A pattern recognition approach to enhancing structures in 3D CT data

Eva van Rikxoort and Bram van Ginneken
Image Sciences Institute, University Medical Center Utrecht
Heidelberglaan 100, 3584 CX Utrecht, The Netherlands
eva@isi.uu.nl bram@isi.uu.nl

Abstract

In medical image processing, several attempts have been made to develop filters which enhance certain structures in 3D data based on analysis of the Hessian matrix. These filters also tend to respond to other structures, e.g. most vessel enhancement filters also enhance nodule-like objects. In this paper, we use pattern recognition techniques to design more optimal filters. The essential difference with previous approaches is that we provide a system with examples of what it should enhance and suppress. These examples are used to train a classifier that determines the probability that a voxel in an unseen image belongs to the desired structures. The advantages of such an approach are excellent performance and flexibility: it can be used for any structure by providing the appropriate examples. We evaluated our approach on enhancing pulmonary fissures, which appear as plate-like structures in 3D CT chest scans. We compared our approach to the results of a recently proposed fissure enhancement filter. The results show that both methods are able to enhance the fissures, but our approach shows better performance; the areas under the ROC curves are 0.9044 and 0.7650, respectively.

Keywords: Hessian matrix, enhancement, supervised, pulmonary fissures

1. INTRODUCTION

Structure enhancement in volumetric data can be important to assist radiologists in the detection of abnormalities or diseases. To enhance structures in 3D CT data, many researchers have employed filters based on eigenvalues of the Hessian matrix or the structure tensor. For example, Frangi et al.[1] and Agam et al.[2] used a tubular structure enhancement filter to enhance vessels, and Wiemker et al.[3] used a plate-enhancement filter to enhance pulmonary interlobar fissures. Most of the enhancement filters proposed so far, are effective in enhancing a certain structure, but not in all cases (e.g. vessel enhancement filters typically show a low response on bifurcations) and they tend to enhance other structures as well (e.g. a vessel filter may yield a higher response on a strong edge than on a low contrast tubular structure). The purpose of this work is to show that supervised filters, that is, filters that are constructed using example input and output data and classifiers from pattern recognition theory,[4] can be superior to conventional structure enhancement filters.

We evaluate the supervised approach on enhancing pulmonary fissures. The human lung is divided into five regions called lobes which are separated by the lobar fissures. The left lung consists of two lobes, the right lung of three lobes. The lobar fissures are important anatomic landmarks in the interpretation of CT scans, radiologic identification of a lesion in relation to the fissures is important for precise localization of the lesion to the anatomic pulmonary lobes.[5] The lobar fissures may be incomplete, in which case the different lobes of the lung are connected. Next to lobar fissures, there are accessory fissures. The accessory fissures do not necessarily separate segments in the lung and are often incomplete. The accessory fissures in the lung are the most common variation of lung specimens.[6]

The pulmonary fissures appear as bright plate like structures on CT chest scans. We compare the results of our method to the results obtained by a fissure enhancement filter proposed by Wiemker et al.[3] In the next section the data used is described. In Section 3, the concept of supervised enhancement is introduced and both methods are described. In Section 4, experiments are described and results of both methods are provided. In Section 5 discussion and conclusions are given.

Medical Imaging 2006: Image Processing, edited by Joseph M. Reinhardt, Josien P. W. Pluim,
Proc. of SPIE Vol. 6144, 61441O, (2006) · 0277-786X/06/$15 · doi: 10.1117/12.651138

2. DATA

In this study, 16 high resolution CT chest scans (of 16 different patients) were used. The scans were acquired on a Philips Mx8000IDT scanner, 16×0.75 mm collimation, $100 - 175$ mAs, voxel resolution typically $0.7 \times 0.7 \times 0.7$ mm. On all scans, the lung fields were segmented with an automatic algorithm comparable to the algorithms described by Hu et al.[7] and Sluimer et al.[8] The set of 16 scans was divided into a test and a training set both containing 8 images.

A human observer manually indicated the lobar fissures in every fourth coronal slice. Since fissures are often hard to distinguish using only 2D information, before actually segmenting the fissures on the coronal slices, a possibility was given to scroll through the scan in the axial and sagittal direction and indicate marker points on each individual fissure. These points were visible in the coronal direction while segmenting the fissures. Segmenting fissures was performed by clicking points on the fissure; between two points, a straight line was automatically drawn.

3. METHOD

We implemented two fissure enhancement methods. A method for unsupervised fissure enhancement as described by Wiemker et al.[3] and a supervised method. Unsupervised, in this context, refers to a method that is not trained with examples but decides which voxels to enhance and suppress using a rule-based strategy. In the remainder of this section first the unsupervised approach is described. Next, the general framework for supervised enhancement is described, followed by the specifications of the system for fissure enhancement.

3.1. Unsupervised fissure enhancement

As comparison to our supervised method, we implemented the unsupervised filter proposed by Wiemker et al.[3] This filter uses second order information to enhance structures. In a plate like structure like a fissure, one strong gray-level curvature is expected perpendicular to the plate. Parallel to the plate small curvatures are expected. To determine these curvatures, the eigenvalues of the Hessian matrix are used. For each voxel, the Hessian matrix is build form the six independent second order derivatives:

$$H = \begin{bmatrix} g_{xx} & g_{xy} & g_{xz} \\ g_{yx} & g_{yy} & g_{yz} \\ g_{zx} & g_{zy} & g_{zz} \end{bmatrix}$$

When performing an eigenvalue analysis of the Hessian matrix, the principal directions in which the local second order structure can be decomposed are extracted. Since building the Hessian matrix requires second order derivatives to be taken from an image, a scale σ for the Gaussian kernel needs to be set.

In the fissure filter, a fixed combination of the two largest eigenvalues and Hounsfield values is used to determine a plateness value P that determines for each voxel the probability of being part of a fissure.

$$\begin{array}{rcl} P = & exp(\frac{I-\mu^2}{2\rho^2}) \cdot \frac{|\lambda_0|-|\lambda_1|}{|\lambda_0|+|\lambda_1|} & \text{, for } \lambda_0 < 0 \\ P = & 0 & \text{, otherwise,} \end{array}$$

where I is the intensity of the voxel, μ the average intensity for a fissure voxel and ρ the the standard deviation of the fissure voxel intensity. The plateness value becomes 1 if the largest eigenvalue is significantly higher than the second largest eigenvalue and 0 if the two largest eigenvalues are equal.

So, in the unsupervised filter, there are three parameters that need to be tuned (i) the scale of the derivative (σ), (ii) the average intensity of the fissure (μ), and (iii) the standard deviation of the fissure intensity (ρ). This tuning of parameters is a general problem of unsupervised methods. The performance of the method highly depends on the settings of these parameters and usually quite some tuning is required for different data or noise-levels.

3.2. Supervised enhancement

Supervised methods are popular for image segmentation. In this paper we propose to use this methodology for structure enhancement as well. Instead of a binary segmentation, the result will be an image in which the probability that each voxel belongs to the desired class is given.

In supervised enhancement, two stages can be distinguished, a training stage in which the system is developed, and a test phase in which the system is applied to previously unseen data. To construct a supervised system, a ground truth is required which gives for each voxel the class label (i.e. belonging to the class to be enhanced or not). In this study, this ground truth is established by a human observer. The training stage consists of the following steps:

1. Sample a number of voxels from the training images
2. For each sample, calculate a set of features and determine the corresponding class label using the ground truth. (This is the input to the system)
3. Train a classifier with the input feature vectors and class labels; after training, the classifier will be able to map unseen feature vectors to output. In order to be able to use this scheme for enhancement, the classifier should perform a soft classification (i.e. return the probability that an unseen pixel belongs to each class).

The test stage consists of the following steps:

1. Compute for each voxel in the test images the feature set.
2. Apply the trained classifier to get the probabilities that each voxel belongs to each class.
3. Construct the enhanced image by assigning to each voxel the probability that it belongs to the class to be enhanced (posterior probability).

A problem of supervised approaches is that they tend to get rather slow on 3D data because of the huge amount of input voxels. The computationally most expensive step is the classification of all the input voxels (step 2 in the test stage), the cost of this step depends on the number of features and the type of classifier used. To speed up the process, supervised methods can be divided into several phases. In the first phase, a cheap classifier (e.g. a linear discriminant classifier) and very few features are used. Voxels which have a very low posterior probability are regarded to be clearly non-structure voxels and are therefore discarded for the next phase(s). In the next phase(s), increasingly complex classifiers and features are used to determine the probabilities for the remaining voxels. After each phase, a certain amount of voxels is discarded for the next phase. The rationale behind this multiple-phase strategy is to use as little voxels as possible in the most complex classifier(s).

3.3. Supervised fissure enhancement

For the supervised fissure enhancement system, a two-step approach was used. As positive examples, voxels indicated by the human observer as being fissure voxels are used. An equal amount of randomly selected voxels elsewhere in the lungs were selected as negative examples.

In the first phase, a linear classifier which uses the eigenvalues of the Hessian matrix and the Hounsfield values on scale $\sigma = 1$ as features was applied. The resulting posterior probability image was thresholded at 0.5 to obtain the final result of this first step. In the posterior probability image, the fissure is again visible as a bright plate (see Figure 1).

In the next phase, a more complex and computationally more expensive K-nearest neighbor classifier with K = 30 is used to determine the probabilities for the remaining voxels. As features, the eigenvalues of the Hessian matrix as well as the gray-value were used at two scales, $\sigma = 1, 3$. Features for the second classifier are extracted from the posterior probability image from the first step.

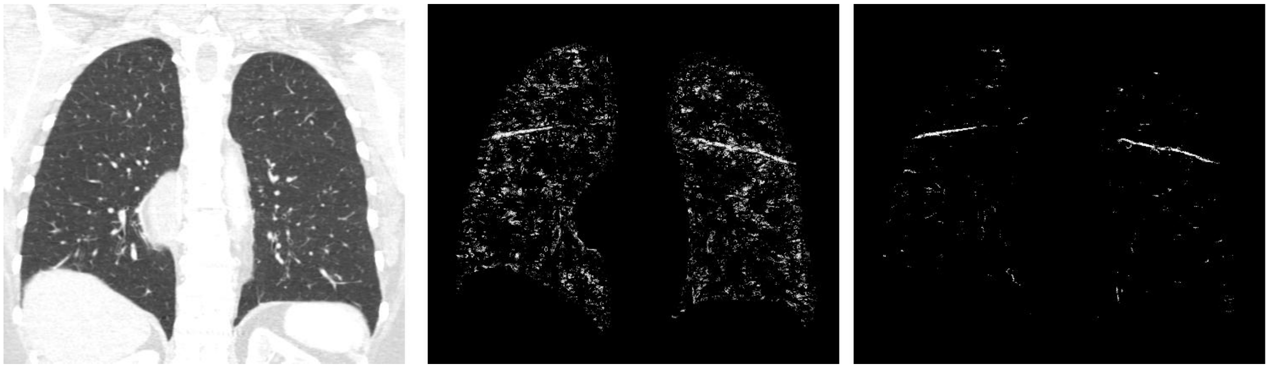

Figure 1. An example of the different steps of the supervised method. On the left, a coronal slice of an input scan. In the middle, the thresholded result of the first step. The fissures show up as bright lines, but there are also a lot of spurious responses. On the right, the final result after the second, final step. The fissures are very clear and there are almost no other responses.

4. EXPERIMENTS AND RESULTS

Both our method and the unsupervised filter were applied to eight independent test scans in which a human observer has indicated fissures to allow for quantitative evaluation. The fissures were drawn in every fourth coronal slice, these slices were used for evaluating the enhancement methods. In Figure 2 results from both approaches are shown.

Both filters are effective in fissure enhancement, the mean output was significantly higher for voxels on fissures then elsewhere in the lung for both methods ($p < 0.001$). Common measures for evaluating enhancement and segmentation results are accuracy, sensitivity and specificity. However, accuracy and specificity are not very useful in the case of fissure enhancement since the majority of pixels is background. For example, if all pixels are assigned to be background, still an accuracy of around 99% will be measured. Therefore we chose to do an ROC-analysis since in an ROC analysis the trade-off between specificity and sensitivity is measured. In Figure 3 the ROC curves for both methods are shown. The area under the curve (A_z) is 0.9044 for our method and 0.7650 for the unsupervised filter, which indicates that our method shows much less response on non-fissure structures. Although only the interlobar fissures were traced by a human observer, visual inspection of the results shows that also the accessory fissures are enhanced quite well (see Figure 4). The computational time required for the supervised method is four times that of the unsupervised filter.

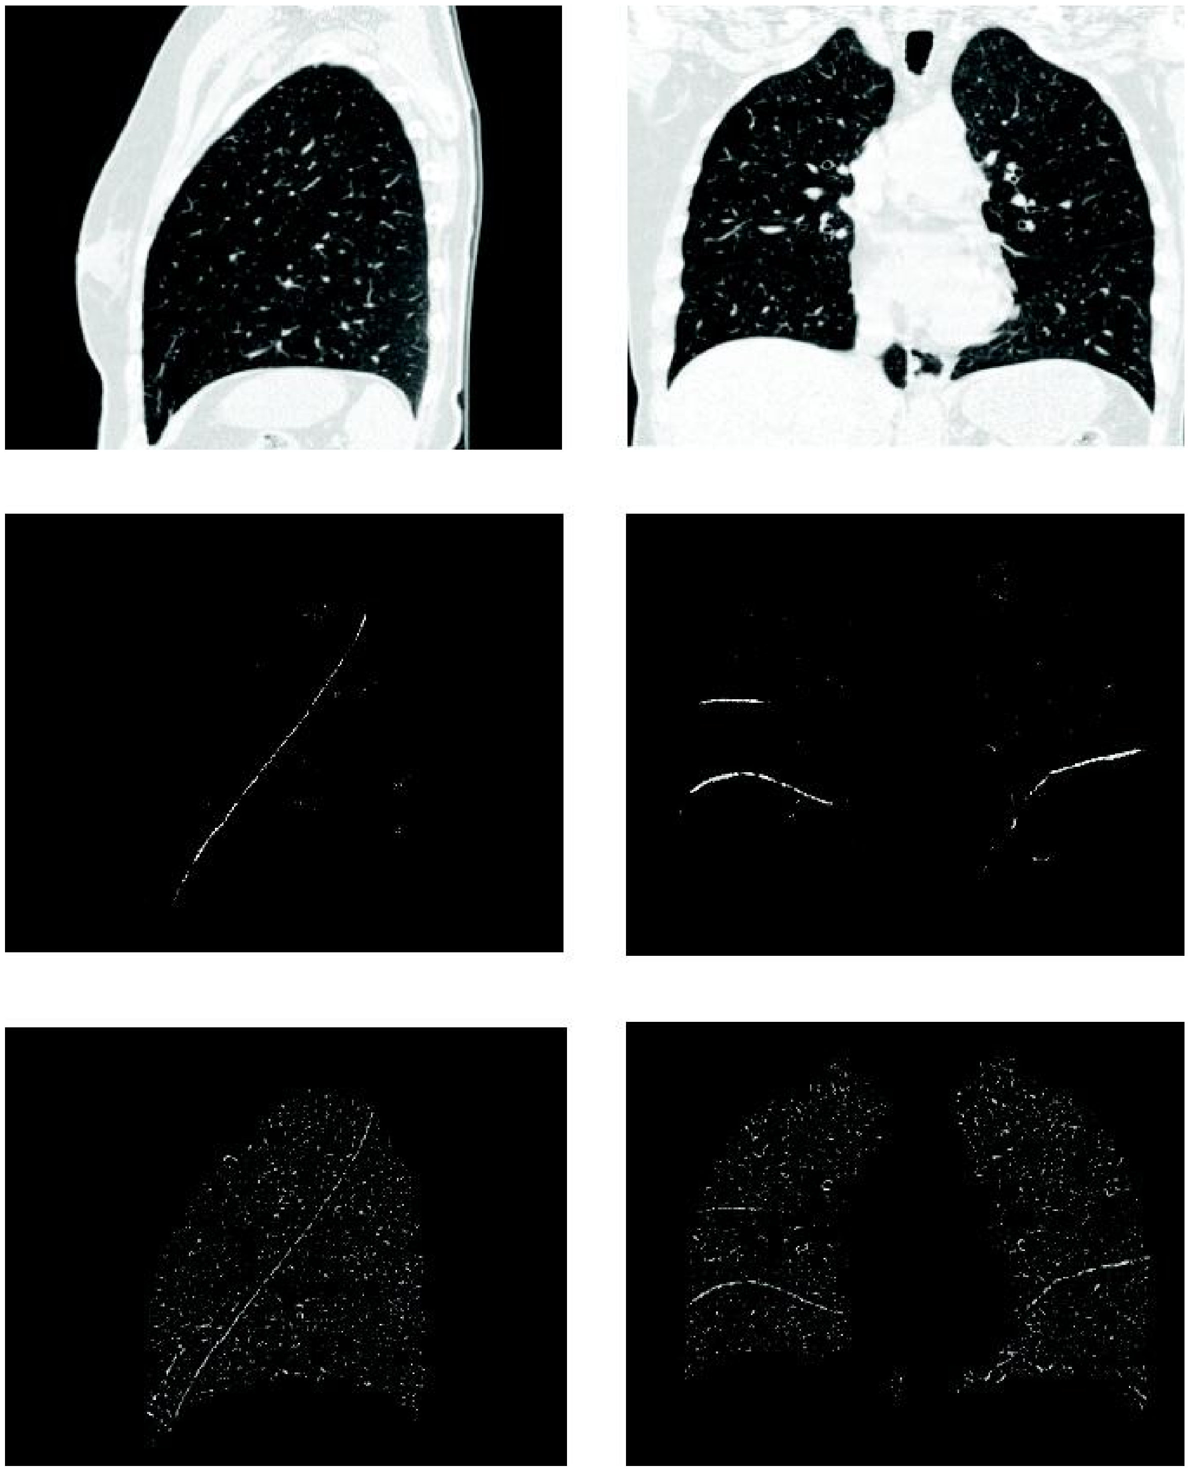

Figure 2. Example results of both enhancement methods. In the top row a sagittal and coronal slice of a test input scan are shown. In the middle row, the results of the supervised method are shown. It is clear that on these slices, there are almost no other responses next to fissure voxels. In the bottom row, results from the unsupervised filter are shown, although the fissures are clearly enhanced, there are a lot of other responses as well.

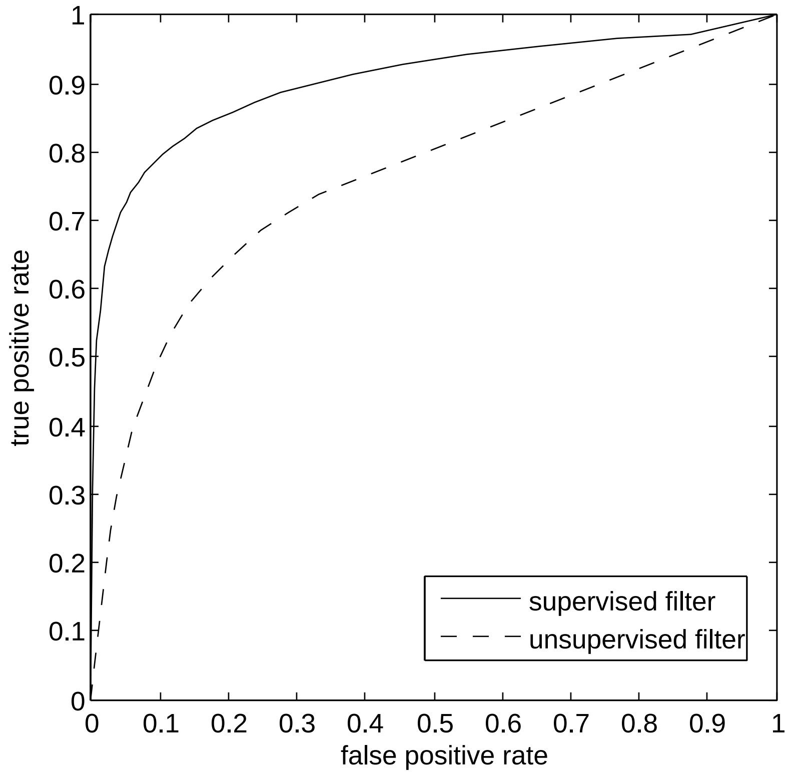

Figure 3. ROC curves for our method (supervised filter) and the filter proposed by Wiemker et al. (unsupervised filter). The areas under the curves are 0.9044 and 0.7650 respectively.

5. DISCUSSION & CONCLUSIONS

In this paper, a pattern recognition approach to enhancing structures was presented and applied to enhancing pulmonary fissures. Although many attempts have been made to enhance structures in 3D CT data using the eigenvalue analysis of the Hessian matrix,[1, 2, 9] this is the first paper in which a pattern recognition approach is applied and compared to a recently published filter approach. The results demonstrate that the supervised approach is superior.

There are multiple advantages of a supervised approach as compared to an unsupervised approach. The supervised approach can be applied to any structure given the appropriate examples, in an unsupervised approach a new system needs to be designed for every structure. In addition, supervised approaches are usually more robust to different appearances of structures since the system learns from examples which variations to expect. In an unsupervised approach, usually the most common appearance is taken as a reference to design the system.

A drawback of unsupervised approaches is the setting of different parameters. In the unsupervised fissure filter used in this paper, three parameters were present; the scale used for calculating the Hessian matrix, the average intensity of the fissure voxels, and the standard deviation of the intensity of the fissure voxels. The last two parameters make sure only voxels with the correct Hounsfield value are enhanced. The setting of these parameters is important to distinguish fissures from other bright plate-like structures in the lung. However, determining the correct settings for those parameters is hard since segmentation of fissures is needed to determine these values.

The other parameter to be set, the scale of the Gaussian kernel used for differentiation, is even more important. Fissures are typically between one and three voxels in width, so when a large scale is used, fissures may not be distinguishable anymore. However, when the scale used is too small, noise will be influencing the results. In supervised approaches, there are of course also some parameters present, however the performance of the system is less sensitive to the setting of these parameters.

The superior results of the supervised filter as opposed to the unsupervised filter was mainly due to the fact that there were much less spurious responses. Of course, post processing can be applied to the unsupervised filter to get rid of these small responses. However, this will introduce another step in which many parameters need to be set, which makes the system more complex and more sensitive to different types of data.

One of the main disadvantages of using a supervised approach is the fact that manual segmentations are required to be able to train the system. Manually segmenting structures on 3D CT data is labor intensive, when new data is used, e.g. from a different scanner or obtained with different acquisition parameters, new training data may be required.

Another disadvantage of supervised approaches is that they tend to become slow on large data sets. To overcome this problem, a two-step approach to supervised enhancement was introduced in this paper. First a simple classifier is used to discriminate between clearly non-fissure voxels and other voxels. Next, a more complex classifier is applied on the result of the first step to determine the probabilities for the remaining voxels. The two-step approach has two advantages, (i) it is six times faster than using only one complex classifier and (ii), performance improves. When only the complex classifier is used with the same features, the area under the ROC curve is 0.8548, which is substantially worse than the area under the ROC curve obtained in the two-step approach (0.9044).

Acknowledgements

This research was supported by the Dutch organization for scientific research (NWO).

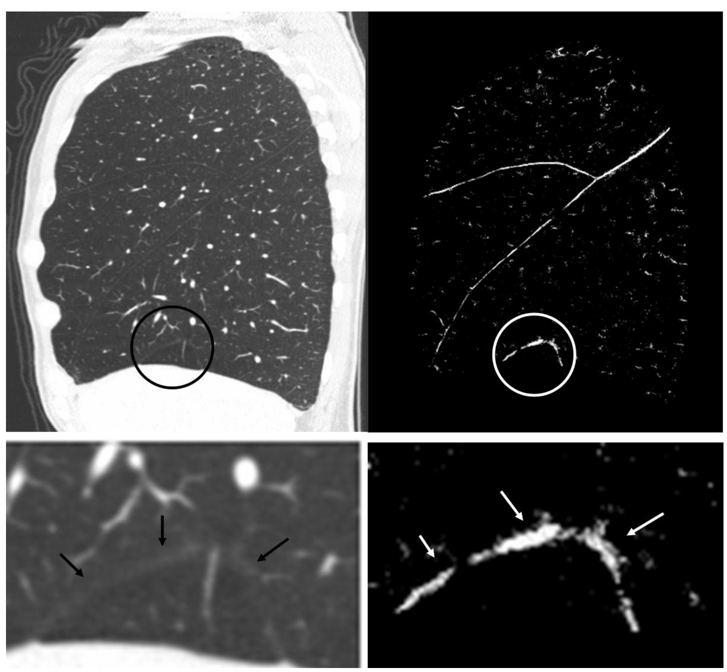

Figure 4. Example result of the supervised enhancement method. In the top row the original slice and the enhanced slice are shown. In the circle an accessory fissure is present, the area inside the circle is enlarged in the bottom row. As can be appreciated, this accessory fissure is also enhanced very effectively.

REFERENCES

1. A. Frangi, W. Niessen, K. Vincken, and M. Viergever, "Multiscale vessel enhancement filtering," in *Medical Image Computing and Computer-Assisted Intervention*, *Lecture Notes in Computer Science* **1496**, pp. 130–137, Springer Verlag, September 1998.
2. G. Agam, S. G. Armato III, and C. Wu, "Vessel tree reconstruction in thoracic CT scans with application to nodule detection," *IEEE Transactions on Medical Imaging* **24**(4), pp. 486–499, 2005.
3. R. Wiemker, T. Bülow, and T. Blaffert, "Unsupervised extraction of the pulmonary interlobar fissures from high resolution thoracic CT data," in *Computer Assisted Radiology and Surgery*, pp. 1121–1126, (Berlin), 2005.
4. R. O. Duda, P. E. Hart, and D. G. Stork, *Pattern Classification*, John Wiley and Sons, New York, NY, 2nd ed., 2001.
5. K. Hayashi, A. Aziz, K. Ashizawa, H. Hayashi, K. Nagaoki, and H. Otsuji, "Radiographic and CT appearances of the major fissures," *Radiographics* **21**(4), pp. 861–874, 2001.
6. A. Yildiz, F. Gölpinar, M. Calikoglu, M. N. Duce, C. Özer, and D. Apaydin, "HCRT evaluation of the accessory fissures of the lung," *European Journal of Radiology* **49**, pp. 245–249, 2003.
7. S. Hu, E. A. Hoffman, and J. M. Reinhardt, "Automatic lung segmentation for accurate quantitation of volumetric x-ray CT images," *IEEE Transactions on Medical Imaging* **20**(6), pp. 490–498, 2001.
8. I. Sluimer, M. Prokop, and B. van Ginneken, "Towards automated segmentation of the pathological lung in CT," *IEEE Transactions on Medical Imaging* **24**(8), pp. 1025–1038, 2005.
9. Q. Li, S. Sone, and K. Doi, "Selective enhancement filters for nodules, vessels, and airway walls in two- and three-dimensional CT scans," *Med Phys* **30**, pp. 2040–2051, 2003.