Segmentation of Intrathoracic Airway Trees from Pulmonary CT Images

E. Høgh-Rasmussen[1], S. M. Aghito[2], B. Toftmann[3]
[1]IMM/DTU, [2]COM/DTU, [3]OFD/RISØ

Abstract— **Computed tomography (CT) image data have a good spatial resolution, and by stacking a number of 2D images it is possible to obtain high quality 3D thorax images. In diagnosing lung diseases it is highly desirable to be able to segment the lung into physiological structures, such as the intra-thoracic airway tree and the pulmonary structure.**

This paper describes a simple algorithm where the airway tree is being identified using a seed growing algorithm, which only requires manual placement of a single point.

I. Introduction

Modern technology is continually improving the treatment of diseases in all medical areas. When it comes to lung-related diseases the situation is no different, and as the development of non-invasive imaging systems keeps improving the devices and drives the cost down at the same time, the amount of data acquired for a patient is greater than ever before.

The use of X-rays for treatment of lung diseases has a long history, and it was used for diagnosing tuberculosis only a few years after Röntgen's discovery in 1895. After the invention of the CT-scanner, where it became possible to acquire X-rays of many thin slices through a whole volume, fig. 1, it became natural to use this tool to improve diagnosing lung diseases.

A modern CT-scanner is capable of scanning the whole thorax in a few seconds. In spite of this, there may still be motion artifacts, e.g. blurring, caused by breathing and heart beating. Several techniques have been developed to alleviate this problem. E.g. advanced CT acquisition routines can use an electrocardiogram (ECG) to synchronize the image acquisition with the heart beat, and similar measures can be used for breathing.

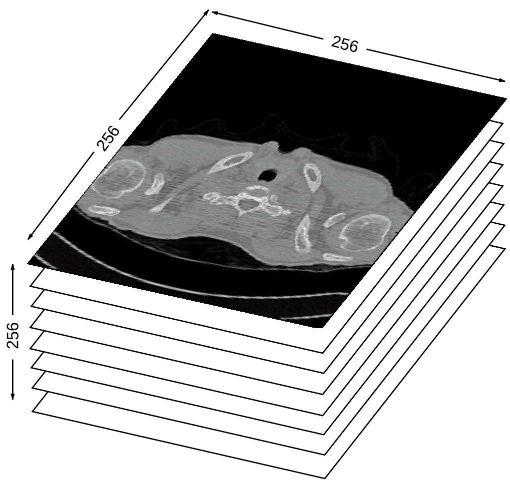

Fig. 1. The CT images as acquired for analysis.

The amount of data produced by a modern CT-scanner is overwhelming, if the full 3D volume must be analyzed as a part of the diagnosis. The analysis would be greatly facilitated if a segmentation of the lungs into separate structures, such as airway trees and pulmonary systems, were available. It is possible to conduct such a segmentation manually. However, since airways as small as 1 mm or smaller can be detected [1] it is clearly a very laborious task that must be performed by specialists. Therefore, it is necessary to automate the process if such methods are to be used in medical practice. The lower the amount of medical work, the higher is the chance that it is clinically useful.

This paper discusses a simple method for identifying the airway tree, which is based on seed-growing that is introduced in section II. We briefly discuss how the parameters were optimized and how an improved algorithm was obtained by considering neighborhoods rather than single voxels.

The algorithm were tested on 5 CT-volumes, and section IV discusses our experience and shows a few typical results as obtained by the algorithm. Also, we briefly describe how a minor modification may be made in order to identify the major lung vessels instead of the airway trees. We also briefly discusses some of the d properties of the modified algorithm. Finally we show how a successful segmentation can lead to an immediate diagnosis.

In section V we briefly mention more complex algorithms and applications, while section VI summarizes our findings during the project.

II. Seed Growing

The idea of segmenting the lungs by means of region growing is rather simple, but in practice several issues arises. It is based on the fact that the absorption coefficient of matter is roughly determined by the density, and since the density of air is low, it follows that X-rays travel almost undisturbed. In photographic X-ray imaging the film is exposed maximally in these areas, and they appear dark after development. This convention is usually adopted in CT-imaging, and is also assumed throughout this paper.

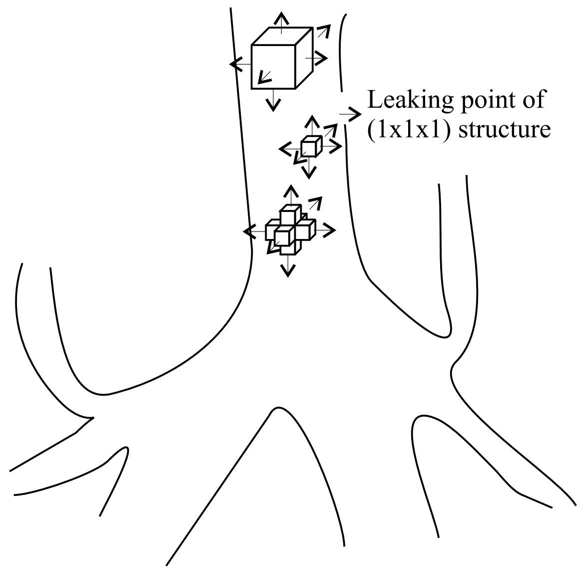

Fig. 2. Growing from single seed-point in the trachea.

This leads to the algorithm illustrated in fig. 2. First a seed point is selected manually and must be guaranteed to belong to the airways. The easiest place to start is in the upper part of the trachea, since the size makes it easy to identify. Since the airways must necessarily by enclosed by surrounding tissue, it is possible to grow the seed point in all directions that contains air, i.e. areas with a very low absorption.

Initially the segmented region consists of this single voxel. The region is then iteratively expanded by checking if any voxels adjacent to the region satisfy a connectivity condition, which can be defined in many different ways as discussed shortly. When the process ceases to include new voxels the result is a contiguous region of voxels that all satisfy the connectivity condition, except possibly for the initial seed point[1]. We decided to use a 6-connected region, i.e. voxels could be connected in 6 directions: up, down, north, east, south, west. In this way the algorithm progressively fills the trachea, including more and more voxels, and finally include the entire airway tree.

The simplest connectivity criteria is simply to accept very dark regions as being air, i.e. growing when the level was below a threshold level T. Looking at the CT-images a reasonable threshold would seem to be about $T = 20$. However, a problem arose because the lung tissue and the airways have a similar (low) attenuation constant. This is not the case for the walls separating the airways from the lungs, so in theory region growing should work. In practice, these walls are quite thin, so because of the limited resolution and the noise it is possible that some holes, as illustrated in fig. 2, may connect the airways to the rest of the lungs, leading to a leaking (fig. 3(b).)

It is therefore necessary to decrease the threshold T significantly in order to avoid this situation. For a pure growing a value of about $T = 2$ to 3 was necessary, depending slightly on the exact images being used. Nevertheless, if the threshold is adjusted properly, this simple algorithm gives surprisingly good results. In fig. 3(a)–3(b) we see the results for $T = 2$ and $T = 3$ for one image.

In the first case there is no leaking, and the airway tree has been identified to a reasonable resolution. In the latter case it is clear that T has been set to a too high value resulting in catastrophic leaking. Even though we can get reasonable results based on a simple threshold, we clearly needed a better method that is less sensitive to the actual images being analyzed.

A. An improved connectivity criteria

To improve on this matter one may use a larger volume than a single voxel, as also illustrated in fig. 2. In this case we require a larger volume to be below the threshold, greatly reducing the sensitivity to the actual value used. This can be compared with filling the airways using a larger pen, which can not slip through holes below that size.

This changed the situation dramatically as shown in fig. 3(c). Here was used a "ball" consisting of 19 voxels nearest to the center[2], and clearly this eliminated the leaking even for $T = 20$. The larger volume also had the side-effect of removing the finer structures of the airway tree, which would be unacceptable for most purposes.

However, there is a simple way to get the resolution of the careful chosen threshold and combine it with the increased robustness of the larger volume. One needs merely to allow the growing using either condition, as shown in fig. 3(d). This used $T = 20$ for the 19-voxel structure and $T = 1$ for single voxel growing, a value sufficiently low to avoid leaking in all the test images.

In fact it is possible to use a different value for any number of volumes used. Thus one can use a very conservative value for growing based on single voxels and use higher values as the volumes becomes larger. We have not pursued this matter further, as our intention with this paper is rather to illustrate the method itself.

B. Implementation

The algorithm was implemented in Matlab. The growing itself was implemented as a stack of edge voxels that had not been grown but had been identified as belonging to the airways. Also a structure was used to mark voxels that had been grown. The iteration then pop'ed a voxel from the stack and checked if any neighbor fulfilled the connectivity criteria without being marked. If any was found, the new voxel was pushed on the stack and the procedure was repeated until the stack was empty. Even though the scheme could have been improved with respect to memory, it never became necessary to do any further. It should be mentioned, however, that the iterative nature of the algorithm

[1] A point on the trachea wall should be allowed to grow inward.

[2] Consider a $3 \times 3 \times 3$ cube and remove the 8 corners.

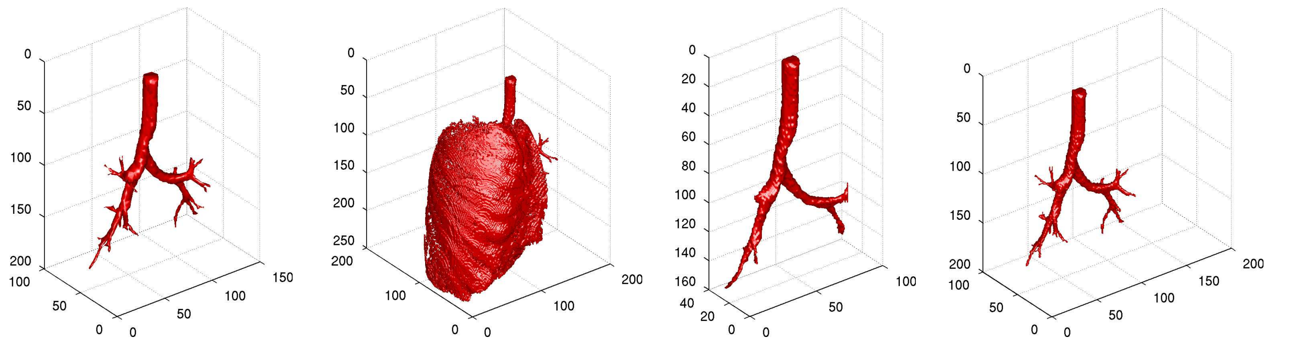

(a) Conservative threshold ($T = 2$.)

(b) Too high threshold causing leaking ($T = 3$.)

(c) When a ball is used, a much larger threshold can be used ($T = 20$.)

(d) Combining the approaches retains the advantages of both.

Fig. 3. The detected airway trees.

made Matlab rather inefficient, with a speed-up of perhaps a factor of 100 between version 6.1 and 6.5.

All the connectivity checking was performed in advance and stored in a Boolean array. To implement the large structure a rank-filter was applied (a maximum filter) followed by thresholding. It was interesting to note that a median-filter, which is normally expected to reduce noise [2], did not work at all.

C. Phantom

The size of the CT-images is very large, and they are of course affected by noise. During the process of designing and debugging an algorithm, it is often useful to test the algorithm on a smaller and simpler dataset. Hence, we created a simple phantom consisting of a $(16 \times 16 \times 16)$ 3D dataset, which could not be identified from a single pass in any direction.

Most of the phantom was filled with zeros and a few isolated gray level voxels. In the center of the volume there were several gray levels, composing a structure with vertical structures going down, branching horizontally, and further up and down. The purpose was to create a figure that was sufficiently complex that a simple algorithm would be unlikely to give the correct result. The algorithm was verified, and we found that it was able to detect the known central structure, so we proceeded to the real CT-images.

III. Vascular trees

With minor modifications the algorithm described in section II can, in principle, be used to identify the vascular tree instead of the airway trees. The vascular tree has the property that it has a relatively large attenuation coefficient compared to the surrounding lung tissue. Thus the upper threshold has to be replaced by a lower threshold. From the algorithmic point of view it becomes necessary to replace the max-filter with a min-filter to obtain the same effect.

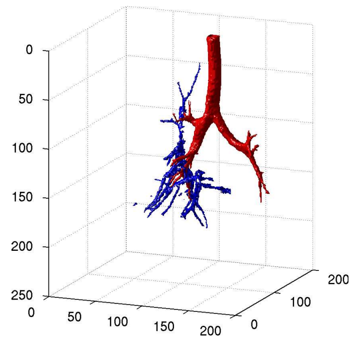

Fig. 4. The vascular tree superimposed on the airway tree using essentially the same algorithm.

A result of running this algorithm on a test-image is shown in fig. 4. The parameters used were a lower threshold of $T = 98$ for single voxels and $T = 90$ for the 19-connected volume described in section II-A.

However, identifying the vascular tree using this method is less useful than the airway trees. This is because the vascular tree is connected to the remaining part of the body and therefore there is a tendency of leaking. This is not something that can be fixed easily, because it reflects the anatomy and thus can not be considered as an error.

IV. Results

Most of the results have already been shown in fig. 3 and discussed in section II in order to explain the algorithm. It was surprising how well the airways tree could be identified using a single threshold. Nevertheless, the threshold necessary to obtain a reasonable resolution of the tree must be chosen very conservatively in order to work on all images, a situation that can not always be tolerated. An incorrect choice results in leaking, as shown in fig. 3(b).

As expected we saw that the algorithm was much less likely to leak when a larger volume was used, or equiva-

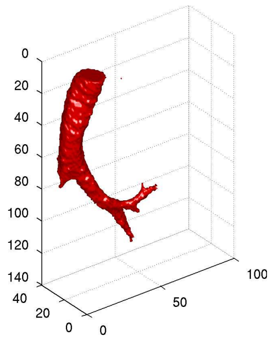

Fig. 5. A lung with the left bronchus obstructed.

lently, a rank filter was applied as a pre-processing step. In itself this resulted in a pruning of the tree (as evident by comparing fig. 3(a) and 3(c)), but combining the two approaches results in a much more versatile algorithm without the need for tuning the parameters for each image.

To show how an algorithm could potentially be useful in diagnosing pulmonary diseases we applied the algorithm to an images of a subject whose left main bronchus were obstructed. The result is shown in fig. 5, clearly demonstrating how easy the condition can be identified from a segmented airway tree.

V. Discussion

Several approaches to 3D segmentation of airway trees have been developed in the past, but with only limited success on *in vivo* data [3]. Two of the more successful approaches on *in vivo* data are rule-based methods [1] and fuzzy logic [3] that uses knowledge about airway morphology and airway inter-relationships with vessels. Good surveys on hitherto research are given in [4] and [5].

In this paper we have presented our results as static 3D pictures, but we have the possibility of rotating the images on the computer screen. However, this kind of visualization is often inadequate for surgical purpose. Virtual bronchoscopy used together with 3D CT images, pre-processed with central guidelines is a very useful technique that gives the physician the required image of the lungs [6].

We do not claim that the algorithm presented in this paper can compete with more advanced algorithms. However, it is often the case that simple methods are more robust than advanced ones, and the algorithm presented here is very easy to use and requires a minimum of manual intervention.

VI. Conclusion

This paper describes an algorithm for segmenting the airway trees from a seeded region growing, using a fixed threshold. This is a very simple method that leads to leaking unless the threshold value have been selected very conservatively.

An improvement of the region growing have been achieved using a larger structure (or ball) to search the airway. This makes the algorithm more tolerant to noise, since they will simply be overlooked.

By combining the methods it is possible to reach the same resolution as what can be obtained with a carefully chosen threshold. This is possible even though the parameters have been relaxed and remain constant across many images.

We finally conclude that even the simple algorithms presented here does a reasonable job with simple means. If better results are required, the risk of leaking must be reduced e.g. by identifying other tissues and using prior information about the structure of the airway trees.

References

[1] Milan Sonka, Wonkyu Park, and Eric A. Hoffman. Rule-based detection of intrathoracic airway trees. *IEEE Transactions on Medical Imaging*, 15(3):314–326, 1996.

[2] Rafael C. Gonzalez and Richard E. Woods. *Digital Image Processing*. Addison-Wesley, Reading, MA, USA, 3rd edition, 1993.

[3] Wonkyu Park, Eric A. Hoffman, and Milan Sonka. Segmentation of intrathoracic airway trees: A fuzzy logic approach. *IEEE Transactions on Medical Imaging*, 17(4):489–497, 1998.

[4] J. M. Reinhardt, R. Uppaluri, W. E. Higgins, and E. A. Hoffman. *Pulmonary Imaging and Analysis*, volume 2 of *Handbook of Medical Imaging: Medical Image Processing and Analysis, Editors: M. Sonka and J. M. Fitzpatrick*, chapter 16, pages 1005–1060. SPIE - The International Society for Optical Engineering, 2000. ISBN 0-8194-3622-4.

[5] M. Sonka, V. Hlavac, and R. Boyle. *Image Processing, Analysis, and Machine Vision*. PWS Publishing, second edition, 1999. Section 16.3, ISBN: 0-534-95393-X.

[6] R. D. Swift, A. P. Kiraly, A. J. Sherbondy, A. L. Austin, E. A. Hoffman, G. McLennan, and W. E. Higgins. Automatic axis generation for virtual bronchoscopic assessment of major airway obstructions. *Computerized Medical Imaging and Graphics*, 26(2):103–118, 2002.