constructed the final KNN classifier using the optimal combination of features and all training samples.

With the constructed KNN classifier, for each voxel we can now estimate the posterior probability of it belonging to the airway class, given a set of optimal features $\vec{x}$, using:

$$p(A|\vec{x}) = \frac{K_A(\vec{x})}{K} \tag{1}$$

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

3. Obtaining vessel orientation similarity

The vessels were first segmented from the lung fields using a multi-scale Hessian eigen analysis approach. The scale for calculating the Hessian matrix was selected for each voxel independently using the scale normalized (Lindeberg, 1998) Frobenius norm of the Hessian matrix:

$$\omega(\sigma_i) = \sigma_i^2 \sqrt{\lambda_1(\sigma_i)^2 + \lambda_2(\sigma_i)^2 + \lambda_3(\sigma_i)^2}$$

where the local vessel scale, σ_v, was then obtained by selecting the smallest scale that corresponds to a local maximum of ω across scales. Using the Hessian eigenvalues at scale σ_v, the following criteria were used to evaluate whether a voxel was part of a vessel or not:

- $\lambda_1,\ \lambda_2 < 0$ (Brightness)
- $\omega \geqslant T_\omega$ (Contrast)
- $(|\lambda_1| - |\lambda_2|)/(|\lambda_1| + |\lambda_2|) < T_1$ (Tubeness 1)
- $(|\lambda_1| - |\lambda_3|)/(|\lambda_1| + |\lambda_3|) > T_2$ (Tubeness 2)

where a voxel was labeled as vessel when all four criteria were satisfied. The brightness criterion ensures that only voxels that were brighter than their surroundings were selected, the contrast crite-

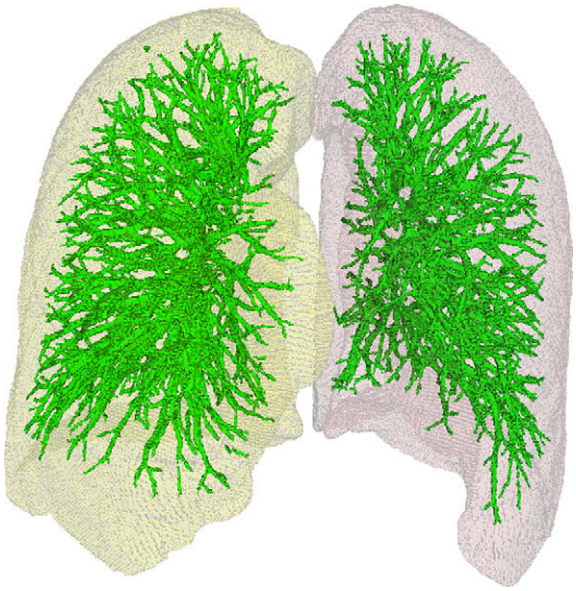

Fig. 2. An example of a segmented vessel tree.

rion reduces the effect of noise by ensuring a certain minimum contrast between the voxel and its surroundings, and finally the two tubeness criteria require vessels to locally resemble bright, solid cylinders. Within a solid bright tubular structure, λ_1 and λ_2 correspond to the principal curvatures along the directions perpendicular to the tube axis, and λ_3 corresponds to the tube axis. Hence, the eigenvalues within a tubular structure have a relationship of $|\lambda_1| \approx |\lambda_2| \gg |\lambda_3|$, resulting in a value near zero for the tubeness 1 criterion and a value near one for the tubeness 2 criterion.

Segmentation using the vessel criteria often results in additional small, isolated regions due to noise. A connected component analysis using a 6-connected neighborhood scheme was employed in order to remove these small isolated regions, where components with volumes less than V_{min} voxels were discarded. Fig. 2 shows an example of the vessel segmentation. Finally, the vessel centerlines were obtained using the 3D thinning algorithm presented in Wang and Basu (2007).

The vessel orientation at the centerline voxels was obtained as the eigenvector corresponding to λ_3 computed at the vessel scale σ_v. This measure is less sensitive to noise and inaccuracies in the vessel segmentation than the orientation obtained directly from the centerline itself. The orientations of the airways were also extracted the same way as the orientations of the vessels, but in contrast to vessels, we calculated the multi-scale Hessian matrix on the airway probability (1), where airways appears as solid, bright tube structures.

Given θ as the angle between the local tube orientation at an airway candidate voxel and the orientation measured at the centerline of a vessel nearest to it, we use $s = |cos(\theta)|$ as vessel orientation similarity measure. When the two orientations are similar, then $s \simeq 1$, and when the orientations are perpendicular, then $s \simeq 0$.

4. Segmentation framework

We obtain the airway tree segmentation using a 3D region growing algorithm, with a connectivity of 6-connected neighbors and a decision function that combines both the airway appearance model of Section 2 and the vessel orientation similarity of Section 3. The trachea, left and right main bronchi were automatically segmented and used as seeds for the region growing algorithm. Fig. 3 shows a block diagram of the proposed segmentation framework.

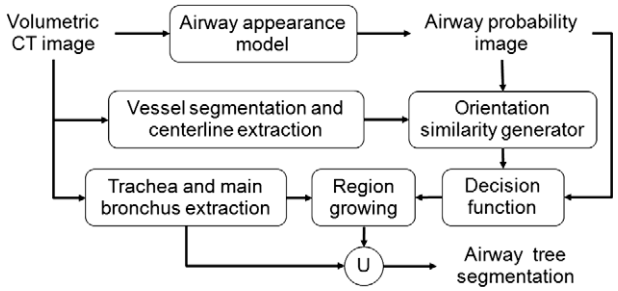

Fig. 3. Block diagram of the segmentation framework.

4.1. Preprocessing

A set of seed points within the trachea was first obtained automatically by searching for a dark elliptical object in the top few slices of a volumetric CT image. This set of seed points was then used to extract the trachea and the main bronchi using a fast marching based algorithm that detects bifurcations. This algorithm is based on the work by Schlathölter et al. (2002), where we use only the bifurcation detection of the original algorithm for preprocessing (refer to Appendix A for details). The algorithm was made to process voxels with intensity value below −900 HU, which was chosen slightly higher than the intensity of air in order to cope with noise and possible artifacts within the trachea. To only extract up to the left and right main bronchi, the algorithm was made to stop after extraction of all first generation branches.

4.2. Airway segmentations

The trachea and the left and right main bronchi obtained in Section 4.1 were used as seed points in a region growing process to extract the remainder of the airway tree, using the airway probability and vessel orientation similarity measures. Here, the vessel orientation similarity was used as a means to lower the threshold on airway probability in regions with low airway probability according to the appearance model, but with local tube orientation that is similar to the orientation of nearby vessels. Three thresholds were introduced for this purpose: upper probability threshold T_u, lower probability threshold T_l, and vessel orientation similarity threshold T_s. All voxels with an airway probability larger than or equal to T_u were accepted automatically. For voxels with airway probability between T_u and T_l, the decision was made based on the vessel orientation similarity and T_s. The decision function for the acceptance of a candidate voxel is defined as

$$D(p(A|\vec{x}), s) = \begin{cases} 1, & p(A|\vec{x}) \geqslant T_u \\ 1, & T_u > p(A|\vec{x}) \geqslant T_l \text{ and } s \geqslant T_s \\ 0, & \text{otherwise} \end{cases} \quad (2)$$

where $p(A|\vec{x})$ is the airway probability computed from (1), and s is the vessel orientation similarity of the candidate voxel. The voxel is labeled as an airway when $D(p(A|\vec{x}), s) = 1$.

5. Experiments and results

Experiments were conducted on 250 low-dose CT images (120 kV, 40 mAs) from 250 different subjects enrolled in the Danish Lung Cancer Screening Trial (DLCST) (Pedersen et al., 2009), where participants were current or former smokers at an age between 50 and 70 years. All images had a slice thickness of 1 mm and in-plane voxel size ranging from 0.72 to 0.78 mm. Three different datasets were used:

Dataset 1: Images from 32 randomly selected subjects. For these images, the leakage free manual segmentations and the leaked segmentations, as described Section 2.1, were made for training purpose.

Dataset 2: Images from 18 subjects with moderate to severe emphysema (an average of 24.25% of the lung volume has intensity below −950 HU). For these images, obvious leakage already occurred using intensity based region growing at the minimum threshold of −1000 HU.

Dataset 3: Images from 200 randomly selected subjects.

The performance of appearance model based region growing was evaluated with and without vessel orientation similarity and was compared to that of region growing based on intensity alone for all three datasets.

The objective of the experiment on dataset 1 is to study the feasibility of the proposed voxel classification approach, which is trained using a manual segmentation that was incomplete but leakage free. Dataset 2 investigates the ability of the proposed method to avoid leakages that are unavoidable using intensity alone. The experiment conducted on the dataset 3 shows the general performance of the proposed method on a large set of images.

5.1. Parameter settings

KNN classification was performed using the ANN library for approximate nearest neighbor searching (Arya et al., 1998). A neighborhood of $K = 21$ was used, and the approximation error *eps* was set to zero to turn off the approximation part of the algorithm. For the extraction of training samples, as described in Section 2.1, a dilation radius $R_{\text{dilate}} = 5$ mm was used in order to include both the airway walls and some surrounding lung tissues of the small airways. The fraction of voxels sampled was set to $f_s = 0.05$, the bin width for the sample extraction process was set to $W = 3$, and the fraction of voxels used as samples for extraction in each bin was $f_b = 0.5$. A total of seven scales, distributed exponentially between 0.5 mm and 3.5 mm were used to compute the features. The same scales were also used to segment the vessels in Section 3. A contrast threshold $T_\omega = 100$, and tubeness measure thresholds $T_1 = T_2 = 0.5$ gave acceptable vessel segmentation results for our application. $V_{min} = 20$ voxels was used to filter off the small isolated regions in the segmented vessel tree.

Suitable settings for the probability thresholds T_u and T_l, and the vessel orientation similarity threshold T_s were selected automatically using the leakage detection algorithm as described in Appendix A. This algorithm, derived from Schlathölter et al. (2002) and van Ginneken et al. (2008), is used to compute the tree length (TL) and percentage of leakage voxels (LVP). We define TL as the total length of all correct branches excluding the trachea. LVP is defined as the percentage of leakage voxels among all labeled voxels, with the trachea and the left and right main bronchi excluded. The probability thresholds T_u and T_l were varied over 21 different values (with 0 excluded), which was equivalent to the number of neighbors K used for the KNN classifier. The vessel orientation similarity 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 threshold combination selected was the one where further increase in the LVP will not give any significant increase in the TL, with a maximum allowed LVP of 5%. Only images from the training set were involved in selecting the threshold settings. The airway probability for each image was obtained from a KNN classifier constructed in a leave-one-out manner.

5.2. Comparison to manual segmentation

A twofold cross-validation experiment was conducted on the 32 images of dataset 1, where the subjects were randomly separated into two groups of 16 images each. The first group was used as training set for selecting the features, constructing the KNN classifier and determining the thresholds that were to be applied to the second group and vice-versa. Because of the lack of a ground truth, true positive and false positive rate based analysis is insufficient to evaluate the results of the experiments. We therefore also detect leakage using the modified Schlathölter algorithm, as described in Appendix A, for evaluation purpose. We report the TL and LVP as described in Section 5.1, as well as the branch count (BC), which is the number of correct branches excluding the trachea. As the modified Schlathölter algorithm is prone to being too sensitive and ended up rejecting true airway branches, detected leakage that overlapped with the manual segmentation was added back into the segmentation to reduce the number of true airway branches rejected.

We compare three different segmentation approaches: region growing on intensity, region growing on airway probability, and the proposed approach of using both airway probability and vessel orientation similarity. The threshold for region growing using intensity was determined based on the training set using the optimal threshold selection procedure described in Section 5.1, from a range of intensities between −1000 HU and −900 HU. The threshold for using airway probability alone was obtained by selecting the best performing threshold settings when $T_u = T_l$ and $T_s = 0$.

Table 1 shows the average results from the cross-validation experiments, given as true positive rate (TPR = TP/(TP + FN)), false discovery rate (FDR = FP/(FP + TP)), BC, TL, and LVP, where TP, FP and FN are true positives, false positives and false negatives respectively. The airway probability based methods have significantly higher TPR than the intensity based region growing. The combination of airway probability and vessel orientation similarity results in significantly longer TL ($p < 0.01$) than using only airway probability, with no significant increase in LVP ($p = 0.33$). The average LVP is the highest for the experiment with intensity based region growing, which is in part caused by two cases of severe leakage with LVP >97%. Figs. 4 and 5 show the surface renderings and slice view of three cases of segmented airways using the different methods, with two common cases and one case where intensity regions growing results in severe leakage.

5.3. Effects of parameter settings

To investigate the impact of the parameters related to the airway appearance model, the cross-validation experiment in Section 5.2 was repeated with varying parameter settings. Three main parameters were considered, which were the number of nearest neighbors K, dilation radius R_{dilate} and sampling bin width W. These parameters were varied from their original values in Section 5.1 one at a time. To simplify comparison and isolate the effects of the appearance model, the thresholds were chosen equal to those used in Section 5.2 for the region growing based on the appearance model and without vessel orientation similarity, thus with $T_u = T_l$ and $T_s = 0$. Table 2 shows the average results of the different parameter settings.

5.4. Leakage avoidance and performance on large dataset

The 32 images of dataset 1 were used to train the classifier and to select the thresholds for processing datasets 2 and 3. As there were no manual segmentations available for the test data, we report the BC, TL and LVP as computed directly using the modified Schlathölter algorithm. Unfortunately, this also means that there

Table 1
Average results from twofold cross-validation experiment with dataset 1.

	TPR (%)	FDR (%)	BC	TL (mm)	LVP (%)
Intensity	90.41	7.43	101.12	1347	6.19
Airway probability	98.40	21.62	150.44	1939	2.40
Airway probability + vessel orientation similarity	98.68	24.94	161.44	2084	2.95

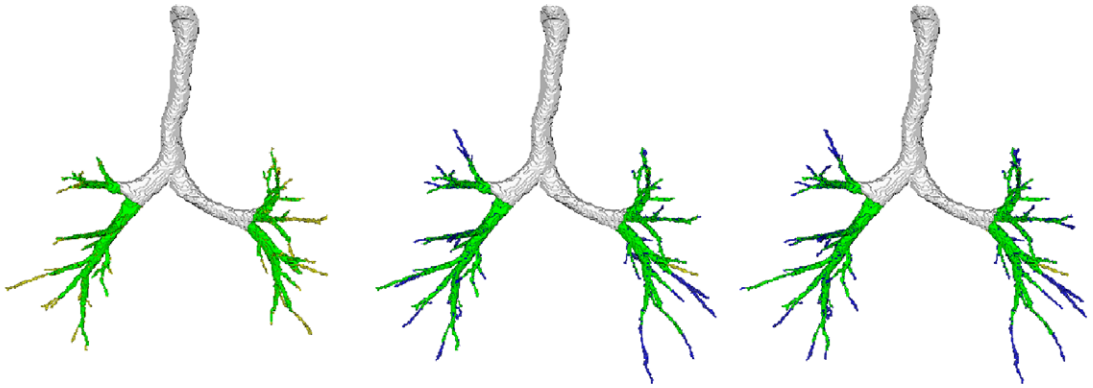

(a) TL: 764 mm, 1414 mm, 1497 mm

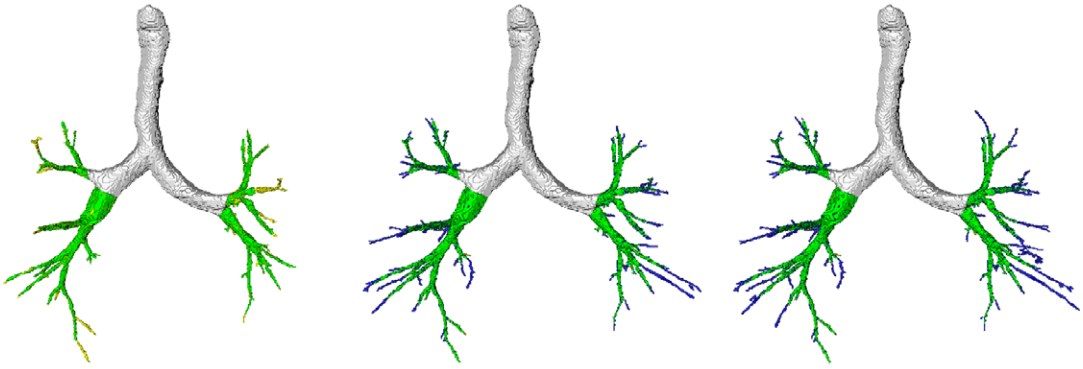

(b) TL: 1008 mm, 1750 mm, 2030 mm

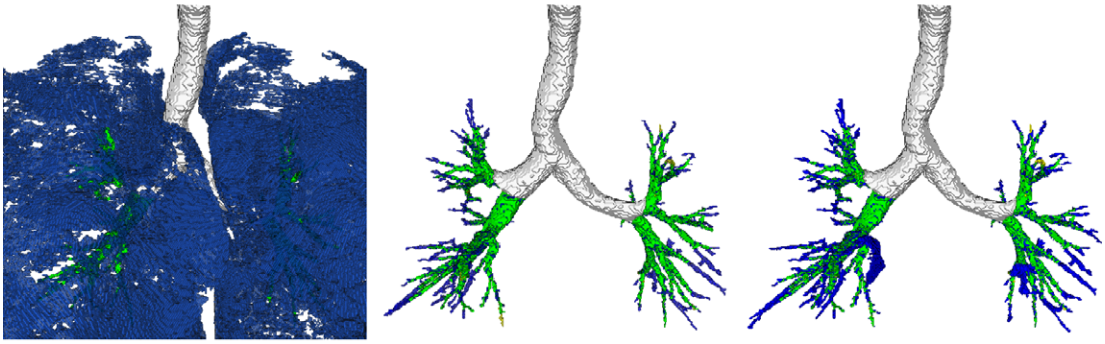

(c) TL: 1953 mm, 2084 mm, 2304 mm

Fig. 4. Surface renderings of results of three different subjects from the twofold cross-validation experiment, along with the TL from the different results. From left to right: surface renderings of segmentation results using intensity, using airway probability, using both airway probability and vessel orientation similarity. Pre-segmented trachea, left and right main bronchi are shown in white, TP, FP and FN are shown in green, blue and yellow respectively. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

will be more overestimation in the LVP as compared to the Sections 5.2 and 5.3, due to the lack of manual segmentation.

Table 3 shows the average results for dataset 2, which contains images from 18 subject with different degree of emphysema, where leakage occurs at −1000 HU. The minimum and maximum LVP, and the number of cases with LVP exceeding 10% are also presented to give a better indication of the amount of leakage occurring in the different methods. Fig. 6 shows an example of the results obtained by the different methods.

All three methods were applied to the 200 screening images in dataset 3, with the average TL, BC and LVP shown in Table 4. We also visually inspected the surface renderings of all results from

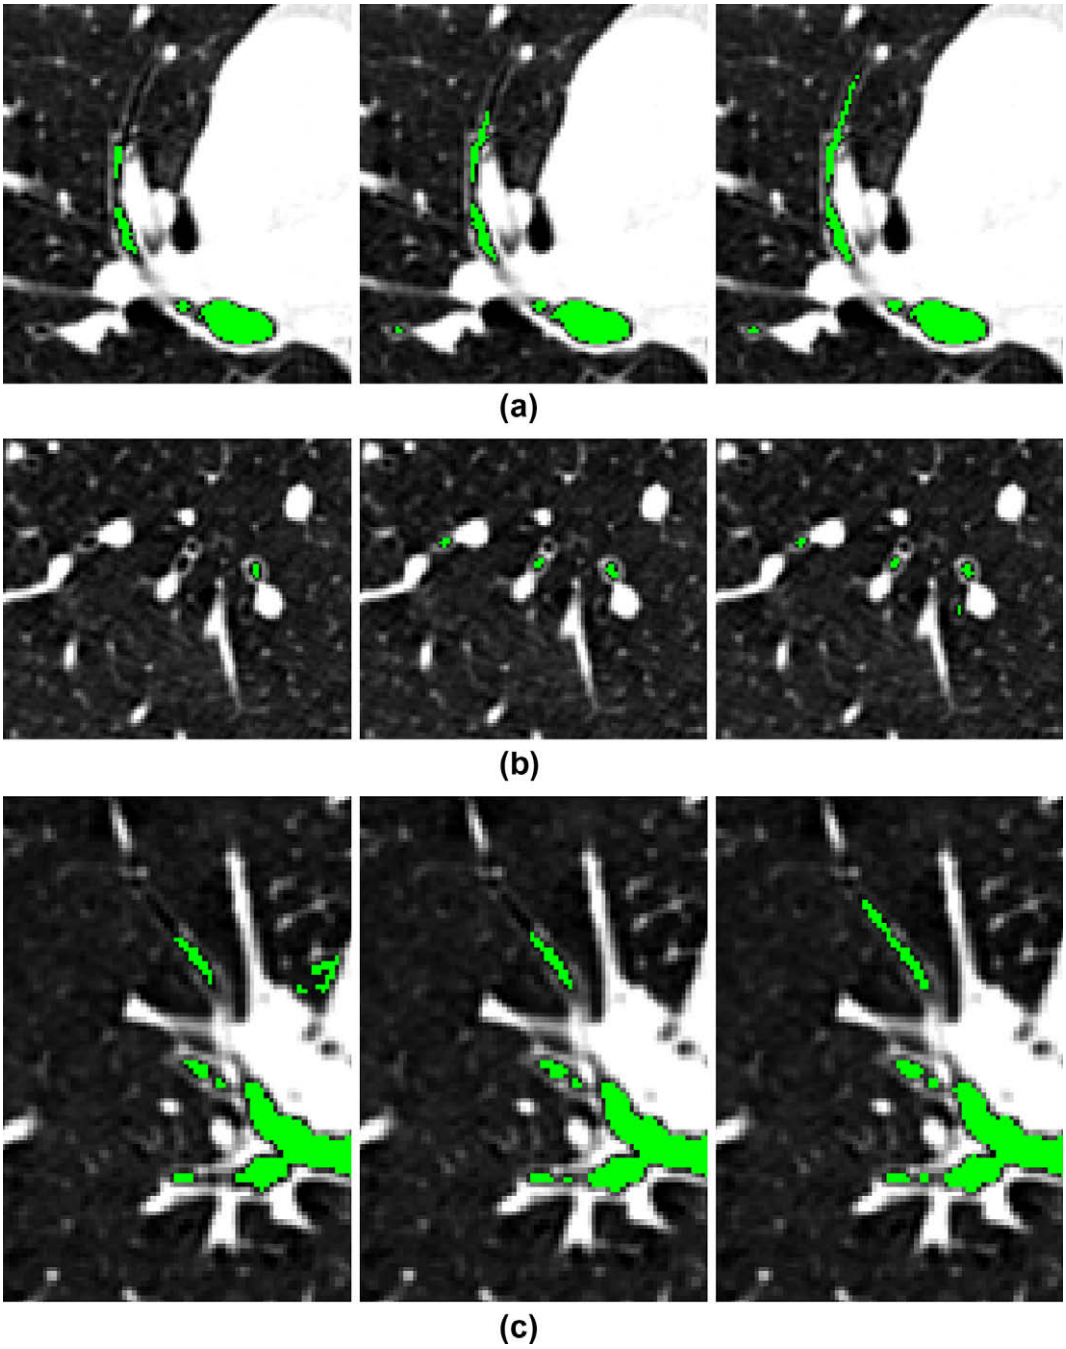

Fig. 5. Axial view images of three different subjects from the twofold cross-validation experiment. From left to right: segmentation results using intensity, using airway probability, and using both airway probability and vessel orientation similarity.

Table 2
Results from a series of experiments conducted using different parameter settings for the construction of the appearance model. Changed parameters are indicated in bold.

K	R_{dilate}	W	TPR (%)	FDR (%)	TL (mm)	LVP (%)
21	5	3	98.40	21.62	1939	2.40
11	5	3	97.78	20.52	1911	1.60
31	5	3	98.48	21.47	1906	2.08
21	**2.5**	3	98.58	18.89	1851	4.92
21	**10**	3	98.58	24.13	1938	2.59
21	5	**10**	98.31	19.13	1852	1.95
21	5	**20**	98.63	25.84	1966	7.35

Table 3
Average results from 18 cases in dataset 2, with severe leakage for intensity based region growing at −1000 HU.

	TL (mm)	BC	LVP (min–max) (%)	LVP > 10%
Intensity	2057	151.33	76.76 (5.22–98.94)	16
Airway probability	2377	174.28	12.77 (0.35–89.65)	4
Airway probability + vessel orientation similarity	2590	193.17	28.79 (0.47–97.20)	6

the proposed method that combines both airway appearance model and vessel orientation similarity. Obvious leakage were observed in only eight cases, all of which were correctly identified by the leakage detection algorithm. We also observed some obvious falsely detected leakages in five cases, where a whole subtree of true

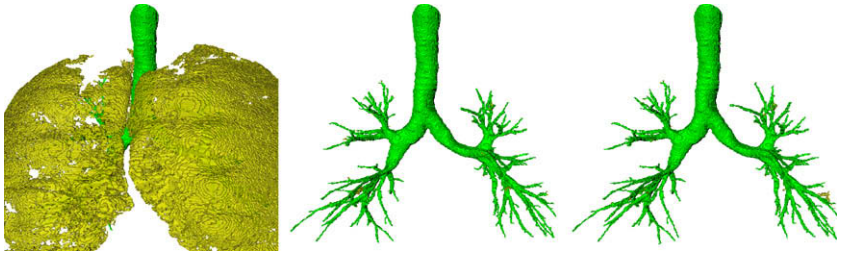

Fig. 6. Surface renderings of results from a subject with emphysema. From left to right: surface renderings of segmentation results using intensity, using airway probability, and using both airway probability and vessel orientation similarity. Leakages are marked in yellow. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

Table 4
Average results from 200 cases in dataset 3.

	TL (mm)	BC	LVP (%)
Intensity	1401	99.26	13.29
Airway probability	2012	146.39	2.23
Airway probability + vessel orientation similarity	2115	154.11	4.54

airway branches was classified as leakage by the modified Schlathölter algorithm. Fig. 7 shows results of four different subjects from this dataset from the proposed method.

6. Discussion

This paper investigates whether:

1. Local image descriptors can improve upon intensity based airway segmentation.
2. A trained appearance model is feasible.
3. Inclusion of co-orientation between vessels and airways can improve results.

A simple 3D region growing framework is used with no further geometrical constraints. Airway probability estimates based on local image descriptors as well as vessel orientation similarity are incorporated in this framework and the performance gain compared to intensity region growing is measured.

Table 1 shows that the methods that use the classifier based appearance model perform better in terms of true positive rate and tree length than intensity based region growing. The voxel classification based methods resulted in higher false discovery rate compared to intensity based method, but it must be noted that a slightly higher false discovery rate may actually be desirable in our case, as new branches found that are not in the manual segmentation are labeled as false positives as well. Upon visual inspection, the majority of 'false positives' from the appearance based methods in our experiments turned out to be valid airways that were missing in the manual segmentation. An example of this is shown in Fig. 5. This is further verified by the modified Schlathölter algorithm, where we observe a large difference between false discovery rate and percentage of leakage voxels in Table 1, with over 20% in false discovery rate compared to less than 5% in percentage of leakage voxels. This suggests that the majority of the false positives in the voxel classification based methods are actually real airways. For the intensity based region growing method, the false discovery rate (7.43%) is only slightly higher than percentage of leakage voxels (6.20%).

Results from our experiments on the three dataset in Tables 1, 3 and 4 all showed a similar trend, with the proposed method having the highest tree length and branch count, followed by region growing with airway probability, and with region growing with intensity having the lowest tree length and branch count. The appearance model based region growing methods are also less prone to leakage, and is capable of avoiding leakages that are unavoidable by region growing with intensity. This is shown in Table 3, where the appearance model based methods clearly have less percentage of leakage voxels and also less cases with severe leakage (LVP > 10%), as compared to region growing with intensity.

6.1. Sensitivity to parameter settings

As mentioned earlier in Section 5.3, the threshold for the airway probability used for each fold of the experiment were fixed to those used in Section 5.2 to simplify comparison and isolate the effects of the appearance model. Therefore, the threshold settings used may not be optimal for the parameter settings used in Table 2. Nonetheless, the overall true positive rate and tree length as observed in Table 2 are still quite high, with a minimum of 97.78% and 1851 mm respectively. However, it was observed that the percentage of leakage voxels varies quite a bit, with a maximum of 7.35%.

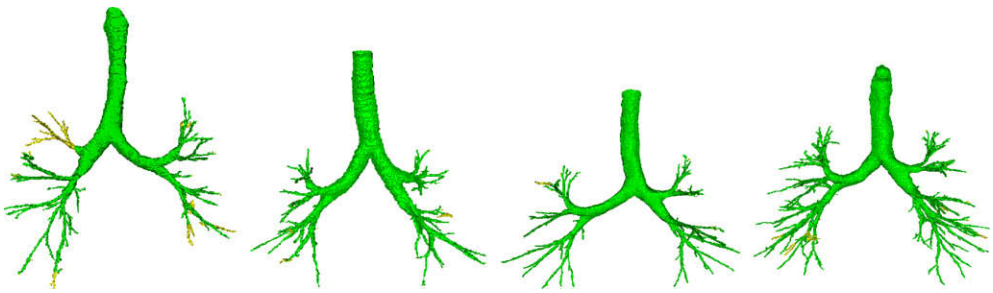

Fig. 7. Surface renderings of results from four different subjects from the large dataset obtained using appearance model and vessel orientation similarity measure. The yellow regions are regions classified as leakages by the modified Schlathölter algorithm. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

It should be noted that all, except for a single case, have percentage of leakage voxels of less than 5%, which to our experience is well within acceptable range. The high tree length and low percentage of leakage voxel over most of the tested parameter settings suggests that the appearance model performs quite well for a wide range of parameters.

Varying the number of nearest neighbors K in the appearance model has very little effect on the performance, resulting in a slight decrease in tree length of no more than 30 mm, and with percentage of leakage voxels remaining well below 3%. Performance degrades when the dilation radius R_{dilate} for extracting non-airway samples becomes too small. Since the over conservative, region growing based segmentation used for training underestimates the lumen, a small value of R_{dilate} results in a large number of actual airway lumen samples being labeled as 'non-airway' in the training. On the other hand, if the R_{dilate} becomes too large, then non-airway samples will be taken from voxels far away from the airway walls, which are less relevant for discriminating between airway lumen and background. A suitable value of R_{dilate} ensures that samples are taken from the airway walls as well as the nearby background. A sharp increase in percentage of leakage voxels is observed when taking a large bin width W in extracting samples along the tree. The reason is that samples from the large airway branches start to dominate the appearance model, causing the appearance model to be less sensitive to the smaller scale features that are required to avoid leakages in the smaller airway branches.

The modified Schlathölter algorithm described in Appendix A contains a large number of parameters and rules, which to our experience are not robust and depended heavily on the quality of the images and conditions of the subjects, e.g. ultra-low dose or scans of patients showing airway pathology. For this particular reason, we only use the modified Schlathölter algorithm for selecting suitable values for T_u, T_l and T_s, which in practice can also be done manually. It should be noted that the proposed method only uses simple region growing for generating the segmentation results. We also use the modified Schlathölter algorithm to estimate the amount of leakage in segmentation results for evaluation purpose, due to the lack of ground truth.

6.2. Effects of vessel orientation similarity measure

The airway probability image from the appearance model provides a good indication of the location of the airways in the image, but it is noisy, with many single voxels or small areas within the airways having relatively low airway probability. In a region growing process this can lead to entire subtrees being discarded due to one broken connection caused by a small cluster of low probability voxels. This happens especially in the thin, peripheral airway branches, where one or two low probability voxels are sufficient to block the entire descending subtree. A simple post processing of the airway probability image is not sufficient to remove these problems; we have experimented with coherence enhancing anisotropic diffusion (Weickert, 1999) to remove the low probability areas, but did not observe a significant improvement. Incorporating vessel orientation similarity as an additional criterion in the region growing helps to overcome these areas, which is supported by the improvements shown in Table 1.

A single scale version of the proposed method was reported in our previous paper (Lo et al., 2008). In comparison, the current, multi-scale approach enables reliable computation of the orientation of airways and vessels of various sizes rather than targeting only the small airways and vessels, which results in an overall improvement in the amount of branches segmented.

A potential weakness of the proposed vessel orientation similarity measure is that it is not well defined near airway or vessel bifurcations. We therefore rely more on the airway probability and accept all voxels where the airway probability is high, and only use the weaker vessel orientation similarity as a second opinion when the airway probability is low. This is implemented using the upper probability threshold T_u, which makes the decision independent of the vessel orientation similarity when the airway probability is high.

6.3. Comparison to results in literature

It is difficult to compare our results to the results obtained by other authors as most airway tree segmentation methods were evaluated by visual inspection only. One of the exceptions is the work by Tschirren et al. (2005), where the number of extracted branches were counted. However, the segmentation results were restricted to the sixth generation and only the number of 'named branches' extracted were counted, which makes a direct comparison with our results impossible. Another exception is the recent work by van Ginneken et al. (2008), in which the average tree length, the number of branches extracted, and the branch length at different generations were reported for several different datasets including a set of 50 low dose images from the NELSON study (van Iersel et al., 2007). Our results on the large dataset in Section 5.4 are similar to the results reported by van Ginneken et al., 2008, who reported a tree length of 2184 mm and an average branch count of 166, as compared to 2127 mm and 155 from our proposed method. A slightly smaller tree length would be expected in our case since our population contains more women (52% compared to 17% for the NELSON study), who typically have smaller lungs and shorter airway branches than men. The scans from DLCST are also of lower resolution as compared to the NELSON study, where the slice thickness is 1 mm for the DLCST study and 0.7 mm for the NELSON study. The method presented in this paper, using a smart airway appearance model with a single, global threshold setting with a simple 3D region growing algorithm, is complementary to the method in van Ginneken et al. (2008). Combining the two approaches may result in even better results.

We participated in the EXACT'09 (Lo et al., 2009b) airway extraction challenge that was organized at MICCAI 2009, where the proposed method was trained and applied to a diverse set of CT scans. Other than that the thresholds T_u, T_l and T_s were selected manually using images from the training set, the parameters for the appearance model and vessel orientation similarity remained same as those reported in Section 5.1. In comparison to other methods evaluated in EXACT'09, the proposed method is capable of obtaining results with very little leakage while still extracting a relatively large number of branches correctly. Among the 15 methods in EXACT'09, seven methods resulted in both a lower tree length an a higher false positive rate. Compared to the remaining seven methods, results from the proposed method stands out mainly by the small amount of leakage, with an average tree length of 1184 mm and a false positive rate of 0.11%. The appearance model constructed on dataset 1 was applied to a scan from EXACT'09, which was acquired using the Siemens Sensation 16 scanner and a very sharp convolution kernel B70s, where we only observed some reduction in the amount of branches extracted as compared to the appearance model trained using the dataset from EXACT'09 (Lo et al., 2009a).

6.4. Possible improvements

One might expect that a trained appearance model, such as is used in this paper, cannot perform much better than the segmentations that were used to train it, which in our case were obtained using a simple intensity based region growing. However, the airway generation and the number of subtrees that can be extracted in such a simple region growing method differ from image to im-

age and from lobe to lobe. Therefore, even though the airway tree segmentations used for training are incomplete individually, it is still possible to obtain sufficient examples of voxels that cover an entire airway tree by using a collection of these incomplete airway trees.

Since decisions are made for each voxel independently, small holes do occur in the segmentation, especially in the larger airways. These holes could, for instance, be removed by applying a morphological closing operator on the segmentation results. Alternatively, more accurate segmentations of the inner and outer airway walls could be sought around the centerlines extracted in this paper using, for instance, graph cut segmentation, as was presented in Pedersen et al. (2009). In the current paper, we aimed to extract the airway trees as complete as possible and were less concerned with the accuracy of the extracted airway lumen. We therefore did not perform such post processing in this work.

Despite the good results, the current setup still misses many small peripheral airways branches. The reason for this is that examples of such small airways are not available in the training segmentations obtained by the interactive intensity based region growing algorithm. A way to approach this problem would be to create a better training set, for instance using manual or interactive segmentation, such as in Tschirren et al. (2009), or by first segmenting high quality, high resolution, clinical dose data, using the intensity region growing method described in Section 2.1, and then train the classifier on lower quality simulated low dose images.

Due to the threshold relaxation introduced by the vessel orientation similarity, the proposed method is more prone to leakage as compared to using airway probability alone, with slightly more number of cases with severe leakage (LVP > 10%) and percentage of leakage voxels, as shown in Tables 3 and 4. This is mainly caused by emphysematous areas around the hilum that locally resemble airway lumen. Leakage into these areas can cause leakage to the lung surface, which usually has high response in the airway appearance model. Fig. 4c (right-most image) shows an example of such case. A way to reduce the severity of this effect would be to prevent the border of the lungs from being evaluated, for instance by restricting the evaluation on regions marked by an eroded lung segmentation.

Our current implementation assumes that the detected vessel nearest to an airway candidate voxel is the accompanying artery. It may be useful to discard airway candidates that are too far away from the detected vessels, in order to further improve segmentation results, similar to Sonka et al. (1996). In addition, an artery–vein separation algorithm (Bülow et al., 2005) could be applied to process only those points that are at a certain distance from the extracted arteries. However, the majority of the false positives we observed in our experiments are caused by high responses from the appearance model at the hilum as mentioned above. Therefore the extent of improvements introduced by adding this extra step may be minimal.

Another approach to improve results would be to use a more flexible decision function, e.g. a second stage classifier that takes into account the airway probability, vessel orientation similarity, distance from nearest vessel, etc. One drawback of a second stage classifier is the increase in computation time. However, a more serious problem would be the continuous usage of the incomplete manual segmentations for training, which may result in a classifier that produces over conservative segmentations. It should be noted that our proposed method performs very well, despite using a relatively simple threshold based approach, as very few branches in the training data are missed, especially when taking into account the number of extra branches found.

Despite of the good results, it is a fact that the proposed method requires a relatively long computation time. On average, the total time needed to segment the airway tree from a single CT scan using the proposed method was 55 min on a single CPU of an Intel Xeon X5355 processor (2.66 GHz). The majority of the time is spent on voxel classification to generate the airway probability image (ca. 25 min), and on computing the seven Gaussian blurred images for the original image and the airway probability image (ca. 20 min). The final segmentation process however is relatively fast, with a total execution time of around 10 min. However, it should be noted that not only most of the processes are independent from each other, the operations within are also highly independent. This makes the proposed approach highly parallelizable, and therefore a large performance gain can be expected from the usage of parallel processing technologies such as multi-core processors and graphics processing unit.

7. Conclusion

An airway tree segmentation method that is based on a trained airway appearance model is presented. It is shown that good results can be obtained even with imperfect training data, with segmentation results that are better than the training data itself. We combine the airway appearance model with a measure of orientation similarity between airways and vessels, which results in an additional significant improvement in segmentation performance with an increase in detected airway tree length and little or no increase in leakage.

Acknowledgments

This work is partly funded by the Danish Council for Strategic Research under the Programme Commission for Nanoscience and Technology, Biotechnology and IT (NABIIT), the Netherlands Organization for Scientific Research (NWO), and AstraZeneca, Lund, Sweden.

Appendix A. Automated detection of leaks and extraction of branch length

The algorithm for leakage detection and branch length extraction is based on the algorithms presented in Schlathölter et al. (2002) and van Ginneken et al. (2008). Their observation was that a wave front propagating through a tree structure remains connected until it encounters a bifurcation, and any side branches can thus be detected as new disconnected components in the propagating front. We use the fast marching algorithm (Tsitsiklis, 1995; Malladi and Sethian, 1996) to propagate the wave front, similar to what was proposed in Schlathölter et al. (2002), with the front monitored through the set of "trial" points. For our purpose of evaluating a segmentation, we use a speed function that gives a value of 1 within the segmented structure and a value of 0 outside. This limits the propagation to only the segmented region.

Besides being able to detect bifurcations, the algorithm is also able to extract centerlines of individual branches and detect leakages. The centerline of a branch is constructed from the centroids of the propagation front at each time stamp. This also makes it possible to obtain the length of each individual branch by measuring the length of its centerline. Leakages were detected through a series of rules based on the geometric properties of the front, as described in Schlathölter et al. (2002) and van Ginneken et al. (2008). Similar to van Ginneken et al. (2008), we divided the leakage detection rules into two levels: a segment level and a tree level.

The segment level rules were applied when a branch was still being segmented by the fast marching algorithm, and the radius and connectivity of the propagation front were being monitored. At this level, three criteria were used for leakage detection:

$$D_{seg}(S) = \begin{cases} 1, & r_c/r_{prev} > 3 \\ 1, & r_c/min(R_{max}) > 1.5 \\ 1, & n_f > 5 \\ 0, & \text{otherwise} \end{cases} \quad (3)$$

where r_c is the radius of the current front, r_{prev} is the average radius of the previous five fronts, R_{max} is a set containing the maximum radius obtained at each of the ancestor branch segments and n_f is the number of disconnected fronts detected. The current segment, S, was discarded if $D_{seg}(S) = 1$. Furthermore, a sub segment was automatically accepted once its length exceeds 5 mm in order to prevent rejection of long segments due to small faults at the end.

The tree level rules were applied after all branch segments have been extracted. The criteria were defined as:

$$D_{tree}(S) = \begin{cases} 1, & n_b > 3 \\ 1, & \theta_{parent} > 100^\circ \\ 0, & \text{otherwise} \end{cases} \quad (4)$$

where n_b is the number of children and θ_{parent} is the angle between a segment and its parent. Segment S along with its children were removed when $D_{tree}(S) = 1$. All connected non-bifurcating segments, due to their children being rejected at some point in the process, were then merged into a single segment. Finally, all end segments with lengths less than 1 mm were considered as noise and were subsequently removed.

It should be noted that the values of the various parameters values were similar to those suggested in Schlathölter et al. (2002) and van Ginneken et al. (2008), and were selected on the basis of pilot experiments on a number of images in dataset 1.

References

Arya, S., Mount, D.M., Netanyahu, N.S., Silverman, R., Wu, A.Y., 1998. An optimal algorithm for approximate nearest neighbor searching fixed dimensions. J. ACM 45 (6), 891–923.

Aykac, D., Hoffman, E., McLennan, G., Reinhardt, J., 2003. Segmentation and analysis of the human airway tree from three-dimensional X-ray CT images. IEEE Trans. Med. Imaging 22 (8), 940–950.

Berger, P., Perot, V., Desbarats, P., de Lara, J.M.T., Marthan, R., Laurent, F., 2005. Airway wall thickness in cigarette smokers: quantitative thin-section CT assessment. Radiology 235 (3), 1055–1064.

Bülow, T., Wiemker, R., Blaffert, T., Lorenz, C., Renisch, S., 2005. Automatic extraction of the pulmonary artery tree from multi-slice CT data. In: Amini, A.A., Manduca, A. (Eds.), Medical Imaging 2005: Physiology, Function, and Structure from Medical Images, vol. 5746. SPIE, pp. 730–740.

Cover, T., Hart, P., 1967. Nearest neighbor pattern classification. IEEE Trans. Inform. Theory 13 (1), 21–27.

Coxson, H.O., Rogers, R.M., 2005. Quantitative computed tomography of chronic obstructive pulmonary disease. Acad. Radiol. 12 (11), 1457–1463.

Duda, R.O., Hart, P.E., Stork, D.G., 2001. Pattern Classification, second ed. Wiley-Interscience. pp. 174–177 (Chapter 4.4).

Fetita, C., Preteux, F., Beigelman-Aubry, C., Grenier, P., 2004. Pulmonary airways: 3-D reconstruction from multislice CT and clinical investigation. IEEE Trans. Med. Imaging 23 (11), 1353–1364.

Gorbunova, V., Lo, P., Ashraf, H., Dirksen, A., Nielsen, M., de Bruijne, M., 2008. Weight preserving image registration for monitoring disease progression in lung CT. In: MICCAI. Lecture Notes in Computer Science, vol. 5242, pp. 863–870.

Graham, M.W., Gibbs, J.D., Higgins, W.E., 2008. Robust system for human airway-tree segmentation. Medical Imaging 2008: Image Processing, vol. 6914. SPIE. p. 69141J.

Hu, S., Hoffman, E., Reinhardt, J., 2001. Automatic lung segmentation for accurate quantitation of volumetric X-ray CT images. IEEE Trans. Med. Imaging 20 (6), 490–498.

Kiraly, A.P., Higgins, W.E., McLennan, G., Hoffman, E.A., Reinhardt, J.M., 2002. Three-dimensional human airway segmentation methods for clinical virtual bronchoscopy. Acad. Radiol. 9 (10), 1153–1168.

Kitasaka, T., Mori, K., Suenaga, Y., Hasegawa, J., Toriwaki, J., 2003. A method for segmenting bronchial trees from 3D chest X-ray CT images. In: MICCAI. Lecture Notes in Computer Science, pp. 603–610.

Kuhnigk, J.M., Hahn, H., Hindennach, M., Dicken, V., Krass, S., Peitgen, H.O., 2003. Lung lobe segmentation by anatomy-guided 3D watershed transform. Medical Imaging 2003: Image Processing, vol. 5032. SPIE, pp. 1482–1490.

Lee, Y.K., Oh, Y.-M., Lee, J.-H., Kim, E.K., Lee, J.H., Kim, N., Seo, J.B., Lee, S.D., KOLD Study Group, 2008. Quantitative assessment of emphysema, air trapping, and airway thickening on computed tomography. Lung 186 (3), 157–165.

Li, B., Christensen, G.E., Hoffman, E.A., McLennan, G., Reinhardt, J.M., 2008. Pulmonary CT image registration and warping for tracking tissue deformation during the respiratory cycle through 3D consistent image registration. Med. Phys. 35 (12), 5575–5583.

Lindeberg, T., 1998. Feature detection with automatic scale selection. Int. J. Comput. Vis. 30 (2), 79–116.

Lo, P., de Bruijne, M., 2008. Voxel classification based airway tree segmentation. Medical Imaging 2008: Image Processing, vol. 6914. SPIE, p. 69141K.

Lo, P., Sporring, J., Ashraf, H., Pedersen, J., de Bruijne., M., 2008. Vessel-guided airway segmentation based on voxel classification. In: Proceedings of the First International Workshop on Pulmonary Image Analysis, pp. 113–124.

Lo, P., Sporring, J., de Bruijne, M., 2009a. Multiscale vessel-guided airway tree segmentation. In: Proceedings of the Second International Workshop on Pulmonary Image Analysis, pp. 323–332.

Lo, P., van Ginneken, B., Reinhardt, J., de Bruijne, M., 2009b. Extraction of airways from CT (EXACT'09). In: Second International Workshop on Pulmonary Image Analysis, pp. 175–190.

Malladi, R., Sethian, J., 1996. Level set and fast marching methods in image processing and computer vision. In: Proceedings of the International Conference on Image Processing, vol. 1, pp. 489–492.

Mayer, D., Bartz, D., Fischer, J., Ley, S., del Río, A., Thust, S., Kauczor, H.-U., Heussel, C.P., 2004. Hybrid segmentation and virtual bronchoscopy based on CT images. Acad. Radiol. 11 (5), 551–565.

Mori, K., Hasegawa, J., Toriwaki, J., Anno, H., Katada, K., 1996. Recognition of bronchus in three-dimensional X-ray CT images with applications to virtualized bronchoscopy system. In: Proceedings of the 13th International Conference on Pattern Recognition, vol. 3, pp. 528–532.

Mori, K., Nakada, Y., Kitasaka, T., Suenaga, Y., Takabatake, H., Mori, M., Natori, H., 2008. Lung lobe and segmental lobe extraction from 3D chest CT datasets based on figure decomposition and Voronoi division. Medical Imaging 2008: Image Processing, vol. 6914. SPIE, p. 69144K.

Murray, C.J., Lopez, A.D., 1996. Evidence-based health policy – lessons from the global burden of disease study. Science 274 (5288), 740–743.

Nakano, Y., Muro, S., Sakai, H., Hirai, T., Chin, K., Tsukino, M., Nishimura, K., Itoh, H., Paré, P.D., Hogg, J.C., Mishima, M., 2000. Computed tomographic measurements of airway dimensions and emphysema in smokers. Correlation with lung function. Am. J. Respir. Crit. Care Med. 162 (3 Pt. 1), 1102–1108.

Ochs, R.A., Goldin, J.G., Abtin, F., Kim, H.J., Brown, K., Batra, P., Roback, D., McNitt-Gray, M.F., Brown, M.S., 2007. Automated classification of lung bronchovascular anatomy in CT using AdaBoost. Med. Image Anal. 11 (3), 315–324.

Pedersen, J., Ashraf, H., Dirksen, A., Bach, K., Hansen, H., Toennesen, P., Thorsen, H., Brodersen, J., Skov, B., Døssing, M., Mortensen, J., Richter, K., Clementsen, P., Seersholm, N., 2009. The Danish randomized lung cancer CT screening trial – overall design and results of the prevalence round. J. Thorac. Oncol. 4 (5), 563–564.

Pisupati, C., Wolff, L., Zerhouni, E., 1996. Segmentation of 3D pulmonary trees using mathematical morphology. In: Proceedings of the Mathematical Morphology and its Applications to Image and Signal Processing, Atlanta, pp. 409–416.

Pudil, P., Novovičová, J., Kittler, J., 1994. Floating search methods in feature selection. Pattern Recogn. Lett. 15 (11), 1119–1125.

Rabe, K.F., Hurd, S., Anzueto, A., Barnes, P.J., Buist, S.A., Calverley, P., Fukuchi, Y., Jenkins, C., Rodriguez-Roisin, R., van Weel, C., Zielinski, J., Global Initiative for Chronic Obstructive Lung Disease, 2007. Global strategy for the diagnosis, management, and prevention of chronic obstructive pulmonary disease: gold executive summary. Am. J. Respir. Crit. Care Med. 176 (6), 532–555.

Schlathölter, T., Lorenz, C., Carlsen, I.C., Renisch, S., Deschamps, T., 2002. Simultaneous segmentation and tree reconstruction of the airways for virtual bronchoscopy. Medical Imaging 2002: Image Processing, vol. 4684. SPIE, pp. 103–113.

Singh, H., Crawford, M., Curtin, J.P., Zwiggelaar, R., 2004. Automated 3D segmentation of the lung airway tree using gain-based region growing approach. In: MICCAI. Lecture Notes in Computer Science, pp. 975–982.

Sonka, M., Park, W., Hoffman, E., 1996. Rule-based detection of intrathoracic airway trees. IEEE Trans. Med. Imaging 15 (3), 314–326.

Tschirren, J., Hoffman, E., McLennan, G., Sonka, M., 2005. Intrathoracic airway trees: segmentation and airway morphology analysis from low-dose CT scans. IEEE Trans. Med. Imaging 24 (12), 1529–1539.

Tschirren, J., Yavarna, T., Reinhardt, J., 2009. Airway segmentation framework for clinical environments. In: Proceedings of the Second International Workshop on Pulmonary Image Analysis, pp. 227–238.

Tsitsiklis, J.N., 1995. Efficient algorithms for globally optimal trajectories. IEEE Trans. Automat. Contr. 40 (9), 1528–1538.

Ukil, S., Sonka, M., Reinhardt, J.M., 2006. Automatic segmentation of pulmonary fissures in X-ray CT images using anatomic guidance. In: Reinhardt, J.M., Pluim, J.P.W. (Eds.), Medical Imaging 2006: Image Processing, vol. 6144. SPIE, p. 61440N.

van Ginneken, B., Baggerman, W., van Rikxoort, E., 2008. Robust segmentation and anatomical labeling of the airway tree from thoracic CT scans. In: MICCAI. Lecture Notes in Computer Science, vol. 5241, pp. 219–226.

van Iersel, C.A., de Koning, H.J., Draisma, G., Mali, W.P., Scholten, E.T., Nackaerts, K., Prokop, M., Habbema, J.D.F., Oudkerk, M., van Klaveren, R.J., 2007. Risk-based selection from the general population in a screening trial: selection criteria, recruitment and power for the Dutch–Belgian randomised lung cancer multi-slice CT screening trial (NELSON). Int. J. Cancer 120 (4), 868–874.

Wang, T., Basu, A., 2007. A note on 'A fully parallel 3D thinning algorithm and its applications'. Pattern Recogn. Lett. 28 (4), 501–506.
Weickert, J., 1999. Coherence-enhancing diffusion filtering. Int. J. Comput. Vis. 31 (2–3), 111–127.
Weickert, J., Ishikawa, S., Imiya, A., 1997. 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, pp. 45–59 (Chapter 4).
Zhou, X., Hayashi, T., Hara, T., Fujita, H., Yokoyama, R., Kiryu, T., Hoshi, H., 2006. Automatic segmentation and recognition of anatomical lung structures from high-resolution chest CT images. Comput. Med. Imaging Graph. 30 (5), 299–313.