The First International Workshop on Pulmonary Image Analysis

New York, September 6, 2008

EDITED BY

Matthew Brown
Marleen de Bruijne
Bram van Ginneken
Atilla Kiraly
Jan-Martin Kuhnigk
Cristian Lorenz
Kensaku Mori
Joseph Reinhardt

Preface

Both the quantity and quality of image data available to study the pulmonary system have increased enormously in the last decade. The goal of this workshop is to bring together researchers in pulmonary image analysis and discuss recent advances in this rapidly developing field. We invited papers on all aspects of image analysis of pulmonary image data, including segmentation, registration, quantification, modeling of the image acquisition process, visualization, statistical modeling, biophysical modeling of the lungs (computational anatomy), and novel applications. In addition, we want to address the effective use of these methodologies for diagnosis and therapy in clinical applications, bringing together theory and practice, by including a hands-on demo session focusing on clinical workstations for pulmonary analysis.

We received many high quality submissions covering a broad spectrum of issues in pulmonary image analysis. All papers underwent a thorough review process with 3-4 reviews per paper by members of the program committee and additional reviewers. We finally accepted 12 papers for oral presentation, 16 poster presentations, and 3 papers describing software systems which will be demonstrated during the poster and demo session.

We would like to take this opportunity to thank the MICCAI 2008 organizers for help with organizing the workshop, David Naidich for agreeing to give the invited lecture, and the reviewers for helping us with the paper selection. We acknowledge the generous contributions of MeVis Medical Solutions, Philips Medical Systems, Siemens Corporate Research, and VIDA Diagnostics, which helped make this workshop possible.

<div align="right">

Matthew Brown
Marleen de Bruijne
Bram van Ginneken
Atilla Kiraly
Jan-Martin Kuhnigk
Cristian Lorenz
Kensaku Mori
Joseph Reinhardt

</div>

Organization

Organizers and Program Committee

Matthew Brown (UCLA Radiological Sciences, Los Angeles, USA)

Marleen de Bruijne (Department of Computer Science, University of Copenhagen, Denmark; Biomedical Imaging Group Rotterdam, Erasmus MC, Rotterdam, The Netherlands)

Bram van Ginneken (Image Sciences Institute, UMC Utrecht, The Netherlands)

Atilla Kiraly (Siemens Corporate Research, Princeton, USA)

Jan-Martin Kuhnigk (MeVis Research, Bremen, Germany)

Cristian Lorenz (Philips Research Europe-Hamburg, Germany)

Kensaku Mori (Department of Media Science, Graduate School of Information Science, Nagoya University, Japan)

Joseph Reinhardt (Department of Biomedical Engineering, The University of Iowa, Iowa City, USA)

Reviewers

Mohammad Aboofazeli

Deborah Aks

Christopher Alvino

Thomas Bülow

Tobias Boehler

Christophe Chefdhotel

Mohammad Dawood

Volker Dicken

Jin Fei

Vladlena Gorbunova

Jens Kaftan

Markus Krebs

Pechin Lo

Jan Moltz

Keelin Murphy

Carol Novak

Benjamin Odry

Eva van Rikxoort

Lauge Sørensen

Jon Sporring

Sotirios Tsaftaris

Thorsten Twellmann

Sponsoring

Table of Contents

Poster Presentations

Software Demonstrations

Keynote Statement

David P. Naidich

New York University - Langone Medical Center, NY, USA

The First International Workshop on Pulmonary Image Analysis provides a unique opportunity for researchers in pulmonary image analysis to explore cutting edge concepts in thoracic imaging. This welcome forum represents the interaction of recent technologic advances in hardware development, providing increasingly large imaging data sets (including the now widespread availability of state-of-the-art MR and multi-detector CT scanners) with ever more sophisticated software applications. The range and depth of the presentations in this Workshop, in particular, stand testament to the creativity of numerous investigators in the field of thoracic imaging.

It is also apparent that we currently stand at a cross-roads in which basic science research must be made more accessible for general clinical applications. Given current limitations in available financial support for scientific research as well as limitations in the timely transfer of sophisticated image processing concepts to clinically accessible imaging formats, it may be anticipated that closer dialogue between basic imaging researchers and clinicians will prove critical in insuring that key advances in diagnosis and management of patients with thoracic disease continues in as rapid a pace as possible. In this regard, identification of key areas of particular clinical concern and most daunting technical challenge may prove especially helpful by focusing the efforts of all on those problems most deserving of attention.

To this end, forums such as the First International Workshop serve an especially crucial function by facilitating an ongoing dialogue between basic researchers, clinicians and industry representatives. It is hoped that this will be the first on many such opportunities to come.

David Naidich, MD
Professor of Radiology and Medicine
New York University - Langone Medical Center

Oral Presentations

Texture Based Emphysema Quantification in Lung CT

Lauge Sørensen[1], Saher B. Shaker[2], and Marleen de Bruijne[1,3]

[1] Department of Computer Science, University of Copenhagen, Denmark,
{lauges, marleen}@diku.dk
[2] Department of Cardiology and Respiratory Medicine, Hvidovre University Hospital,
Copenhagen, Denmark
[3] Biomedical Imaging Group Rotterdam, Erasmus MC, Rotterdam, The Netherlands

Abstract. In this paper we propose to use texture based pixel classification in lung computed tomography (CT) for measuring emphysema. Two quantitative parameters for emphysema, based on the pixel classification, are suggested; relative class area and mean class posterior.
The approach is evaluated on a group of 39 patients, of whom 20 have been diagnosed with chronic obstructive pulmonary disease, using two different feature groups, local binary patterns and a filter bank based on Gaussian derivatives. The pixel classification based quantitative parameters correlate well with lung function ($r = 0.80$, $p < 10^{-5}$ for the parameter with the highest correlation) and correlate significantly better than the most commonly used CT based emphysema quantification method, namely relative area of low attenuation.

1 Introduction

Chronic obstructive pulmonary disease (COPD) is a major cause of death and a growing health problem worldwide. In the United States it is the fourth leading cause of morbidity and mortality and it is estimated to be ranked the fifth most burdening disease worldwide by 2020 [1]. COPD is a chronic lung disease characterized by limitation of airflow in the airway and it comprises two components: Chronic bronchitis, which is an inflammation of the small airways, and emphysema, which is characterized by gradual loss of lung tissue.

The primary diagnostic tools for COPD are lung function tests (LFT). Another diagnostic tool that is gaining more and more attention is computed tomography (CT) imaging. CT is a sensitive method for diagnosing emphysema and both visual and quantitative CT are closely correlated with the pathological extent of emphysema [2]. This makes CT suitable for both early detection and study of COPD as well as for monitoring the effect of different treatments.

We focus on the assessment of emphysema, which is thought to be the main cause of shortness of breath and disability in COPD. Emphysema is usually classified into three subtypes, or patterns, and we will adopt the naming and definitions from Webb et al. [3]. These subtypes are: Centrilobular emphysema (CLE) defined as multiple, small, spotty lucencies, that may have thin walls,

paraseptal emphysema (PSE) defined as multiple, lucencies in a single layer along the pleura, commonly with thin walls visible, and panlobular emphysema (PLE) defined as a lucent lung with small pulmonary vessels.

In CT emphysema lesions, or bullae, are visible as areas of abnormally low attenuation values, close to that of air. Different objective quantitative measures of emphysema can be derived from the histogram of CT attenuation values. The most common measure is the emphysema index or relative area of low attenuation (RA) [2], which measures the amount of lung parenchyma pixels that have values below a certain threshold relative to the total amount of lung parenchyma pixels. Common for the quantitative methods based on the attenuation histogram is that they ignore the possibly valuable information inherent in the emphysema disease patterns, such as subtype, shape, and size distribution.

One way to objectively analyze the properties of the disease patterns is to use texture analysis [4]. Several publications exist on characterizing emphysema and other disease patterns in regions of interest (ROI) in lung CT images using texture features [5–9]. In [5, 7, 8] the entire lung is labelled by subdividing the lung into adjacent ROIs followed by a classification step that assigns the same label to all pixels within a ROI. In [5, 7] the labelled result is evaluated by comparing the agreement between the output of the classification and that of human expert readings of the same ROIs. [8] Reports the percentage of different disease patterns present.

In this paper we propose a quantitative measure for emphysema, based on a pattern classification approach that utilizes local texture information. Compared to RA, a pattern classification approach allows for more than the two classes healthy and emphysematous, making it possible to quantify different subtypes of emphysema, which may be related to prognosis of the patient. Further, texture may be less influenced by inspiration level, compared to using intensity alone. We perform full lung classification by computing the posterior class probability for each pixel in the lung based on the local neighborhood around the pixel. Two ways of deriving a quantitative measure for emphysema, from the posterior, are investigated and evaluated. The first approach is to perform a hard classification and compute the relative area of the classes, the second approach is to compute the average posterior probability of each class in the lung. Thus, we obtain two measures per class. To our knowledge the use of the mean class posterior probability for quantifying emphysema is novel. We experiment with using two different kinds of features for this purpose, local binary patterns [10] and a filter bank based on Gaussian derivatives. These two feature groups were previously tested and evaluated on a set of hand picked ROIs, achieving an accuracy of 95.2% and 94.0% respectively in discriminating between normal tissue (NT), CLE, and PSE [9]. The experimental results reported in [9] are based on a subset of the CT images used in the experiments in this paper.

2 Methods

In the following the lung pixel classification system is described. Sections 2.1
and 2.2 describe the features we use for characterizing the lung texture, namely
local binary patterns and a filter bank based on Gaussian derivatives, Section
2.3 describes the classification framework, and Section 2.4 describes how the
classification framework is used for pixel classification and how the classification
result is turned into a quantitative measure for emphysema.

2.1 Local Binary Patterns

The first group of features is based on the local binary patterns (LBP) proposed
by Ojala et al. [10]. LBP measures the local structure at a given pixel by thresh-
olding P samples on a circle of radius R around the pixel using the intensity
in the pixel as threshold. The resulting thresholded samples are interpreted as
a binary number, that provides a unique code for each kind of local structure
or pattern. The operator is highly non-linear and detects microstructures in the
image at different resolutions governed by the parameter R, for example spots,
edges, corners, etc. Applying the LBP operator to an ROI results in an LBP
code image. Based on this an LBP histogram is formed by accumulating the
LBP codes directly into a histogram. We use the rotation invariant formulation
of LBP, see [9, 10] for more details. LBP are by design gray-scale invariant, and
this is not a desired property when dealing with CT images, where values are
measurements of a physical property of the tissue displayed. Therefore the dis-
tribution of the intensities is included, by forming the joint histogram between
the LBP and the intensities in the center pixels.

2.2 Gaussian function and its derivatives

The second group of features is based on the Gaussian function $G(\mathbf{x}; \sigma)$ and
combinations of derivatives of $G(\mathbf{x}; \sigma)$. We use $\mathbf{x} = [x, y]^T$ to denote the pixel
position. By varying the standard deviation σ of the function in a discrete man-
ner we obtain a whole bank of filters that can be applied to the image by
convolution. The Gaussian function itself is included to make the filter bank
sensitive to offsets in absolute intensity. The filters that we use are similar to
those used in [7], except that the filters we use are all rotation invariant. The
filter bank comprise the following filters: $G(\mathbf{x}; \sigma)$; the Laplacian of the Gaus-
sian $\nabla^2 G(\mathbf{x}; \sigma)$; the gradient magnitude $|\nabla G(\mathbf{x}; \sigma)|_2$; the Gaussian curvature
$K(\mathbf{x}; \sigma) = \partial^2 G(\mathbf{x}, \sigma)/\partial x^2 + \partial^2 G(\mathbf{x}, \sigma)/\partial y^2 - 2\partial^2 G(\mathbf{x}, \sigma)/\partial xy$. Feature histograms
are obtained by convolving the ROIs with each filter and making histograms of
each filter response.

2.3 Classification framework

Classification is done using the kNN classifier [11] with combined histogram
similarities as distance measure. We use histogram intersection as the similarity

Model	k	Window size	Features	Feature specific
Model 1	1	31×31 pixels	Joint LBP and intensity histogram, see Section 2.1	One resolution: R = 1, P = 8
Model 2	1	31×31 pixels	Histograms of filter responses of the Gaussian derivative based filters, see Section 2.2	Two filters: $G(\mathbf{x}, \sigma = 0.5)$ $\|\nabla G(\mathbf{x}, \sigma = 1)\|_2$

Table 1. Parameter settings for the two models used in the pixel classification experiments. k is the number of neighbors used in the kNN classifier. The unit of R and σ is pixels and the unit of P is samples.

measure between histogram H and histogram K

$$L_{hist}(H, K) = 1 - \sum_{b=1}^{N_b} \min(H_b, K_b), \tag{1}$$

where H_b denotes bin b of histogram H, N_b is the number of histogram bins and the histograms are assumed normalized to sum to one. In the case of measuring combined histogram similarity based on different histograms, e.g. different Gaussian based filters, the similarities are computed individually for each feature histogram using (1) and summed afterwards

$$L(\mathbf{x}, \omega_i, m) = \sum_{n}^{N_f} L_{hist}(f_n(\mathbf{x}), M_{m,n}^{\omega_i}), \tag{2}$$

where $M_{m,n}^{\omega_i}$ is the n'th feature histogram of prototype m from class ω_i, $f_n(\mathbf{x})$ denotes some function that extracts the local neighborhood around pixel $\mathbf{x}$ in the current image and computes the n'th feature histogram, and N_f is the number of feature histograms used in the combination. The histograms of the intensities and the filter responses are constructed using non-linear binning, where the binning is found by employing two rules on the total distribution of the ROIs in the training set: The total distribution should be approximately uniform and the number of bins is $\lfloor \sqrt[3]{N_p} \rfloor$, where N_p is the number of pixels in the ROI.

In [9] the accuracy of both feature groups is estimated as an average of a number of leave-one-patient-out experiments, in which the optimal filters and parameters are selected on the training set and can vary in between the experiments. Table 1 summarizes the parameter settings for the two feature groups that we will be using in this paper, which are those that were most often selected in the experiments in [9].

2.4   Quantification by pixel classification

Prior to classification, the lung parenchyma is extracted using a combination of thresholding, connected component analysis, and manual editing. The posterior probability of class ω_i given pixel $\mathbf{x}$ is computed based on the combined histogram similarity (2) with the closest prototype histograms of each class and is given by

$$P(\omega_i | \mathbf{x}) = \frac{\min_m L(\mathbf{x}, \omega_i, m)}{\sum_{j=1}^{3} \min_m L(\mathbf{x}, \omega_j, m)}. \tag{3}$$

The size of the local neighborhood is the same as the window size used in the model, i.e. 31×31 pixels, see Table 1, and the pixel being classified is in the center of this window. The classification result is a posterior class probability for each pixel in the lung.

A hard classification can be obtained by using the maximum a posteriori (MAP) rule in each pixel [11]. It should be noted that pixels that are not part of the lung segmentations are not classified, but they can still contribute to the classification, e.g. part of the exterior of the lung is in the local neighborhood when classifying a pixel at the border of the lung. In this way all potential relevant structural information is included, like being at the border of the lung or near large vessels and airways.

The quantitative measures for emphysema that we propose are the relative class area (RCA) and the mean class posterior (MCP). RCA is defined by the relative amount of pixels with a given class label, obtained using the MAP rule, divided by the total number of lung pixels N_l

$$\text{RCA}_{\omega_i} = \frac{1}{N_l} \sum_{j}^{N_l} \delta(\arg\max_c P(\omega_c|\mathbf{x}_j) - i),$$ (4)

where δ denotes the Kronecker delta function. MCP is given by averaging the posterior class probability of a given class, obtained using (3), across all pixels in the lung

$$\text{MCP}_{\omega_i} = \frac{1}{N_l} \sum_{j}^{N_l} P(\omega_i|\mathbf{x}_j).$$ (5)

3 Experiments and Results

3.1 Data

The data used for the experiments consists of a set of thin-slice CT images of the thorax. CT was performed using GE equipment (LightSpeed QX/i; GE Medical Systems, Milwaukee, WI, USA) with four detector rows, using the following parameters: In-plane resolution 0.78×0.78 mm, 1.25 mm slice thickness, tube voltage 140 kV, and tube current 200 milliampere (mA). The slices were reconstructed using a high spatial resolution (bone) algorithm. A population of 39 individuals, 9 healthy non-smokers, 10 smokers without COPD, and 20 smokers diagnosed with moderate or severe COPD according to LFT [1] were scanned in the upper, middle, and lower lung, resulting in a total of 117 CT slices. Visual assessment of the leading pattern, either NT, CLE, PSE, or PLE, and severity, ranging from 0 to 5, in each of these slices was done individually by an experienced chest radiologist and a CT experienced pulmonologist. In cases of disagreement, consensus readings were obtained. 216 non-overlapping ROIs were annotated in the slices representing the three classes: NT (107 observations, of which 48 were near the lung border or hilum area), CLE (50 observations), and PSE (59 observations). PLE was excluded due to underrepresentation in the data

	Model 2 NT	Model 2 CLE	Model 2 PSE
Model 1 NT	49.6	2.1	3.9
Model 1 CLE	2.4	18.8	3.5
Model 1 PSE	2.3	0.9	16.6

Table 2. Confusion matrix showing the percentage of pixel labels that the two models agree/disagree on.

set as only 2 out of the 20 individuals diagnosed with COPD had PLE as the leading pattern. These 216 ROIs were used as prototypes in the kNN classifier.

The 39 individuals also underwent LFT, performed according to the European Respiratory Society recommendations, prior to the CT scanning of the lungs. One widely used LFT is forced vital capacity in one second (FEV_1) which is the amount of air in liters that you can forcibly blow out in one second. FEV_1 can be adjusted for age, sex, and height by dividing by a predicted value according to these three parameters, thereby obtaining FEV_1%pred, and it is this LFT that we will use in the evaluation.

3.2   Lung pixel classification

Each of the 39 individuals were in turn measured using our proposed approach by classifying the pixels in each of the three CT slices using either Model 1 or Model 2, while leaving the prototypes coming from that individual out of the kNN classifier. The classification result is then used as a quantitative measure for emphysema by applying (4) or (5) to the posterior. Figure 3.2 shows the CT slices from two different patients, along with obtained pixel classifications and NT pixel posterior, when using Model 1. For comparison the RA below -910 HU (RA910) is shown in Figure 3.2 bottom-left.

3.3   Comparison of Model 1 and Model 2

The confusion matrix in Table 2 reveals that the two models generally are in good agreement; in 85% of the pixels, the two models agree on the class label. The highest level of disagreement is in the cases where Model 1 labels a pixel as NT or CLE whereas Model 2 labels that pixel as PSE, which happens in 3.9% and 3.5% of the cases. Correlating the class posteriors also shows a high degree of agreement between the two models, with $r = 0.93$ for NT, $r = 0.93$ for CLE, and $r = 0.91$ for PSE.

Two specific cases of disagreement are shown in Figure 2. In the first case Model 1 labelled half the pixels as CLE whereas Model 2 labelled these pixels mostly NT and partly CLE. Interestingly RA910 does not label one single pixel as emphysematous. Perhaps the models have picked up on a texture pattern and really found some emphysema, which was not possible using a simple threshold of -910 HU. Increasing the threshold to -860 HU reveals some low attenuation pixels partly in the areas labelled CLE by Model 1. Thus it seems that there is
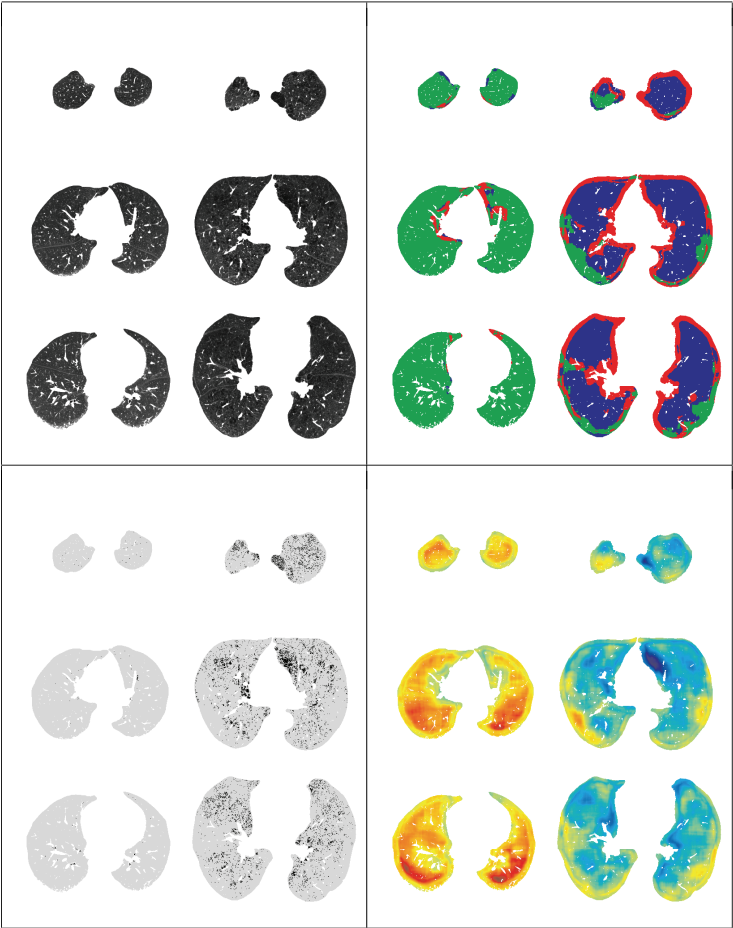
Fig. 1. Example images from two patients. The images are organized as follows: First column of each sub-box is a healthy non-smoker and second column is a patient diagnosed with moderate COPD according to LFT [1]. First row of each sub-box corresponds to the upper scan, second row to the middle scan, and third row to the bottom scan. **Top-left:** The segmented images used in the classification. The images are shown with the window setting -600/1500 HU. **Top-right:** The pixel classification result obtained when applying Model 1 to the lung segmentations shown top-left. Green corresponds to NT, blue to CLE, and red to PSE. **Bottom-left:** A threshold of -910 HU applied to the lung segmentation shown top-left. The areas below the threshold are indicated in black and the lung segmentations are indicated in light gray. **Bottom-right:** Posteriori NT probabilities. Dark red means that a NT prototype ROI is very similar to the given pixel's neighborhood in histogram feature space and corresponds to a NT probability of 0.48 when using (3). Dark blue means that all NT prototype ROIs are dissimilar and corresponds to a NT probability of 0.16. Refer to the electronic version for colors.
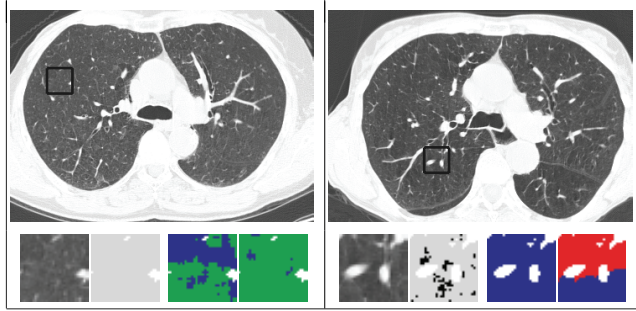
Fig. 2. Visual inspection of the classification results obtained using the two models in Table 1. **Left:** A non-smoker. **Right:** A smoker diagnosed with COPD. **Top:** The CT slice overlayed with a square indicating where the particular case is taken from. **Bottom:** Zoom in on the cases, from left to right; the original image, a threshold of -910 HU applied to the image with the areas below the threshold indicated in black and the lung segmentations indicated in light gray, the classification result obtained using Model 1, and the classification result obtained using Model 2. Green corresponds to NT, blue to CLE, and red to PSE. Refer to the electronic version for colors.

some emphysematous pattern present that Model 1 picks up. In the second case in the right part of Figure 2 Model 1 has labelled all lung pixels as CLE whereas Model 2 has labelled some CLE and some PSE. The ROI in the second case is from within the lung and thus by definition it should not be PSE. An explanation of why Model 2 labels many of the pixels PSE could be the large vessels seen within the ROI. Since we are using histograms, the spatial information in the ROI is ignored, and thus emphysematous regions with large vessels can share similarities with emphysema at the boundary in histogram feature space.

3.4 Relation to lung function

We evaluate the quantitative measures obtained from the pixel classification by correlating them with two other measures for emphysema, namely $FEV_1\%$pred, representing the classical objective way of measuring COPD by LFT, and an emphysema score (ES) computed by summing the visually assessed emphysema severity across the three slices. ES represents the subjective way of measuring emphysema by human visual assessment. For reference, we also compute RA910 for each of the patients and correlate that with $FEV_1\%$pred and ES. The Correlation results are reported in Table 3, where the correlations with $FEV_1\%$pred are computed using the Pearson correlation coefficient r and the correlations with ES are computed using the Kendall tau correlation coefficient τ.

RCA correlate well with $FEV_1\%$pred and generally also MCP. All texture based quantitative parameters for emphysema show significant correlation with

Measure	FEV$_1$%pred	ES	Separation
Model 1 RCA$_{NT}$	0.80 ($< 10^{-5}$)	-0.56 (0.006)	$< 10^{-5}$
Model 1 RCA$_{CLE}$	-0.77 ($< 10^{-5}$)	0.55 (0.006)	$< 10^{-5}$
Model 1 RCA$_{PSE}$	-0.77 ($< 10^{-5}$)	0.49 (0.067)	$< 10^{-5}$
Model 1 MCP$_{NT}$	0.73 ($< 10^{-5}$)	-0.56 ($< 10^{-5}$)	$< 10^{-5}$
Model 1 MCP$_{CLE}$	-0.58 (0.0001)	0.48 (0.0001)	0.0005
Model 1 MCP$_{PSE}$	-0.73 ($< 10^{-5}$)	0.50 ($< 10^{-5}$)	$< 10^{-5}$
Model 2 RCA$_{NT}$	0.79 ($< 10^{-5}$)	-0.56 (0.004)	$< 10^{-5}$
Model 2 RCA$_{CLE}$	-0.76 ($< 10^{-5}$)	0.53 (0.013)	$< 10^{-5}$
Model 2 RCA$_{PSE}$	-0.74 ($< 10^{-5}$)	0.46 (0.158)	$< 10^{-5}$
Model 2 MCP$_{NT}$	0.73 ($< 10^{-5}$)	-0.54 ($< 10^{-5}$)	$< 10^{-5}$
Model 2 MCP$_{CLE}$	-0.63 ($< 10^{-5}$)	0.48 (0.0001)	0.0001
Model 2 MCP$_{PSE}$	-0.69 ($< 10^{-5}$)	0.49 (0.0001)	$< 10^{-5}$
RA910	-0.62 ($< 10^{-5}$)	0.61 ($< 10^{-5}$)	$< 10^{-5}$
FEV$_1$%pred	-	-0.44 ($< 10^{-5}$)	$< 10^{-5}$

Table 3. Correlations with FEV$_1$%pred and ES as well as ability to separate patient groups according the a rank sum test. p-values of the correlations are shown in parenthesis next to the correlation coefficients.

FEV$_1$%pred and all except one, MCP$_{CLE}$ of Model 1, have a higher correlation with FEV$_1$%pred than RA910. RCA$_{NT}$ measured using Model 1 achieves the highest correlation of $r = 0.80$, which is significantly better than the correlation of $r = -0.62$ for RA910 ($p = 0.006$) according to a Hotelling/Williams test [12]. Note that we in the Hotelling/Williams test inverted the signs of the RA910 measures so that the two correlations being compared have the same sign. Looking at the correlations with ES, RA910 achieves the highest correlation coefficient ($\tau = 0.61$). Only two correlation coefficients are not significantly different from zero, and that is RCA$_{PSE}$ using both models. All the texture based quantitative parameters for emphysema, as well as FEV$_1$%pred and RA910, can separate non-smokers/healthy smokers from smokers diagnosed with COPD according to a rank sum test ($p \leq 10^{-4}$).

4 Discussion and Conclusion

It is not surprising that the texture based emphysema parameters perform differently than RA910, since the two approaches are very different in the amount of information they utilize. RA910 is based on a single threshold, -910 HU, and makes a decision for each pixel based only on the information in that particular pixel. The texture based emphysema parameters on the other hand, base the decision on all pixels in a local neighborhood and thus incorporate much more information. Further, the decision is not based on a specific threshold parameter, but is based on the distribution of the attenuation values as well as measurements of local structure. A consequence of this is that the proposed texture based approaches are expected to be less sensitive to differences in inspiration level. It is known that RA is sensitive to changes in inspiration level, since inspiration level influences the lung density, and thereby the CT attenuation values [13].

The insensitivity to inspiration level is something that is not investigated in this paper, but could be interesting to evaluate in the future on longitudinal data.

There is a tendency that RCA correlate better with FEV_1%pred than MCP, however the difference is not significant in all cases. This tendency could be due to uncertainty at the boundaries, evident in the bottom-right part of Figure 3.2, causing boundary effects as seen in the top-right part of Figure 3.2.

To conclude, we have proposed new parameters for quantifying emphysema in lung CT using texture based pixel classification. The proposed measures generally correlate well with lung function and the highest correlation, $r = 0.80$, is achieved by the relative normal tissue area.

References

1. Rabe et al.: Global strategy for the diagnosis, management, and prevention of chronic obstructive pulmonary disease: GOLD executive summary. Am J Respir Crit Care Med **176**(6) (Sep 2007) 532–555
2. Mller et al.: "density mask". an objective method to quantitate emphysema using computed tomography. Chest **94**(4) (Oct 1988) 782–787
3. Webb et al.: High-Resolution CT of the Lung, Third Edition. Lippincott Williams & Wilkins (2001)
4. Tuceryan, M., Jain, A.K.: Texture analysis. In: The Handbook of Pattern Recognition and Computer Vision (2nd Edition). World Scientific Publishing (1998) 207–248
5. Uppaluri et al.: Computer recognition of regional lung disease patterns. Am J Respir Crit Care Med **160**(2) (Aug 1999) 648–654
6. Chabat et al.: Obstructive lung diseases: texture classification for differentiation at CT. Radiology **228**(3) (Sep 2003) 871–877
7. Sluimer et al.: Automated classification of hyperlucency, fibrosis, ground glass, solid and focal lesions in high resolution ct of the lung. Medical Physics **33**(7) (2006) 2610–2620
8. Xu et al.: MDCT-based 3-D texture classification of emphysema and early smoking related lung pathologies. IEEE Trans Med Imaging **25**(4) (Apr 2006) 464–475
9. Srensen, L., Shaker, S.B., de Bruijne, M.: Texture classification in lung ct using local binary patterns. MICCAI **11**(Pt. 1) (2008) 934–941
10. Ojala et al.: Multiresolution gray-scale and rotation invariant texture classification with local binary patterns. IEEE Trans Pattern Anal Mach Intell **24**(7) (2002) 971–987
11. Jain et al.: Statistical pattern recognition: a review. IEEE Trans Pattern Anal Mach Intell **22**(1) (Jan. 2000) 4–37
12. Sickle, J.V.: Analyzing correlations between stream and watershed attributes. Journal of the American Water Resources Association **39**(3) (June 2003) 717–726 Errata: 41(3) 741–741.
13. Shaker et al.: Volume adjustment of lung density by computed tomography scans in patients with emphysema. Acta Radiol **45**(4) (Jul 2004) 417–423

Robust lung nodule growth measurement by combining registration and segmentation

Sven Kabus[1], Florian Müller[2], Rafael Wiemker[1], and Bernd Fischer[3]

[1] Philips Research Europe – Hamburg, Germany
sven.kabus@philips.com
[2] Institute for Signal Processing, University of Lübeck, Germany
[3] Institute of Mathematics, University of Lübeck, Germany

Abstract. Growth assessment for lung nodules is known to be the most relevant clinical parameter to distinguish malignant nodules from benign ones. The assessment of growth is usually done by multiple segmentations in follow-up CT examinations of a patient. Each segmentation is performed separately yielding a volume number assigned to the nodule at a certain point in time. Experiments have shown that the segmented volume may change due to small variations in certain voxel values. To circumvent this ill-posed problem, we present a growth assessment scheme which combines the result of an elastic registration and of a segmentation method. Regarding the absence of publicly available ground truth data, a method for the creation of artificial nodules is developed. Experiments show that the described combined growth assessment method leads to similar or better results than a growth assessment method achieved by a segmentation method only.

1 Introduction

The differential diagnosis of lung nodules is an essential step in the early detection of lung cancer. Besides shape properties, the most important criterion for the distinction between malignant and benign nodules is the growth rate. The growth assessment is usually done by comparing follow-up CT examinations of a patient, typically with a difference of three to six months. Currently, growth assessment schemes consist of a matching of the nodule positions in a first step, followed by a (semi-)automatic segmentation of the (same) nodule in both examinations as a second step, and, finally, a comparison of the two segmented volumes from which a change in volume can be computed. Of particular importance is the fact that each segmentation of the same nodule is performed *independently* from the other one. Experiments with follow-up data have shown that the segmented volumes may vary due to small variations in certain voxel values, e.g., if a nodule is attached to a vessel in the follow-up scan but not in the baseline scan. In this case a leakage of the follow-up segmentation into the vessel is reported [1]. To circumvent unwanted leakage into adjacent structures, detection of lung wall and vessel structures can be employed [2]. However, from the view of a vessel structure, the segmentation scheme may be faced again with
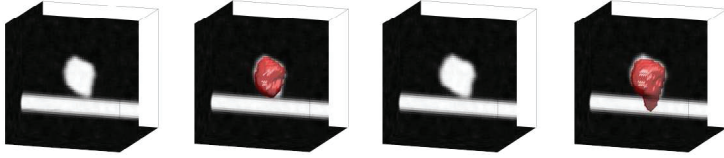
Fig. 1. Illustration of segmentation as an ill-posed problem: While the segmentation of an artificial vessel-attached nodule (left) is successful (center left), a slight modification at the nodule-vessel interface (center right) changes the segmentation result significantly (right).

a leakage problem, here from the vessel into the nodule. In particular, nodules which are heavily vascularized by vessels are affected. The leakage problem is even more prominent for small nodules, since a vessel-attached nodule with a similar or smaller diameter than the vessel diameter is likely to be characterized as protuberance of the vessel. From a mathematical point of view the segmentation of follow-up scans can therefore be expressed as an ill-posed problem, since small variations in the input data (the examinations themselves) may lead to large variations in the output data (the segmented volumes), cf. Figure 1.

In contrast to a segmentation, a regularized registration scheme is known to be a well-posed problem [3]. Its application to lung nodules enables for discrimination between expanding and shrinking lesions, for instance, by visually interpreting the displacement field and/or the according Jacobian map [4]. For a quantitative analysis we propose to combine the registration with a segmentation of the nodule which is required for the baseline scan only. The displacement field is then evaluated within the nodule, resulting in a single growth factor. Here, we compare the popular Jacobian map with two alternative criteria: the divergence of the displacement field and the quotient of the volume of the original and the deformed segmentation.

Validation should be based on ground truth data which, at first, are generated by multiple readers, at second, deal with follow-up data, and at third, are publicly available. Multiple reader delineations are, for instance, provided by the lung image database consortium (LIDC) [5]. However, no follow-up data are included in the LIDC database. Existing follow-up data known to us either originate from so-called coffee-break examinations (inter-examination time range from several minutes up to a few hours, therefore the nodule volume is assumed to not have changed between the scans) or are delineated by a single radiologist only, cf., e.g., [2]. Nevertheless, multiple reader delineations are advantageous, since inter-observer variability can be taken into account. For the LIDC data, an assessment revealed a median volumetric deviation of each delineation from the averaged delineation by 14% [6]. To our knowledge, there is no publicly available screening study including both follow-up data and delineations by multiple readers. Therefore, for validation and for comparison with a standard segmen-

tation method, artificial nodules are generated – either embedded into noisy background or into real-world data.

For the registration we use an elastic model based on the Navier-Lamé equations. Its principles and numerical implementation are introduced in Section 2. The second section finishes with the description of three growth assessment methods which combine the result of the registration with the result of a segmentation method. Section 3 describes the settings of the experiments. A method for generating artificial nodules as test cases is introduced here. The results are presented in Section 4 and concluded in Section 5.

2   Methods

Basic ideas. Given a reference and a template image, image registration tries to find a *displacement field* $u : \mathbb{R}^3 \to \mathbb{R}^3$ such, that the displaced template image minimizes both a certain similarity measure $\mathcal{D}$ and a regularizing term $\mathcal{S}$. Within this work, the sum of squared differences is used as similarity measure. By adding a regularizing term, the registration problem is a well-posed problem [3]. In general, a regularizer introduces assumptions about the likelihood of certain transformations. Within this context, a perfect regularizer would privilege transformations, which are likely to be the result of growth or shrinkage of a nodule. However, no general growth model for pulmonary nodules is known until today so that a more general regularizer has to be chosen. *Elastic registration* assumes that the images can be characterized as an elastic and compressible material and leads to the elastic regularizer [7]. The elastic properties are modelled by the so-called Lamé constants λ, μ. By using calculus of variations a solution for the regularized registration problem is calculated by computing a stationary point u of its Gâteaux derivative. This leads to a system of non-linear partial differential equations, known as Navier-Lamé equations,

$$\mathcal{A}[u] := \mu \Delta u + (\mu + \lambda) \nabla \cdot \nabla u = f \; , \tag{1}$$

where f is the Gâteaux derivative of $\mathcal{D}$.

Discretization & numerical solution. For the discretization of (1) finite differences in conjunction with Neumann boundary conditions have been chosen. A *time-marching iteration scheme* is employed resulting in the iterative solution of

$$(k\mathtt{A} + \mathtt{Id})\,vec(u^{(l+1)}) = k \cdot vec(f(u^{(l)})) + vec(u^{(l)}) =: \hat{f}(u^{(l)}), \, l \in \mathbb{N}_0 \; , \tag{2}$$

with k as temporal step size, $vec(\cdot)$ as lexicographically ordering operator and $\mathtt{A}\,vec(u)$ as discretization of $\mathcal{A}[u]$. For details, we refer to [3, 8]. This leads to a linear system of equations (LSE) with $3N$ unknowns, where N is the number of voxels of a single image. Since $(k\mathtt{A} + \mathtt{Id})$ does not depend on $u^{(l)}$, it has to be computed only once. Opposed to this, the term $\hat{f}$ depends on $u^{(l)}$ and needs

to be recomputed in every iteration step. For the calculation of the displaced template image a linear interpolation scheme turned out to be sufficient. The matrix $(k\mathtt{A} + \mathtt{Id})$ is sparse, symmetric [8] and highly structured with a typical size of $3 \cdot 40^3$ entries. For the solution of the LSE a conjugate gradient method has been used.

Starting with $u^{(0)} \equiv 0$ the step size k is initialized such, that the maximum displacement in the first iteration step is approximately the voxel grid width and set to be $k_{\max}$. Then, for each iteration step, the step size k is *adapted* under certain conditions. If $\mathcal{D}^{(l+1)} \geq \mathcal{D}^{(l)}$ then the step size is halved and the current iteration step repeated. If $\mathcal{D}^{(l+1)} < \mathcal{D}^{(l)}$ and $k + \Delta_k < k_{\max}$ then the step size is increased by Δ_k. This procedure works well within all our experiments. The iteration loop is stopped when the relative update of $\mathcal{D}$ is below a threshold, i.e., $(\mathcal{D}^{(l)} - \mathcal{D}^{(l+1)})/\mathcal{D}^{(l)} < 10^{-3}$. Throughout this work, the regularizing parameter λ has been set to zero and μ has been chosen as 0.01 being large enough to avoid any folding within the displacement u.

Estimating the growth rate. The displacement field u as result of the described registration scheme together with a segmentation image allows the growth assessment of a nodule. Here, a segmentation image is defined as a binary mask indicating the affiliation of an image element to the object of interest. The set of its image elements with a value of 1 is called foreground in the following. Assuming a segmentation of the reference image, three different criteria are employed to estimate the growth rate.

The first criterion is based on a *transformation of the segmentation image* according to the displacement field u. With a given sampling rate, the growth rate is assessed by comparing the volume of the original segmentation image with the volume of the deformed segmentation image.

The second criterion calculates the determinant of the *Jacobian* of u. It represents the local volume difference between reference and template image [9] and is given by $\det(Id + \nabla u)$. Taking the mean value over the foreground of the underlying segmentation image results in the estimated growth rate.

The third criterion is inspired by vector calculus and utilizes the *divergence* of the displacement field. Again, the growth rate is estimated by taking the mean value with respect to the segmentation.

3   Experiments

A general problem for the evaluation part of this work is the lack of ground truth data. Though, common public databases as the lung image database provided by the LIDC [5] contain delineations by various radiologists, they do not contain delineations for follow-up cases. Single-reader delineations, on the other hand, do not allow for assessing the inter-observer variability. Since experiments on the LIDC data revealed that this variability is not negligible (comparing the volume of each single radiologist's delineation to the mean radiologists' delineation, a median deviation of 14% could be observed [6]), a quantitative comparison is

difficult on the basis of a single delineation. Finally, phantoms currently do not reflect any vascularization around a nodule and therefore do not provide suitable ground truth for our purpose.

Consequently, an algorithm for *generating artificial nodules* and *simulating growth or shrinkage* has been developed. The advantage of this procedure is the availability of ground truth segmentations and the possibility to influence certain characteristics of the generated nodules, e.g., their size or surrounding. Each test case consists of two images, and each image contains one artificially generated nodule. These images may be seen as follow-up images of a CT scan. Two kinds of test cases have been created. While the A-cases are embedded into noisy background, the B-cases are embedded into real-world image data of lung tissue. Category A contains 246 test cases and category B contains 68 test cases. The volumes of the artificial nodules, which are used for the evaluation of the described growth assessment methods, range from 30mm^3 to 1529mm^3 according to a mean diameter ranging from about 4 to 14mm. The simulated growth factors range from 0.7 to 1.4. The voxel resolution was set to $0.6 \times 0.6 \times 1\text{mm}^3$.

Generating test cases. The generation of artificial nodules is inspired by [10] and utilizes the parametric representation as described in [6]. Briefly, the surface of a nodule is spanned by sample vectors which depend on a spherical basis model. With their origin on the sphere the sample vectors point into the direction of the surface normals. Radial basis function interpolation is used to describe the shape profile. For further details we refer to [6].

The creation of the test cases is explained next. Starting with an artificial nodule described by a set of parameters, a second nodule with the same center is generated by changing these parameters such that growth or shrinkage is simulated. This is done by randomly choosing a sample vector as deformation center on the surface of the nodule model followed by an in- or decrease of its length randomly. Proportional to the distance to the deformation center the according neighboring vectors are modified. The two models will be denoted as P_1 and P_2 in the following and contain the parameters needed for an explicit surface description. Once the models are calculated, corresponding binary masks are created and smoothed with a Gaussian kernel with variance σ which results in image masks $Q_1^\sigma(P_1), Q_2^\sigma(P_2)$ with range $[0, 1]$. Within our experiments, σ was set to 1. Given now a foreground image F with normally distributed values ($\mu_F = 0\text{HU}$, $\sigma_F = 60\text{HU}$) and a background image B, the final test case images $R_{1,2}$ are calculated by $Q_i^\sigma(P_i) \cdot F + \overline{Q_i^\sigma(P_i)} \cdot B$, $i = 1, 2$. Note that in case of category A, the background image B consists of normally distributed intensity values ($\mu_B = -860\text{HU}$, $\sigma_B = 60\text{HU}$), whereas cases of category B each use a subvolume of real CT data. Thus, category A contains only isolated nodules, whereas category B contains mainly nodules located next to vessels and/or the pleura or other pulmonary structures. Since the parametric model is given for each of the test cases, an according ground truth segmentation can be calculated for each of these. Examples of test cases of category B created with the described method are shown in Figure 2.
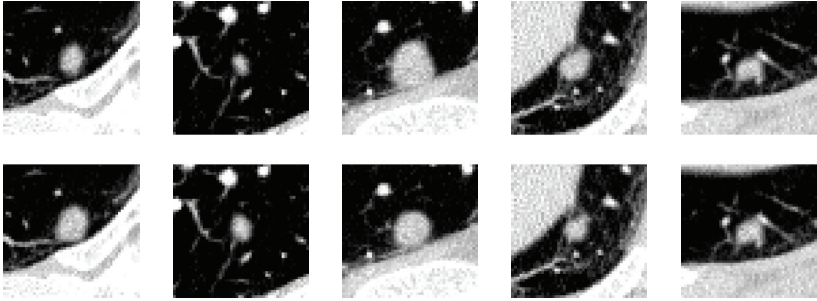
Fig. 2. Exemplary axial views (top row: baseline, bottom row: follow-up) of test cases of category B, i.e., artificial nodules embedded into real world image data.

Evaluating the methods. For the evaluation of the presented growth assessment methods the test cases of each category have been subdivided into three sets. The mean volume of both nodules of each test case has been used as leading quantity for this.

As described in Section 2 the proposed growth assessment requires a segmentation image. To distinguish the accuracy of the registration step from that of the segmentation step, we assume a correct segmentation of the baseline scan and, therefore, employ the ground-truth segmentation image for this specific scan.

To compare the three *combining* assessment methods, i.e., the methods which utilize one segmentation image and a displacement as result of an elastic registration, with a *non-combining* growth assessment scheme, a segmentation of both the baseline and the follow-up scan is computed separately with a state-of-the-art segmentation method [11]. This segmentation method has been evaluated on 33 nodules in a study at the Hematology-Oncology Clinic of Little Rock, Arkansas. Comparisons were made of manual segmentations between two radiologists, and manual vs. automatic volume segmentation. The agreement between manual and automatic volumetry proved to be equally good as the agreement between the two human readers [11].

4   Results

For the quantification of the performances of the growth assessment methods the <u>m</u>agnitude of the difference between the <u>r</u>eal and the <u>a</u>ssessed growth rate (MRA) has been chosen. The results for cases from category A and B are shown in Table 1. Arranged according to the nodule volume, mean and standard deviation of MRA are given for the standard segmentation method (SEGM) as well as for the registration method evaluated by deforming the segmentation image (DEF), calculating the mean of the Jacobian map (JAC), and determining the mean of the divergence (DIV).

Table 1. Mean (standard deviation) of the MRA in terms of the percentage for categories A and B.

		SEGM	DEF	JAC	DIV
A	$30 - 70\,\text{mm}^3$	4.9 (3.8)	7.1 (4.4)	7.0 (4.6)	7.0 (5.1)
	$71 - 172\,\text{mm}^3$	2.3 (2.2)	4.4 (3.5)	4.5 (3.6)	5.1 (3.9)
	$173 - 1529\,\text{mm}^3$	1.6 (1.3)	3.3 (2.7)	3.1 (2.6)	3.9 (3.1)
	all	2.9 (2.7)	4.6 (3.7)	4.6 (3.8)	5.1 (4.1)
B	$30 - 70\,\text{mm}^3$	12.3 (12.8)	5.3 (3.6)	5.8 (3.8)	5.8 (4.3)
	$71 - 172\,\text{mm}^3$	17.0 (27.7)	4.5 (3.6)	4.7 (3.8)	4.8 (4.1)
	$173 - 1529\,\text{mm}^3$	8.4 (13.4)	4.1 (3.3)	4.2 (3.1)	4.1 (3.3)
	all	13.1 (18.9)	4.6 (3.5)	4.8 (3.6)	4.8 (3.9)

In category A it can be observed that all methods perform better with an increasing mean volume of the nodules. Overall, SEGM leads to a MRA of 2.9%. In comparison to this, the MRA of the combining methods is about two percentage points larger. This can be explained by the fact that segmentation of an isolated nodule is a straightforward task. Consequently, its accuracy is expected to be higher than for the combining methods which involve two algorithmic steps and thus two sources of potential errors. In contrast, growth estimation for embedded nodules as in category B is a more complicated task from a segmentation point of view since leakage into adjacent vessels may occur (cf. Figures 1,2).

For category B, SEGM results in a MRA of 13.1% which is about four times larger than the MRA for category A cases but still in the range of inter-observer variability [6]. In contrast, the MRA of the combining methods stays below 5% and is nearly the same as for the category A cases.

Compared to each other, the three combining methods lead in both categories to similar results. Since the numerical calculation of DEF does not include the approximation of derivatives, this method may be superior to JAC and DIV from a numerical point of view. However, our results show only minor better results for DEF.

Even though the category A cases are supposed to be less challenging, the combining methods lead to slightly worse results for the set of nodules with a volume of 30 to 70mm^3 than the according category B cases. Further investigation will clarify if this observation is significant or if it is caused by the low number of test cases.

5   Conclusion

We presented a growth assessment method for lung nodules which combines registration and segmentation. Since small variations in the examination data may lead to large variations in the segmented volumes, segmentation of follow-up

scans can be seen as an ill-posed problem. In contrast, a regularized registration scheme is a well-posed problem. It is based on an elastic model that characterizes the image as an elastic and compressible material. In combination with a segmentation image of the baseline scan, the displacement field as result of the registration is evaluated and a growth factor estimated. Three different evaluation methods have been analyzed. Motivated by the lack of ground truth, artificially generated test cases have been used for validation – either as isolated nodules or embedded into real-world data. While for isolated nodules the proposed combined method performs slightly worse than a state-of-the-art segmentation method, the situation changes for the more challenging embedded nodules: Here, the proposed combined method outperforms the non-combining method by a factor of four. Moreover, this result is independent of the choice of the evaluation method.

References

1. Reeves, A.P., Chan, A., Yankelevitz, D.F., Henschke, C.I., Kressler, B., Kostis, W.J.: On measuring the change in size of pulmonary nodules. IEEE Transactions on Medical Imaging **25**(4) (2006) 435–450
2. Ko, J., Betke, M.: Chest CT: Automated nodule detection and assessment of change over time – preliminary experience. Radiology **218**(1) (2001) 267–273
3. Modersitzki, J.: Numerical Methods for Image Registration. Oxford University Press (2004)
4. Kawata, Y., Niki, N., Ohmatsu, H., Kusumoto, M., Kakinuma, R., Mori, K., Nishiyama, H., Eguchi, K., Kenko, M., Moriyama, N.: Analysis of evolving processes in pulmonary nodules using a sequence of three-dimensional thoracic images. Proceedings of SPIE: Image Processing **4322** (Apr 2003) 1890–1901
5. Armato III, S.G., McLennan, G., McNitt-Gray, M.F., Meyer, C.R., Yankelevitz, D., Aberle, D.R., Henschke, C.I., Hoffman, E.A., Kazerooni, E.A., MacMahon, H., Reeves, A.P., Croft, B.Y., Clarke, L.P.: Lung Image Database Consortium: Developing a Resource for the Medical Imaging Research Community. Radiology **232**(3) (2004) 739–748
6. Opfer, R., Wiemker, R.: A new general tumor segmentation framework based on radial basis function energy minimization with a validation study on LIDC lung nodules. Volume 6512., San Diego, CA, USA, SPIE (2007) 651217–10
7. Broit, C.: Optimal Registration of Deformed Images. PhD thesis, University of Pennsylvania (1981)
8. Kabus, S.: Multiple-Material Variational Image Registration. PhD thesis, University of Lübeck (2006)
9. Leow, A., Yanovsky, I., Chiang, M.C., Lee, A., Klunder, A., Lu, A., Becker, J., Davis, S., Toga, A., Thompson, P.: Statistical Properties of Jacobian Maps and the Realization of Unbiased Large-Deformation Nonlinear Image Registration. IEEE Transactions on Medical Imaging **26**(6) (2007) 822–832
10. Zhang, X., Chui, H., Olcott, E., Raffy, P., Yu, N., Giger, M.L., Karssemeijer, N.: Evaluation of lung nodule growth measurement for MDCT exams with different dosages using synthetic nodules. Medical Imaging 2007: Computer-Aided Diagnosis **6514**(1) (Mar 2007) 651438–8

11. Wiemker, R., Rogalla, P., Blaffert, T., Sifri, D., Hay, O., Shah, E., Truyen, R., Fleiter, T.: Aspects of computer-aided detection (CAD) and volumetry of pulmonary nodules using multislice CT. The British Journal of Radiology **78** (2005) 46–56

Segmentation of Juxtapleural Lung Nodules in CT Scans Based on Ellipsoid Approximation

Jan Hendrik Moltz, Jan-Martin Kuhnigk, Lars Bornemann and
Heinz-Otto Peitgen

MeVis Research GmbH – Center for Medical Image Computing, Bremen, Germany
jan.moltz@mevis.de

Abstract. This article presents a new algorithm for segmenting juxta-
pleural lung nodules in CT scans. Segmentation is an essential part of
volumetric therapy monitoring for cancer patients. Pulmonary nodules
that have extensive contact to the chest wall or other structures of similar
density are a special challenge for automatic segmentation. We propose
a ray casting approach to identify points at the visible boundary of the
nodule and then approximate its shape by an ellipsoid that is a least
squares fit of these points. The adjacent structures are cut off by mor-
phological processing within a dilated version of this ellipsoid. Evaluation
on 333 juxtapleural nodules showed that this method yields good results
and can be integrated easily into a general segmentation algorithm for
lung nodules with no substantial increase in computation time.

1   Introduction

In oncological therapy monitoring, the estimation of tumor growth from consec-
utive CT scans is an important aspect in deciding whether the given treatment
is adequate for the patient. Traditionally, this is done by measuring and compar-
ing the largest axial diameter of each lesion manually, but this approach implies
several problems. First, manual examinations are always subjective, error-prone
and time-consuming. Second and even more importantly, a 3d quantity (volume)
is estimated based on a 1d measurement (diameter). This simplification would be
valid if tumors were perfectly spherical and grew symmetrically but in practice
it leads to inaccurate results.

Although volumetry has the potential to enhance the accuracy and repro-
ducibility of growth estimation, measuring the lesion volume manually would
take too much time in the workflow of a radiologist. This is the motivation for
employing software assistants in oncological therapy monitoring since they are
able to perform automatic volume measurements in the 3d data. In order to be
accepted in clinical routine, they have to work both fast and accurately. Lesion
segmentation is an essential prerequisite for volumetry and efficient algorithms
are needed for different kinds of tumors.

In this article, we focus on the segmentation of juxtapleural lung nodules.
Pulmonary nodules are mostly located centrally within the lung parenchyma,
but they can also be attached to the pleura, a thin membrane that covers the

lungs. In CT scans of the thorax, the voxels can basically be divided into two density classes: while the dark ones represent the lung parenchyma, bright voxels may be nodules, but also blood vessels or structures adjacent to the lungs such as the chest wall, the heart, or the diaphragm (Fig. 1(a)). Since the pleura itself is invisible in CT images the boundary between a juxtapleural nodule and any of these structures shows little or no contrast and it is sometimes impossible even for a radiologist to determine the exact boundary of a nodule. Therefore the segmentation of this kind of lesions is particularly challenging; a mere threshold-based method or simple morphological processing is not sufficient.

In the following, all computations are restricted to a region of interest (ROI) whose center and size are determined by a user-defined stroke that is drawn across the nodule. We assume that the ROI contains the lesion completely and its center is close to the center of the lesion. Furthermore, we supersample the ROI to isotropic voxels if necessary.

2 State of the Art

Several authors that worked on solid lung nodule segmentation have also proposed solutions for juxtapleural nodules. An obvious idea is to compute a lung segmentation in order to separate the nodule from structures outside the lungs as it is done by Fetita *et al.* [1]. However, we decided not to incorporate any global information in order to keep computation times as low as possible and to be able to integrate the algorithm into an existing workstation. Van Ginneken [2] uses a local 2d lung field segmentation but no evaluation for juxtapleural nodules is given.

One of the first dedicated segmentation algorithms for juxtapleural nodules was presented by Shen *et al.* [3]. Assuming that the chest wall is physiologically smooth and that a nodule creates a "bump" with a high local curvature, the nodule can be separated by smoothing the wall surface. This is implemented by projecting the surface to a plane whose normal is the mean of all surface voxel normals. On the projection image, the nodule appears as a region of high values which are replaced by a cubic polynomial interpolation of the other values. The smoothed 3d surface is then computed via backprojection. This idea is promising but there will be problems when the wall itself has points of high curvature as in Fig. 2(c).

An algorithm proposed by Wiemker *et al.* [4] starts with region growing and determines the optimal cut-off value retrospectively by means of an objective function that separates the wall at its strongest inflection. Unfortunately, no results for juxtapleural nodules are shown.

Reeves *et al.* [5] approximate the pleural surface by a clipping plane that is iteratively refined until a leap in volume change is observed when it actually reaches the pleura. This works for small nodules where the actual convex or concave shape of the lungs can be ignored in a local view, but in other cases a plane is not suitable and the algorithm will fail.

Okada *et al.* [6] presented an approach that uses morphological opening similar to ours, but applied it in an inverse way: the size of the structure element is chosen such that the nodule is removed and the adjacent structure is retained. This presumes, however, that in the ROI the wall region is significantly larger than the nodule. Therefore the approach will fail when the nodule is large or when the wall has a concave shape. Furthermore, it is implemented such that segmenting a juxtapleural nodule takes more than twice as long as for a central one.

In a recent publication, Dehmeshki *et al.* [7] describe an algorithm that uses sphericity-constrained region growing on a fuzzy connectedness map. Successful results are reported for nodules "very close to lung wall or diaphragm" but the method always needs a visible contrast between the nodule and the adjacent structure.

The methods described in this article are based on the algorithm by Kuhnigk *et al.* [8] which has been designed to segment small and large nodules even with extensive vascularization. This is done by an initial region growing, followed by morphological opening with an adaptive erosion strength to remove blood vessels. As a preliminary solution for juxtapleural nodules, the algorithm makes use of the fact that the lungs are convex in most parts and that juxtapleural nodules create a concavity in this shape. The idea is to reconstruct the lung shape of the tumor-free state by computing the convex hull of the lung parenchyma within the ROI and cut off the nodule along the boundary of the convex hull. However, in regions where the lungs are not convex, such as the boundaries to the heart or the diaphragm, the convex hull is not suitable for this purpose since it does not remove convex parts of the attached structure completely. Furthermore, the results depend on the size of the ROI. This leads to poor segmentation results as shown in Fig. 2(d,e,f). In our contribution, we present an improvement to the algorithm that can handle this case as well. Parts of this work have previously been published in German [9].

3   Segmentation Method

The goal of our extension of Kuhnigk's algorithm [8] was to improve the segmentation of nodules located at concave parts of the pleura while changing the original method as little as possible in order to get consistent results. The convex hull operation is obviously not suitable for reconstructing the shape of a concave object but we observed that the error decreases when the ROI is made smaller since the convex hull is basically determined by the most distant points of the lung boundary that are contained in the ROI. Therefore our approach is to make the ROI as small as possible so that the disturbing effect of the concavity is minimized. As a minimal ROI, we choose a dilated ellipsoid that is computed as an approximation of the nodule shape. Ellipsoid approximation of lung nodules has been used with different goals and methods in the literature [6, 10].

(a) (b) (c)

(d) (e) (f)
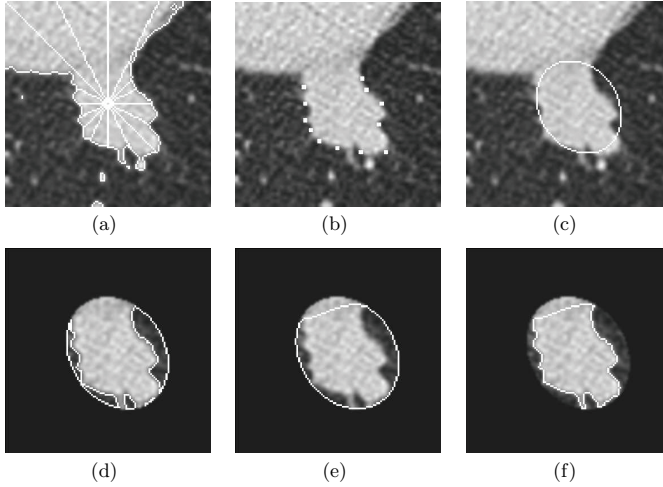
Fig. 1. Step-by-step illustration of the segmentation algorithm for juxtapleural lung nodules, exemplified by the central axial slice. (a) Result mask of the initial region growing and rays cast towards the boundary of the mask. (b) Valid ray endpoints on the nodule boundary. (c) Ellipsoid fitted to the boundary points. (d) Part of the mask within the dilated ellipsoid. (e) Convex hull of the inverse mask within the dilated ellipsoid. (f) Difference of convex hull and inverse mask.

Our method is a preprocessing step for Kuhnigk's algorithm [8] and consists of three parts which are described in the following sections and illustrated in Fig. 1:

1. identification of points on the nodule boundary by region growing and subsequent ray casting from the seed point;
2. calculation of an ellipsoid that approximates the shape of the nodule;
3. convex hull operation as in [8], but restricted to the dilated ellipsoid.

3.1 Region Growing and Ray Casting

Initially, region growing is performed, using the ROI center as a seed point. Since we only need to separate the nodule and attached high-density structures from the lung parenchyma in this first step, we can use -400 HU as a fixed threshold. In order to find points on the boundary between the nodule and the parenchyma, we apply a ray casting approach. Starting from the seed point, rays are sent out through all surface voxels of a $5 \times 5 \times 5$ cube around the seed point. This ensures a symmetric distribution of rays and an alignment to the voxel grid. These $5^3 - 3^3 = 98$ rays are traced until they reach either the boundary

of the region growing mask or leave the ROI (Fig. 1(a)). In the former case, the
endpoints are stored, otherwise discarded (Fig. 1(b)). Since some false boundary
points may be found due to noise or other structures in the outer parts of the
ROI, ray endpoints above a certain distance from the seed point should also be
discarded. We found the 95% quantile of the distances of all points to provide a
good threshold.

3.2   Ellipsoid Approximation

Typically, the points found by the ray casting procedure cover a major part
of the actual nodule surface. Assuming that the nodule has approximately an
ellipsoid shape, we aim at reconstructing this shape by fitting an ellipsoid to the
points (Fig. 1(c)).

A 3d ellipsoid is defined as a conic section

$$\{\mathbf{x} \in \mathbb{R}^3 \mid \mathbf{x}^T \mathbf{A} \mathbf{x} + \mathbf{b}^T \mathbf{x} + c = 0\}, \tag{1}$$

where the symmetric matrix $\mathbf{A} \in \mathbb{R}^{3\times3}$ is positive or negative definite. Due to
its symmetry, $\mathbf{A}$ has only six degrees of freedom, plus a total of four for $\mathbf{b} \in \mathbb{R}^3$
and $c \in \mathbb{R}$. From the valid endpoints of the 98 rays we want to determine those
ellipsoid parameters which are optimal in a least squares sense. This establishes
a non-linear equation system which can be reduced to a generalized eigenvalue
problem and solved efficiently with a method proposed by Grammalidis and
Strintzis [11]. It does not guarantee $\mathbf{A}$ to be definite, but our experiments showed
that this is almost always the case. If the points are distributed in a way such
that it is not possible to fit an ellipsoid to them – if the nodule, for example, has
an irregular shape or very extensive contact to other structures – a sphere can
be computed instead with the radius as the only free parameter. Although this
is a coarser approximation, it can still yield acceptable results in most of these
rare cases.

It should be noted that the center of the ellipsoid is included in the optimiza-
tion. The user-defined seed point influences only the distribution of the boundary
points. Since the equation system is highly overdetermined it is robust against
variations caused by different user interactions.

3.3   Convex Hull

For the following computations we use a slightly dilated version of the ellipsoid
as a new minimal ROI (Fig. 1(d)). The dilation is necessary to ensure that
the nodule is covered completely. At its margin, the ellipsoid contains some
parenchyma voxels as well, so the convex hull operation can now be applied to
reconstruct the original lung shape within this ROI (Fig. 1(e)). This is sufficient
for determining the boundary of the nodule and it works in concave parts of the
lungs as well because these concavities are no longer visible inside the ellipsoid.
For performance reasons, the convex hull has been implemented as the union of
slice-wise convex hulls in axial, sagittal, and coronal views.

Subsequently, the original algorithm [8] is executed and the adaptive opening
procedure removes attached vessels. The final result is shown in (Fig. 1(f)).
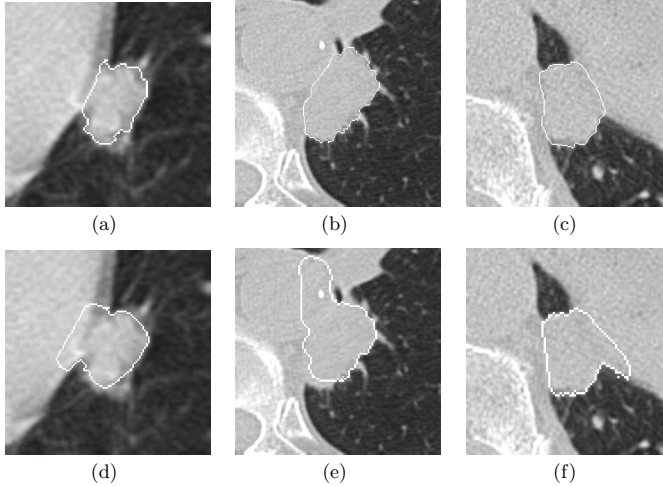
Fig. 2. Upper row: Examples of successful segmentations with the extended algorithm. Lower row: Corresponding results of the original version.

4 Results and Discussion

For our evaluation, we used a database of 333 ROIs of juxtapleural lung nodules from various patients, clinics and CT scanners with seed points set manually by radiologists. Since extensive studies have been conducted for the original version [8], we focussed on the effects of the extensions presented above. It is often impossible to determine the exact boundary between a nodule and a structure it is attached to. In cases like these where no reliable ground truth is available it is most important to produce consistent results so that volumes can be compared over time. Therefore we evaluated the segmentation results visually and examined if they were consistent with our approach to reconstruct the tumor-free shape of the lung parenchyma.

While in 71% of the cases the result of the original algorithm was classified as good, our extension could increase this proportion to 89%. For an additional 5%, a good result was obtained after applying the interactive correction procedure of [8]. Fig. 2 shows some examples of successful segmentations and reveals a significant improvement over the previous results. Most of the nodules that could not be segmented had complex shapes or a very extensive connection to high-density structures that made it impossible to fit ellipsoids to them. Examples can be seen in Fig. 3.

The experiments show that we have developed an algorithm which is able to segment most juxtapleural lung nodules, even if the surrounding parenchyma
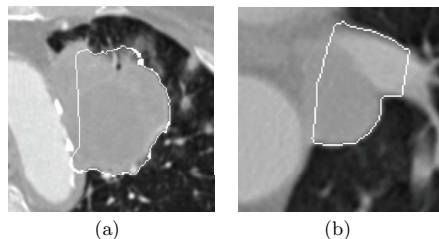
Fig. 3. Examples of difficult cases where the new algorithm did not succeed. (a) The ellipsoid does not cover the entire lesion since its 3d shape is hard to estimate. (b) The extensive vasculature is not cut off completely.

is not convex. Thus we closed a gap in Kuhnigk's algorithm [8] that had systematic problems in these cases. Our proposed method can easily be integrated there and has no effect on central nodules. No significant increase in runtime could be found. Although some additional operations are performed, subsequent computations can be restricted to the dilated ellipsoid. The computation time for the complete segmentation is around 2 s for large nodules and often below 1 s for small ones on a PC with a 2 GHz DualCore processor.

All of the nodules in this study had direct contact to the pleura and constituted a sample of very difficult cases. Considering that [8] already reported a success rate of 91% on a representative collection of central and juxtapleural lung nodules, we can conclude that with our extension the algorithm allows a successful segmentation in almost all cases.

References

1. Fetita, C.I., Prêteux, F., Beigelman-Aubry, C., Grenier, P.: 3d automated lung nodule segmentation in HRCT. In: Proc. MICCAI. (2003) 626–634
2. van Ginneken, B.: Supervised probabilistic segmentation of pulmonary nodules in CT scans. In: Proc. MICCAI. (2006) 912–919
3. Shen, H., Goebel, B., Odry, B.: A new algorithm for local surface smoothing with application to chest wall nodule segmentation in lung CT data. In: Proc. SPIE. Volume 5370. (2004) 1519–1526
4. Wiemker, R., Rogalla, P., Blaffert, T., Sifri, D., Hay, O., Shah, E., Truyen, R., Fleiter, T.: Aspects of computer-aided detection (CAD) and volumetry of pulmonary nodules using multislice CT. Br. J. Radiol. **78** (2005) S46–S56
5. Reeves, A.P., Chan, A.B., Yankelevitz, D.F., Henschke, C.I., Kressler, B., Kostis, W.J.: On measuring the change in size of pulmonary nodules. IEEE Trans. Med. Imag. **25**(4) (2006) 435–450

6. Okada, K., Ramesh, V., Krishnan, A., Singh, M., Akdemir, U.: Robust pulmonary nodule segmentation in CT: Improving performance for juxtapleural cases. In: Proc. MICCAI. (2005) 781–789
7. Dehmeshki, J., Amin, H., Valdivieso, M., Ye, X.: Segmentation of pulmonary nodules in thoracic CT scans: A region growing approach. IEEE Trans. Med. Imag. **27**(4) (2008) 467–480
8. Kuhnigk, J.M., Dicken, V., Bornemann, L., Bakai, A., Wormanns, D., Krass, S., Peitgen, H.O.: Morphological segmentation and partial volume analysis for volumetry of solid pulmonary lesions in thoracic CT scans. IEEE Trans. Med. Imag. **25**(4) (2006) 417–434
9. Moltz, J.H., Kuhnigk, J.M., Bornemann, L., Peitgen, H.O.: Segmentierung pleuraständiger Lungenrundherde in CT-Bildern mittels Ellipsoidapproximation. In: Proc. BVM. (2008) 173–177
10. Fan, L., Qian, J., Odry, B.L., Shen, H., Naidich, D., Kohl, G., Klotz, E.: Automatic segmentation of pulmonary nodules by using dynamic 3d cross-correlation for interactive CAD systems. In: Proc. SPIE. Volume 4684. (2002) 1362–1369
11. Grammalidis, N., Strintzis, M.G.: Head detection and tracking by 2-d and 3-d ellipsoid fitting. In: Proc. CGI. (2000) 221–226

A Novel Approach for Accurate Estimation of the Growth Rate of the Detected Lung Nodules

Ayman El-Baz[1], Georgy Gimel'farb[2], Robert Falk[3], David Heredia[1], and
Mohamed Abo El-Ghar[4]

[1]Bioimaging Laboratory, Bioengineering Department, University of Louisville,
Louisville, KY, USA.
[2] Department of Computer Science, University of Auckland, Auckland, New Zealand.
[3]Director, Medical Imaging Division, Jewish Hospital, Louisville, KY, USA.
[4] Urology and Nephrology Department, University of Mansoura, Mansoura, Egypt.

Abstract. Our long term research goal is to develop a fully automated,
image-based diagnostic system for early diagnosis of pulmonary nodules
that may lead to lung cancer. This paper focuses on monitoring the de-
velopment of lung nodules detected in successive chest low dose (LD) CT
scans of a patient. We propose a new methodology for 3D LDCT data
registration which is non-rigid and involves two steps: (i) global align-
ment of one scan (target) to another scan (reference or prototype) using
the learned prior appearance model followed by (ii) local alignment in
order to correct for intricate deformations. After equalizing signals for
two subsequent chest scans, visual appearance of these chest images is
modeled with a Markov-Gibbs random field with pairwise interaction.
We estimate the affine transformation that globally register the target
to the prototype by gradient descent maximization of a special Gibbs en-
ergy function. To handle local deformations, we deform each voxel of the
target over evolving closed equi-spaced surfaces (iso-surfaces) to closely
match the prototype. The evolution of the iso-surfaces is guided by an
exponential speed function in the directions that minimize distances be-
tween the corresponding voxel pairs on the iso-surfaces in both the data
sets. Preliminary results on the 135 LDCT data sets from 27 patients
show that our proper registration could lead to precise diagnosis and
identification of the development of the detected pulmonary nodules.

1 Introduction

Because lung cancer is the most common cause of cancer deaths, fast and accu-
rate analysis of pulmonary nodules is of major importance for medical computer-
aided diagnostic systems (CAD). We have already introduced the following three
successive pre-processing stages of such a system: a fully automatic segmentation
algorithm to separate lung regions from LDCT images [1], a fully automatic nod-
ule detection algorithm showing the accuracy up to 93.3% on the experimental
database containing 200 real LDCT chest data sets with 36,000 2D slices [3], and
an accurate segmentation algorithm to separate the detected pulmonary nodules
from the lung regions in the LDCT images [4]. This paper focuses on the next

stage, namely, on accurate registration of the detected nodules for subsequent volumetric measurements to monitor how the nodules are developing over the time.

(a)   (b)   (c)   (d)
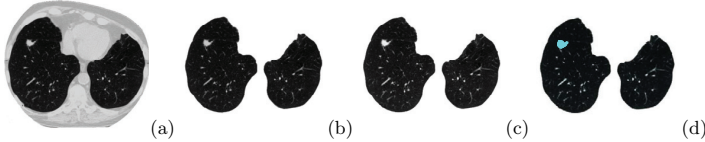
Fig. 1. Pre-processing steps: (a) an initial LDCT slice, (b) the segmented lung regions [1], (c) the normalized segmented lung regions, and (d) the segmented pulmonary nodules [4].

Figure 1 shows the results of the above-mentioned three pre-processing stages of the proposed CAD system for monitoring detected pulmonary nodules (these stages are not discussed in this paper): (*i*) an initial LDCT slice in Fig. 1(a) is segmented with the algorithms in [1] in order to isolate lung tissues from the surrounding structures in the chest cavity as shown in Fig. 1(b), (*ii*) data normalization as shown in Fig. 1(c), and (*iii*) the nodules in the isolated lung regions are segmented by evolving deformable boundaries under forces that depend on the learned current and prior appearance models as shown in Fig. 1(d) (see [4]). This paper focuses on details of the proposed global and local registration models being the core of our approach to monitoring the nodule development.

Previous work. Tracking the temporal nodule behavior is a challenging task because of changes in the patient's position at each data acquisition, as well as effects of heart beats and respiration. In order to accurately measure how the nodules are developing in time, all these motions should be compensated by registering LDCT data sets taken at different time. Many methods have been proposed for solving medical image registration problems (see e.g. [5]) and to exclude the lung motions (see [6]). Moreover, it has been reported that the computer-assisted volume measurement is more reliable for small pulmonary nodules than the measurement by human experts [7]. Therefore, the remaining principal difficulty in monitoring and evaluating the nodule growth rate is automatic identification (or registration) of corresponding nodules in the follow-up scans. Registration of the two successive CT scans determines transformation of one image with respect to the other [8]. Some examples of previous works on registration of CT lung images are overviewed below.

Most of them exploit corresponding local structural elements (features) in the images. For the follow-up of small nodules, Brown et al. [9] developed a patient-specific model with 81% success for 27 nodules. Ko et al. [10] used centroids of local structures to apply rigid and affine image registration with 96% success for 58 nodules of 10 patients. To account for non-rigid motions and deformations of the lung, Woods et al. [11] developed an objective function using an anisotropic smoothness constraint and a continuous mechanical model. Feature points required by this algorithm are detected and registered as explained

in [12], and then the continuous mechanical model is used to interpolate the image displacement. In the Wood's experiments, the difference between the estimated and actual volumes was about 1.6%. Later on, Dougherty et al. [13] developed an optical flow and model based motion estimation technique for estimating first a global parametric transformation and then local deformations of the images. This method aligned sequential CT images with a 95% correlation. Naqa et al. [14] combined the optical flow analysis with spirometric data (measurements of the airflow into and out of lungs) in order to track the breathing motion automatically. The spirometry in this study was obtained by using the reconstruction of free breathing from the 4D CT data proposed in [15].

In several studies CT lung images are matched directly for pulmonary registration. Zhang et al. [16] used a standard lung atlas to analyze the pulmonary structures in CT images. The atlas is registered to a new image by combining global rigid and local elastic transformations of a 3D surface. Li et al. [17] still used feature points to search for correspondence but exploited landmark and intensity based registration algorithms to warp a template image to the rest of the lung volumes. Okada et al. [18] proposed an anisotropic intensity model fitting with analytical parameter estimation to evaluate the nodule volume without explicit image segmentation. Zhao et al. [19] and Kostis et al. [20] proposed to segment 2D and 3D nodules by thresholding the voxel intensity followed by a connectivity filter. Their algorithms accurately segment well-defined solid nodules with similar average intensities but become unreliable on cavities or non-solid nodules. Reeves et al. [21] proposed a framework for measuring changes of the nodule size from two CT scans recorded at different times. This approach is based on using rigid registration to align the scans followed by adaptive thresholding to segment the nodules.

Nonetheless, all the existing computational methods for monitoring the pulmonary nodules detected in the CT scans do not account for large deformations of the lung tissues due to breathing and heart beating. These methods are not suitable for some types of pulmonary nodules such as cavities and ground glass nodules. Also, these methods require significant user interaction which is difficult for a clinical practitioner.

Our approach to nodule monitoring is simpler for the clinical use because all the computations to follow up the pulmonary nodules detected in the LDCT images are performed fast and in a fully automatic mode. To achieve the goal of early diagnosis of malignant nodules, our system follows up the detected pulmonary nodules using the following image analysis steps:

1. location of the nodule either by an expert (radiologist) or using our previous lung CAD system [3];
2. segmentation of the lung in the LDCT images;
3. registration (global and local alignment) of two successive LDCT scans to correct motion artefacts caused by breathing and other patient motions and establish correspondence between the nodules;
4. segmentation of the corresponding pulmonary nodules; and

5. measurement of volumetric changes between the corresponding nodules.

2 Lung Motion Correction Models

2.1 Global Alignment

Basic notation. Let $\mathcal{Q} = \{0, \ldots, Q-1\}$; $\mathbf{R} = [(x, y, z) : x = 0, \ldots, X-1; y = 0, \ldots, Y-1; z = 0, \ldots, Z-1]$, and $\mathbf{R_p} \subset \mathbf{R}$ be a finite set of scalar image signals (e.g. gray levels), a 3D arithmetic lattice supporting digital LDCT image data $g : \mathbf{R} \to \mathcal{Q}$, and an arbitrary-shaped part of the lattice occupied by the prototype, respectively. Let a finite set $\mathcal{N} = \{(\xi_1, \eta_1, \zeta_1), \ldots, (\xi_n, \eta_n, \zeta_n)\}$ of the (x, y, z)-coordinate offsets define neighboring voxels, or neighbors $\{((x + \xi, y + \eta, z + \zeta), (x - \xi, y - \eta, z - \zeta)) : (\xi, \eta, \zeta) \in \mathcal{N}\} \wedge \mathbf{R_p}$ interacting with each voxel $(x, y, z) \in \mathbf{R_p}$. The set $\mathcal{N}$ yields a 3D neighborhood graph on $\mathbf{R_p}$ that specifies translation invariant pairwise interactions between the voxels with n families $\mathcal{C}_{\xi,\eta,\zeta}$ of second-order cliques $c_{\xi,\eta,\zeta}(x, y, z) = ((x, y, z), (x + \xi, y + \eta, z + \zeta))$. Interaction strengths are given by a vector $\mathbf{V}^\mathsf{T} = \left[\mathbf{V}_{\xi,\eta,\zeta}^\mathsf{T} : (\xi, \eta, \zeta) \in \mathcal{N} \right]$ of potentials $\mathbf{V}_{\xi,\eta,\zeta}^\mathsf{T} = \left[V_{\xi,\eta,\zeta}(q, q') : (q, q') \in \mathcal{Q}^2 \right]$ depending on signal co-occurrences; here T indicates transposition.

Data normalization: To account for possible monotone (order -preserving) changes of signals (e.g. due to different sensor characteristics), every LDCT data set is equalized using the cumulative empirical probability distribution of its signals (see Fig. 1(c)).

MGRF based appearance model: In a generic MGRF with multiple pairwise interaction [1], the Gibbs probability $P(g) \propto \exp(E(g))$ of an object g aligned with the prototype g° on $\mathbf{R_p}$ is specified with the Gibbs energy $E(g) = |\mathbf{R_p}|\mathbf{V}^\mathsf{T}\mathbf{F}(g)$ where $\mathbf{F}^\mathsf{T}(g)$ is the vector of scaled empirical probability distributions of signal co-occurrences over each clique family: $\mathbf{F}^\mathsf{T}(g) = [\rho_{\xi,\eta,\zeta}\mathbf{F}_{\xi,\eta,\zeta}^\mathsf{T}(g) : (\xi, \eta, \zeta) \in \mathcal{N}]$ where $\rho_{\xi,\eta,\zeta} = \frac{|\mathcal{C}_{\xi,\eta,\zeta}|}{|\mathbf{R_p}|}$ is the relative size of the family and $\mathbf{F}_{\xi,\eta,\zeta}(g) = [f_{\xi,\eta,\zeta}(q, q'|g) : (q, q') \in \mathcal{Q}^2]^\mathsf{T}$; here, $f_{\xi,\eta,\zeta}(q, q'|g) = \frac{|\mathcal{C}_{\xi,\eta,\zeta;q,q'}(g)|}{|\mathcal{C}_{\xi,\eta,\zeta}|}$ are empirical probabilities of signal co-occurrences, and $\mathcal{C}_{\xi,\eta,\zeta;q,q'}(g) \subseteq \mathcal{C}_{\xi,\eta,\zeta}$ is a subfamily of the cliques $c_{\xi,\eta,\zeta}(x, y, z)$ supporting the co-occurrence $(g_{x,y,z} = q, g_{x+\xi,y+\eta,z+\zeta} = q')$ in g. The co-occurrence distributions and the Gibbs energy for the object are determined over $\mathbf{R_p}$, i.e. within the prototype boundary after an object is affinely aligned with the prototype. To account for the affine transformation, the initial image is resampled to the back-projected $\mathbf{R_p}$ by interpolation.

The appearance model consists of the neighborhood $\mathcal{N}$ and the potential $\mathbf{V}$ to be learned from the prototype.

Learning the potentials: The MLE of $\mathbf{V}$ is proportional in the first approximation to the scaled centered empirical co-occurrence distributions for the prototype [1]:

$$\mathbf{V}_{\xi,\eta,\zeta} = \lambda \rho_{\xi,\eta,\zeta} \left(\mathbf{F}_{\xi,\eta,\zeta}(g^\circ) - \frac{1}{Q^2}\mathbf{U} \right); \; (\xi,\eta,\zeta) \in \mathcal{N} \qquad (1)$$

where $\mathbf{U}$ is the vector with unit components. The common scaling factor λ is also computed analytically; it is approximately equal to Q^2 if $Q \gg 1$ and $\rho_{\xi,\eta,\zeta} \approx 1$ for all $(\xi,\eta,\zeta) \in \mathcal{N}$. In our case it can be set to $\lambda = 1$ because the registration uses only relative potential values and energies.

Learning the characteristic neighbors: To find the characteristic neighborhood set $\mathcal{N}$, the relative Gibbs energies $E_{\xi,\eta,\zeta}(g^\circ) = \rho_{\xi,\eta,\zeta}\mathbf{V}_{\xi,\eta,\zeta}^{\mathsf{T}}\mathbf{F}_{\xi,\eta,\zeta}(g^\circ)$ for the clique families, i.e. the scaled variances of the corresponding empirical co-occurrence distributions, are compared for a large number of possible candidates.

To automatically select the characteristic neighbors, we consider an empirical probability distribution of the energies as a mixture of a large "non-characteristic" low-energy component and a considerably smaller characteristic high-energy component: $P(E) = \pi P_{\mathrm{lo}}(E) + (1 - \pi)P_{\mathrm{hi}}(E)$. Both the components $P_{\mathrm{lo}}(E)$, $P_{\mathrm{hi}}(E)$ are of arbitrary shape and thus are approximated with linear combinations of positive and negative discrete Gaussians (efficient EM-based algorithms introduced in [1] are used for both the approximation and the estimation of π).

Appearance-based registration: The desired affine transformation of an object g corresponds to a local maximum of its relative energy $E(g_{\mathbf{a}}) = \mathbf{V}^{\mathsf{T}}\mathbf{F}(g_{\mathbf{a}})$ under the learned appearance model $[\mathcal{N}, \mathbf{V}]$. Here, $g_{\mathbf{a}}$ is the part of the object image reduced to $\mathbf{R}_{\mathbf{p}}$ by the 3D affine transformation $\mathbf{a} = [a_{11}, \ldots, a_{23}]$: $x' = a_{11}x + a_{12}y + a_{13}z + a_{14}$; $y' = a_{21}x + a_{22}y + a_{23}z + a_{24}$; $z' = a_{31}x + a_{32}y + a_{33}z + a_{34}$. The initial transformation step is a pure translation with $a_{11} = a_{22} = 1$; $a_{12} = a_{21} = 0$, ensuring the most "energetic" overlap between the object and prototype. In other words, the chosen initial position $(a_{14}^*, a_{24}^*, a_{34}^*)$ maximizes the Gibbs energy. Then the gradient search for the local energy maximum closest to the initialization selects all the 12 parameters $\mathbf{a}$.

Figures 2(c,d) show the results of the global alignment of two segmented lungs. It is clear from Fig. 2(d) that the global alignment is not perfect due to local deformation.

2.2 Local motion model

To handle local deformations, we propose to deform the object over evolving closed equi-spaced surfaces (distance iso-surfaces) so that it closely matches the prototype. The evolution is guided by an exponential speed function and intends to minimize distances between corresponding voxel pairs on the iso-surfaces in
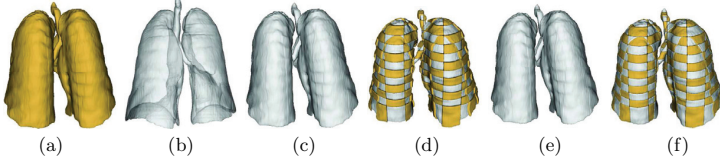
Fig. 2. 3D global and local registration: (a) reference data, (b) target data, (c) target data after 3D affine transformation, (d) checkerboard visualization to show the motion of lung tissues, (e) results of our non-rigid registration, and (f) checkerboard visualization to show the quality of the proposed local deformation model.

both the images. The normalized cross correlation of the Gibbs energy is used to find correspondences between the iso-surfaces.

Our approach involves the following steps. First, a distance map inside the object is generated using fast marching level sets [22]. Secondly, the distance map is used to generate iso-surfaces (Fig. 3). Note that the number of iso-surfaces is not necessarily the same for both the images and depends on the accuracy and the speed required by the user. The third step consists in finding correspondences between the iso-surfaces using the normalized cross correlation of the Gibbs energy. Finally, the evolution process deforms the iso-surfaces in the first data set (the target image) to match the iso-surfaces in the second data set (the prototype). The following notation is used below for defining the evolution equation:

- $\mathbf{b}_{\mathbf{g}_1}^h = [\mathbf{p}_k^h : k = 1, \ldots, K] - K$ control points on a surface h on the reference data such that $\mathbf{p}_k = (x_k, y_k, z_k)$ form a circularly connected chain of line segments $(\mathbf{p}_1, \mathbf{p}_2), \ldots, (\mathbf{p}_{K-1}, \mathbf{p}_K), (\mathbf{p}_K, \mathbf{p}_1)$;
- $\mathbf{b}_{\mathbf{g}_2}^\gamma = [\mathbf{p}_n^\gamma : n = 1, \ldots, N] - N$ control points on a surface γ on the target data such that $\mathbf{p}_n = (x_n, y_n, z_n)$ form a circularly connected chain of line segments $(\mathbf{p}_1, \mathbf{p}_2), \ldots, (\mathbf{p}_{N-1}, \mathbf{p}_N), (\mathbf{p}_N, \mathbf{p}_1)$;
- $S(\mathbf{p}_k^h, \mathbf{p}_n^\gamma)$ – the Euclidean distance between a point on the surface h in the image $\mathbf{g}_1$ and the corresponding point on the surface γ in the image $\mathbf{g}_2$;
- $S(\mathbf{p}_n^\gamma, \mathbf{p}_n^{\gamma-1})$ – the Euclidean distance between a point on the surface γ in the image $\mathbf{g}_1$ and the nearest point on the surface $\gamma - 1$ in $\mathbf{g}_1$, and
- $\nu(.)$ – the propagation speed function.

The evolution $\mathbf{b}_\tau \rightarrow \mathbf{b}_{\tau+1}$ of a deformable boundary $\mathbf{b}$ in discrete time, $\tau = 0, 1, \ldots$, is specified by the system $\mathbf{p}_{n,\tau+1}^\gamma = \mathbf{p}_{n,\tau}^\gamma + \nu(\mathbf{p}_{n,\tau}^\gamma)\mathbf{u}_{n,\tau}$; $n = 1, \ldots, N$ of difference equations where $\nu(\mathbf{p}_{n,\tau}^\gamma)$ is a propagation speed function for the control point $\mathbf{p}_{\mathbf{n},\tau}^\gamma$ and $\mathbf{u}_{n,\tau}$ is the unit vector along the ray between the two corresponding points. The propagation speed function

$$\nu(\mathbf{p}_{n,\tau}^\gamma) = \min \left\{ S(\mathbf{p}_k^h, \mathbf{p}_{n,\tau}^\gamma), S(\mathbf{p}_{n,\tau}^\gamma, \mathbf{p}_{n,\tau}^{\gamma-1}), S(\mathbf{p}_{n,\tau}^\gamma, \mathbf{p}_{n,\tau}^{\gamma+1}) \right\}$$

satisfies the condition $\nu(\mathbf{p}_{n,\tau}^\gamma) = 0$ if $S(\mathbf{p}_k^h, \mathbf{p}_{n,\tau}^\gamma) = 0$ and prevents the current point from cross-passing the closest neighbor surfaces. The latter restriction is known as the smoothness constraint.

Again, the checkerboard visualization (Fig. 2(d)) of the data set in Fig. 2(a) and the aligned data set in Fig. 2(c) highlights the effect of the motion of lung tissues. It can be seen that the connections at the lung edges between the two volumes are not smooth when using only the global registration model. This is due to the local deformation which comes from breathing and heart beats. The connections of the lung edges between the two volumes are considerably smoother when using the proposed local deformation model (see Fig. 2(f)).

Validation of the proposed local deformation model: To validate the local registration, we simulated local deformations on the real LDCT data set using the free form deformation (FFD) [23] (it simulates local displacement with the 3D cubic spline). To measure the accuracy of the proposed local registration, three different types of the deformation fields were generated with the FFD: (1) small deformation, (2) moderate deformation, and (3) large deformation as shown in Table 1. Our registration model has been applied to each type of deformation, and the accuracy of our approach has been quantitatively assessed by comparing the simulated and recovered voxel displacements (see Table 1).
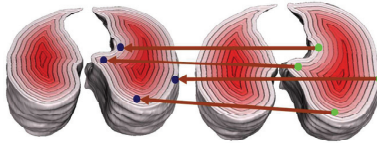
Fig. 3. Equi-spaced surfaces.

Table 1. Registration accuracy for simulated displacements (all units in mm).

	Simulated displacement		
	Type 1	Type 2	Type 3
Maximum displacement	1.7	10.8	19.9
Mean ± standard deviation	0.6 ± 0.4	2.3 ± 0.7	9.1 ± 1.1
	Alignment error		
Maximum error	0.6	1.4	2.1
Mean ± standard deviation	0.4 ± 0.3	1.0 ± 0.4	1.2 ± 1.6

3 Experimental Results

The proposed registration models were tested on the clinical datasets collected from 27 patients. Each patient has five LDCT scans, with the three months period between each two successive scans. This preliminary clinical database was collected by the LDCT scan protocol using a multidetector GE Light Speed Plus scanner with the following scanning parameters: slice thickness of 2.5 mm reconstructed every 1.5 mm, scanning pitch 1.5, pitch 1 mm, 140 KV, 100 MA, and F.O.V 36 cm.

After the two volumes at different time instants are registered, the task is to find out if the nodules are growing or not. For this purpose, the lung nodules were segmented after registration using our previous approach [4]. Once the nodules

are segmented in the original and the registered image sequences, the volumes of the nodules are calculated using the Δx, Δy, and Δz values from the scanner (in our case, 0.7, 0.7, and 2.5 mm, respectively). Figure 4 shows the estimated growth rate for the two detected pulmonary nodules (for two different patients over one year) before and after data alignment.

Fig. 4. Results of our registration for two patients over one year.
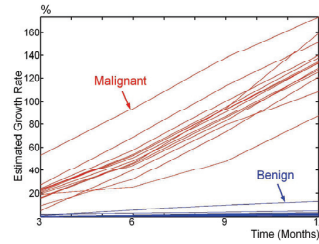
Fig. 5. Estimated volumetric changes for 14 malignant and 13 benign nodules.

It is clear that our alignment algorithm facilitates accurate evaluations of temporal changes in the nodule size. Moreover, the proposed alignment would help doctors and radiologists to track the nodule growth direction which is crucial for surgical or radiation treatment. Also, it is apparent that the malignant nodule doubles in size for 360 or less days, while the volumetric changes in the benign nodule are very small (maximum 6% over one year, see Figure 5).

Our statistical analysis using the unpaired t-test shows that the difference between the average growth rate of malignant nodules and the average growth rate of benign nodules found with the proposed approach is statistically significant (as shown in Table 2). Also, Table 2 shows that no significant difference is found if the growth rate is measured without the data alignment step. Figure 5 shows volumetric changes for 14 malignant and 13 benign nodules. It is obvious that the growth rate of the malignant nodules is considerably higher than the growth rate of the benign nodules, and this encourages to use the estimated growth rate as a discriminatory feature.

A traditional Bayes classifier based on the analysis of the growth rate of both benign and malignant nodules for 27 patients diagnosed 14 and 13 patients as malignant and benign, respectively. For simplicity, this classifier used a multivariate Gaussian model of the growth rate with the rates at 3, 6, 9, and 12 months as four discriminant features. The same patients were diagnosed by biopsy (the ground truth) showing that the classification was 100% correct. Therefore, the

Table 2. Growth rate statistics for 14 patients with malignant nodules and 13 patients with benign nodules (p – statistical significance; μ – average rate, %; σ – standard deviation, %).

Scanning period	With the proposed registration					Without the registration				
	Malignant		Benign			Malignant		Benign		
	μ_M	σ_M	μ_B	σ_B	p	μ_M	σ_M	μ_B	σ_B	p
3 months	22	16	0.9	0.7	10^{-4}	5.6	4.8	2.8	1.9	0.1
6 months	49	20	2.9	2.3	10^{-4}	11	6.6	8.4	5.1	0.3
9 months	91	29	4.5	3.8	10^{-4}	24	9.3	17	11	0.1
12 months	140	32	5.4	4.3	10^{-4}	30	11	20	16	0.1

proposed image analysis techniques could be a promising supplement to the current technologies for diagnosing lung cancer.

4   Conclusions

We introduced a new approach for registering 3D spiral LDCT images that combines an initial affine global alignment of one scan (the target) to another scan (the reference) using the learned prior appearance model and subsequent local alignments that account for more intricate deformations. Preliminary results on 27 patients show the registration could lead to accurate diagnosis and identification of temporal development of detected pulmonary nodules. Our present C++ implementation on the Intel dual processor (3GHz each) with 8 GB memory and 1.5 TB hard drive with RAID technology takes about 330 sec for processing 182 LDCT slices of size 512×512 pixels each, i.e about 1.8 sec per slice. Our future work will focus on testing the proposed approach on more diverse data sets. We have already started to collect the data from additional 200 patients with different types of pulmonary nodules (e.g., ground glass, cavity, etc), in order to better measure the accuracy and limitations of the proposed framework.

References

1. A. Farag, A. El-Baz, and G. Gimel'farb, "Precise Segmentation of Multi-Modal Images," *IEEE Transactions on Image Processing*, vol. 15, no. 4, pp. 952–968, 2006.
2. A. El-Baz, "Novel Stochastic Models for Medical Image Analysis," *Ph.D. dissertation, University of Louisville*, Louisville, KY, Chapter 5, pp. 115–150, 2006.
3. A. Farag, A. El-Baz, and G. Gimelfarb, "Quantitative Nodule Detection in Low Dose Chest CT Scans: New Template Modeling and Evaluation for CAD System Design," *Proc. Int. Conf. on Medical Image Computing and Computer-Assisted Intervention (MICCAI'05)*, Palm Springs, California, USA, October 26–29, 2005, pp. 720–728, 2005.

4. A. El-Baz, et al., "A Framework for Automatic Segmentation of Lung Nodules from Low Dose Chest CT Scans", *Proc. IAPR Int. Conf. on Pattern Recognition (ICPR'06)*, Hong Kong, August 20–24, 2006, vol. 3, pp. 611–614, 2006.

5. J. Maintz and M. Viergever, "A Survey of Medical Image Registration," *Journal of Medical Image Analysis*, vol. 2, pp. 1–36, 1998.

6. J. Ko and D. Naidich, "Computer-Aided Diagnosis and the Evaluation of Lung Disease," *Journal of Thoracic Imaging*, vol. 19, no. 3, pp. 136–155, 2004.

7. W. Kostis, et al., "Small Pulmonary Nodules: Reproducibility of Three-Dimensional Volumetric Measurement and Estimation of Time to Follow-Up CT," *Radiology*, vol. 231, no. 2, pp. 446–452, May 2004.

8. B. Horn, "Closed-Form Solution of Absolute Orientation using Unit Quaternions," *Journal of the Optical Society of America B*, vol. 4, no. 4, pp. 629–642, 1987.

9. M. Brown, et al., "Method for Segmenting Chest CT Image Data using an Anatomical Model: Preliminary Results," *IEEE TMI*, vol. 16, no. 6, pp. 828–839, 1997.

10. J. Ko, and M. Betke, "Chest CT: Automated Nodule Detection and Assessment of Change over Time-Preliminary Experience," *Radiology*, vol. 218, pp. 267–273, 2001.

11. K. Woods, et al., "Model Supported Image Registration and Warping for Change Detection in Computer-Aided Diagnosis," *Applied Imagery Pattern Recognition (AIPR) Annual Workshops*, Washington DC, 2000.

12. L. Fan and C. Chen, "An Integrated Approach to 3D Warping and Registration from Lung Images," *Proceedings of SPIE Conf. Developments in X-Ray Tomography II*, July 1999.

13. L. Dougherty, et al., "Alignment of CT Lung Volumes with an Optical Flow Method," *Academic Radiology*, vol. 10, no. 3, pp. 249–254, 2003.

14. I. Naqa, et al., "Automated Breathing Motion Tracking for 4D Computed Tomography," *Nuclear Science Symp. Conference Record*, vol. 5, pp. 3219–3222, 2003.

15. D. Low, et al., "A Method for the Reconstruction of Four-Dimensional Synchronized CT Scans Acquired During Free Breathing," *Medical Physics*, vol. 30, pp. 1254–1263, 2003.

16. L. Zhang and J. Reinhardt, "3D Pulmonary CT Image Registration with a Standard Lung Atlas," *Proc. SPIE Conf. Medical Imaging*, vol. 4322, pp. 67–77, 2000.

17. B. Li, et al., "3-D Inter-Subject Warping and Registration of Pulmonary CT Images for a Human Lung Model," *Proc. SPIE Conf. on Medical Imaging*, vol. 4683, pp. 324–335, San Diego, CA, 2002.

18. K. Okada, et al., "Robust Anisotropic Gaussian Fitting for Volumetric Characterization of Pulmonary Nodules in Multislice CT," *IEEE TMI*, vol. 24, no. 3, pp. 409–423, 2005.

19. B. Zhao, et al., "Two-Dimensional Multicriterion Segmentation of Pulmonary Nodules on Helical CT Images," *Medical Physics*, vol. 26, pp. 889-895, 1999.

20. W. Kostis, et al., "Three-Dimensional Segmentation and Growth-Rate Estimation of Small Pulmonary Nodules in Helical CT Images," *IEEE TMI*, vol. 22, pp. 1259–1274, 2003.

21. A. Reeves, et al., "On Measuring the Change in Size of Pulmonary Nodules," *IEEE TMI*, vol. 25, no. 4, pp. 435–450, 2006.

22. J. Sethian, "Fast Marching Level Set Method for Monotonically Advancing Fronts," *Proc. National Academy of Sciences, USA*, vol. 93, pp. 1591–1595, Feb. 1996.

23. D. Rueckert, et al., "Nonrigid Registration using Freeform Deformations: Application to Breast MR Images," *IEEE TMI*, vol. 18, no. 8, pp. 712–721, 1999.

Computerized generation of realistic pulmonary nodule phantoms in helical CT images

Xiangwei Zhang[1], Eric Olcott[2], Philippe Raffy[1], Naichang Yu[1], Haili Chui[1]

[1] Hologic/R2, Santa Clara, CA, 95054 USA
[2] Stanford School of Medicine, Palo Alto, CA, 94087 USA

Abstract. We develop a framework of computerized generating solid pulmonary nodules in real chest helical computed tomography (CT) images. Spheres with random deformations are used to model nodule shapes. The nodule density is represented by a uniform signal with additive zero mean Gaussian noise. The insertion of synthetic nodules into real CT images is formulated as α blending between foreground nodules and background pulmonary tissues. The blending factors reflect not only the partial volume effect but also the smoothing effect in the filtered back projection (FBP) CT reconstruction. A new lesion insertion scheme based on dual source blending is proposed to blend the image noise and the lesion object separately for a better noise model. A subjective evaluation is performed by a human expert; and statistics of simulated nodules and real nodules are compared to give a quantitative analysis. These validations demonstrate a high level of similarity between the synthetic nodules and real nodules. An evaluation study of a commercial Computer-aided detection (CAD) system on an objective database created using this framework is also presented.

1   Introduction

Computer-aided detection (CAD) is a promising tool to assist in lung nodule detection, and to assess lesion size change over time based on computed tomography (CT) scans. The use of CAD may improve the performance of radiologists in helical CT lung screening. Many techniques for automated nodule detection and characterization have been developed [1], [2], [3], [4], [5], [6], [7].

Evaluating these methods is difficult due to the lack of database with large number of nodules/gold standards. Building real nodule database with expert opinions as gold standard [8] suffers from several drawbacks, especially inaccurate volume definition. Physical nodule phantoms can give accurate volume definition, but it is difficult to create large amount of nodules with different characteristics.

Compared to building real nodule CT database and making physical phantoms, computerized generation of synthetic lesions with known characteristics offers a powerful tool for CAD evaluation. A computerized nodule generation method was reported by Raffy et al, where the nodules were modelled as ellipsoids, and the insertion was a direct replacement followed by a smoothing [9]. In our previous work [10], we simulated nodules using deformed spheres, and
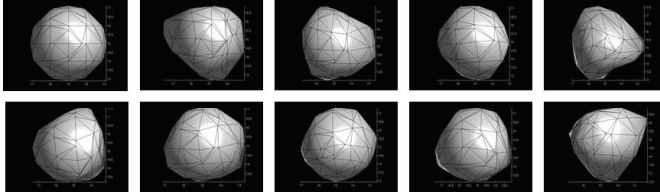
Fig. 1. various nodule shapes created using TPS deformations.

inserted the synthetic nodules into real CT images using 'weighted averaging' between nodules and the background images. Multiple sclerosis lesion phantoms in the magnetic resonance (MR) images were created by Rexilius et al: three different shapes were created to model lesions, each lesion was inserted into MR images using a 'linear combination' of the lesion and the MR scans [11].

In this work, we develop a framework for simulating solid nodules in helical CT images. The simulation method is detailed in Sect. 2, the model validation is described in Sect. 3, an application for CAD evaluation is presented in Sect. 4, and the results are given in Sect. 5.

2 Materials and Methods

A typical pulmonary nodule is about $2mm$ to $30mm$, takes a sphere-like shape. The density of the major core part is similar to muscle and vessel; near the nodule border, the density attenuates gradually until merging into the background. We concentrate on modeling the following characteristics of pulmonary nodules: 1) shape; 2) size; 3) core density; 4) density attenuation on the border.

2.1 Nodule shape/size modeling

Nodule shapes are modeled using unit spheres with randomly generated high dimensional deformations. We adopt the thin-plate spline (TPS) [12] as the non-rigid mapping to do the deformations. TPS deformations can be expressed as:

$$T^{'} = T \cdot d + k \cdot c \qquad (1)$$

Here the T is $n \times 4$ matrix in which each row is homogeneous coordinate representation of an original control point. d is 4×4 affine transform matrix. k is a $n \times n$ matrix depending on control points. c is $n \times 4$ matrix in which the first column elements are all zeros, and the remaining three columns consist of randomly distributed deformation coefficients.

Experiments showed that this model is able to generate highly realistic nodule shapes, by tuning the random distribution parameters. In this work, normal
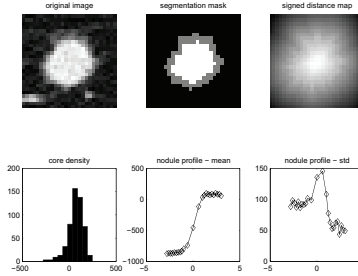
Fig. 2. Nodule density analysis. First row: original image; segmented nodule; signed distance. Second row: histogram of core density; density mean of different layer (indicated by distance values); density variance of different layer.

distribution is used to create deformation coefficients in c. Some examples of synthetic nodule shapes created are illustrated in Fig. 1.

The deformed unit sphere can be easily scaled, rotated, skewed by applying further affine transform by introducing d.

2.2 Nodule density modeling

Density distribution is analyzed for selected real solid nodules with varying sizes, shapes, CT protocols. Only isolated nodules are used, as the non-isolated nodules are difficult to segment, thus, affect the estimation. 100 real nodules from 18 CT exams (0.6-1.3mm collimations, 20-120mAs exposure) are selected.

It is important to distinguish the nodule core part and the volume average layer (due to partial volume effect, reconstruction smoothing effect). Each nodule is analyzed layer by layer from background to the center using a 3D distance transform [13], with the border as feature points, see Fig. 2. The exterior part is set to negative. For each pixel, the distance value indicates the layer – how far it is from the nodule border, and in which direction – toward or away. Note that the variance estimated for the background and volume average layer in Fig. 2 is not valid, due to the structure noise in the background, and the fast changing density in the transition layer.

It has been showed that the density of each individual core follows an approximate Gaussian distribution, but the means and variances vary across nodules and cases. For filtered back projection (FBP), the noise variance in CT images is a slowly varying spatial function [14]. In this work, stochastic noise for each simulated nodule is represented by additive Gaussian noise, and the noise variance is estimated from a neighboring structure in real CT images with water-like attenuation, either nearby vessels or chest walls. We represent the density of nodule core as $I_o = \bar{I}_o + N_o(0, \sigma_o)$, with I_o being the nodule density, $\bar{I}_o$ being

object density, N_o being the zero mean Gaussian noise. For volume average layers, there is a gradual density reduction from interior layers to exterior layers. The thickness of volume average layers is related to reconstruction kernels, and is also estimated. The modeling of volume average layer is described later.

2.3 Nodule insertion using dual source α blending

In this work, inserting synthetic lesions into real images is newly formulated as α blending, a common technique in computer graphics [15]. Given source and destination images, we can control blending on a pixel by pixel basis. Specifically, each pixel of the synthetic nodule is the source with blending factor α, and pixel of the CT scan at corresponding location is the destination with a blending factor $1 - \alpha$. Note that the blending factor is a function of pixel positions.

The first step is the rasterization of the continuous nodule shape in digital space with higher resolution (for better accuracy) than the original scan. The cuboid region of interest (ROI) need to include both the core and the volume averaging layer. This step leads to a binary representation of the synthetic shape with 1 for interior, and 0 for exterior. The volume of the nodule is the total number of the interior pixels multiplied by the volume of each pixel.

The second step is to calculate the α. We simulate the partial volume effect using α channel antialiasing. α value for a pixel is set to be a number between 0 and 1 that is the percentage of that pixel covered by the nodule. A pixel in ROI with the original resolution corresponds to a larger cuboid in the high resolution binary image, and the α value should be the percentage of that cuboid covered by the nodule – this is implemented by rasterization of the cuboid in the high resolution space, trilinear interpolation in the binary image, and averaging over the cuboid. This results in volume average layer with a single pixel thickness. The volume average layer can be more than a single pixel layer, depending on the reconstruction kernel used in FBP. To simulate this smoothing effect, we add a Gaussian smoothing on the binary image, so that interpolation occurs on a smoothed gray scale image. The kernel size is chosen roughly as the thickness of the transition layer.

The third step is to do the blending. Given the α value for each pixel of the ROI in the original CT image resolution, the blending can be described as

$$I = \alpha I_o + (1 - \alpha)I_b \qquad (2)$$

where I is the final density, I_o the nodule density, I_b the background density. Similar technique were reported [10], [11], but not formulated as α blending. This single source α blending gives us a better solution than direct smoothing [9], but the volume average layer is still visually artificial – obviously less noisy than both the nodule core part, and the background. The explanation comes from the following analysis.

Similar to the representation of nodule density as $I_o = \bar{I}_o + N_o(0, \sigma_o)$; we can also describe the background as $I_b = \bar{I}_b + N_b(0, \sigma_b)$, with I_b the background density, $\bar{I}_b$ the background density without noise, N_b the zero mean Gaussian

noise. (Note that this formula is just for convenience of analysis, it does not mean we can explicitly separate the true signal and the noise for the background. The reason we can model a solid nodule explicitly using a constant density with additive Gaussian noise is that we can reasonably assume that a solid nodule has a similar density at the major core; this assumption obviously does not hold for complicated lung field.) Then formula (2) becomes

$$I = \bar{I} + n; \quad with \quad \bar{I} = \alpha \bar{I}_o + (1 - \alpha)\bar{I}_b, \quad n = \alpha N_o + (1 - \alpha)N_b \qquad (3)$$

with $\bar{I}$ being the blending results without noise, and n is the combined noise after the blending. The mean and the variance of the combined noise are

$$E(n) = 0; \quad \sigma_n^2 = \alpha^2 \sigma_o^2 + (1 - \alpha)^2 \sigma_b^2 \qquad (4)$$

By assuming $\sigma_o = \sigma_b = \sigma$ (due to the slow changing of noise variance spatially), the above formula becomes

$$\sigma_n^2 = [\alpha^2 + (1 - \alpha)^2]\sigma^2 \qquad (5)$$

With $\alpha \in [0, 1]$, we have $\sigma_n^2 \leq \sigma^2$. This means that the noise level will be reduced using the above single source blending method. The worst case is that $\sigma_n^2 = \sigma^2/2$ when $\alpha = 0.5$.

To compensate for this undesired effect, we propose a new way of inserting synthetic nodule by using separate blending for object and noise, so called 'dual source blending' in computer graphics [15]. In this scheme, the nodule object and the noise part are treated as two separate source images, and have different but dependent blending factors. This new dual source blending for lesion insertion can be described as

$$I = \alpha \bar{I}_o + \acute{\alpha} N_o + (1 - \alpha)I_b \qquad (6)$$

where α the original source blending factor, $\acute{\alpha}$ the new source blending factor introduced specific for noise part. This formula can be extended as

$$I = \bar{I} + n; \quad with \quad \bar{I} = \alpha \bar{I}_o + (1 - \alpha)\bar{I}_b, \quad n = \acute{\alpha} N_o + (1 - \alpha)N_b \qquad (7)$$

with $\bar{I}$ part being same as before, but the combined noise part being changed, compared to formula (3). By assuming $\sigma_o = \sigma_b = \sigma$, the mean and variance of the blended noise are

$$E(n) = 0; \quad \sigma_n^2 = [\acute{\alpha}^2 + (1 - \alpha)^2]\sigma^2 \qquad (8)$$

To attain same noise levels across core part, volume average layer, and background, we need $\sigma_n^2 = \sigma^2$, i.e., $\acute{\alpha}^2 + (1 - \alpha)^2 = 1$. We can set

$$\acute{\alpha} = \sqrt{1 - (1 - \alpha)^2} \qquad (9)$$

By using dual source blending with blending factors satisfying formula (9), we can maintain the noise level before and after the insertion, under the assumption $\sigma_o = \sigma_b = \sigma$.
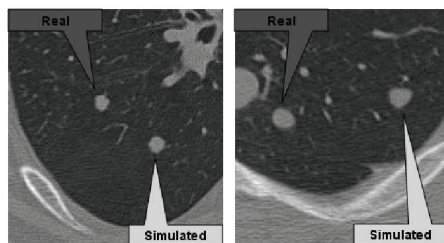
Fig. 3. Illustration of nodule simulation result.

This object insertion method can be easily applied to nodule 'cut and paste' procedure in which a real nodule is 'cut' from its original position, and 'paste' into desired positions. This 'cut' step is actually a segmentation processing; and the 'paste' step is an object insertion procedure. The proposed dual source α blending technique is very suitable for 'pasting' nodules. This 'cut and paste' technique can create 'new' nodules, but is relatively restricted compared to the synthetic model presented in this work.

3 Model Validation

The synthetic nodules simulated from this model were evaluated qualitatively and quantitatively with three studies. First, an expert radiologist carefully reviewed a randomly mixed set of 200 nodules (100 real + 100 synthetic) and provided, for each nodule, a 1 – 10 scale visual rating (1 = not real, 10 = definitely real). Second, the radiologist then rated another random mixed set of 20 nodule cases (10 real + 10 synthetic, with 2 – 10 nodules in each case) on a case by case basis. The ratings were compared between the simulated nodule group and real nodule group. Third, the correlation coefficients of the density attenuation profiles for 20 real nodules and 20 similar sized synthetic nodules were calculated to provide a quantitative measurement of the similarities.

4 Applications for CAD evaluation

A commercial CAD system for automatic lung nodule detection and volume measurement was tested using computer-simulated nodules of various sizes and different contexts (isolated and juxta-vascular) inserted into helical CT exams with different dosages. Two normal CT chest cases (1 low dose = $20mAs$ / 1 regular dose = $120mAs$; both $1.25mm$ collimation) were used as bases for simulation. For each case, two groups (isolated and juxta-vascular) of nodules

with a certain diameter and random shape were inserted. Each group included 50 nodules. Different nodule diameters $(4, 5, 6, 8, 10mm)$ were used. The CAD detection and volume measurement results were scored against the truth.

5   Results

In the model validation studies, 89/100 real nodules and 86/100 synthetic nodules were rated highly realistic (ratings ≥ 7) in the first study, as were 9/10 real cases and 8/10 synthetic cases in the second study, showing no significant statistical difference (Fisher exact test, $p \geq 0.674$). The correlation coefficients of attenuation profiles between simulated and real nodules showed a mean of 0.95 and standard deviation of 0.03.

With regard to CAD evaluation, the CAD detection rates were all 100% except one sub-group ($4mm$ low dose, 98%). The volume measurement errors were similar for different size groups ($6 - 10mm$: mean $< 2.0\%$, std. $< 2.3\%$; $4 - 5mm$: mean $\leq 3\%$, std. $< 5\%$). For simulated juxta-vascular nodules, the detection rate gradually deteriorated as nodule size decreased ($6 - 10mm$: $96+\%$; $5mm$: $86 + \%$; and $4mm$: $73 + \%$), as did the volume measurement errors ($6 - 10mm$: mean $< 4.2\%$, std. $< 10\%$; $4 - 5mm$: mean $< 6.8\%$, std. $< 14\%$). Overall, the CAD performance (detection or volume measurement) was not significantly affected by the different dosages.

6   Discussion

In the literature, the shape of nodules (lesions) was usually represented by simple shapes; the density distribution and attenuation profiles have not been thoroughly investigated; and how to merge the simulated nodule with the background has not been properly solved, especially the noise synthesis.

In this work, by applying randomly generated high dimensional deformations on spheres, we can create very realistic nodule shapes; dissecting nodules layer by layer using distance transform to investigate density distribution and attenuation profiles forms a solid base for simulating the density and the volume average layer; the newly formulated α blending framework for lesion insertion gives a well-understood description of the merging between nodule and background for each pixel. Additionally, the newly designed dual source α blending technique for lesion insertion is able to maintain the noise level across the lesion, the volume average layer, and the background, thus, makes the inserted nodule look more realistic. The Validation demonstrates a high level of similarity between the synthetic nodules and real nodules.

Using the proposed technique, it is possible to conduct a flexible performance evaluation of a CAD system on nodules with different sizes, contexts, shapes, densities, in CT images with different dosages.

The techniques presented in this work, including shape simulation using randomly deformed sphere, object density analysis using distance transform, object

insertion based on dual source α blending, can be easily extended for simulating other lesions, such as colon polyps, mass or calcs in digital mammography/tomosynthesis.

Although the deformed nodule shapes presented in this paper are very realistic, they only represent a small fraction of the possible shapes of real nodules, for example, modeling of highly spiculated nodules, partly solid nodules are not discussed. The main target of this work is to simulate nodules that are most suitable for early detection and accurate estimation of growth – small, somewhat round, solid nodules probably occur most often for early stage pulmonary nodules.

7   Conclusions

We proposed a new simulation model to insert synthetic lung nodules, with shapes and density statistics similar to real nodules, into normal CT chest exams. Nodule shapes were modeled using spheres with added random non-linear deformations. Nodule density and attenuation profiles were analyzed on real nodule samples. The volume average layers were simulated using a dual source α blending between synthesized nodule and real CT background. The synthetic nodules simulated from this model were evaluated qualitatively and quantitatively. These Validation studies demonstrated a high level of similarity between the synthetic nodules and real nodules.

References

1. Giger, M.L., Bae, K.T., MacMahon, H.: Computerized detection of pulmonary nodules in computed tomography images. Investigate. Radiol. **29** (1994) 459–465
2. Armato, S.G., Giger, M.L., Moran, C.J., Blackburn, J.T., Doi, K., MacMahon, H.: Computerized detection of pulmonary nodules on CT scans. Radiographics **19** (1999) 1303–1311
3. Kostis, W.J., Reeves, A.P., Yankelevitz, D.F., Henschke, C.I.: Three-dimensional segmentation and growth rate estimation of small pulmonary nodules in helical CT images. IEEE Transactions on Medical Imaging **22** (2003) 1259–1273
4. Paik, D.S., Beaulieu, C.F., Rubin, G.D., Acar, B., Jeffrey, R.B., Jr., Yee, J., Dey, J., Napel, S.: Surface normal overlap: a computer-aided detection algorithm with application to colonic polyps and lung nodules in helical CT. IEEE Transactions on Medical Imaging **23** (2004) 661–675
5. Zhang, X., McLennan, G., Hoffman, E.A., Sonka, M.: Automated detection of small-size pulmonary nodules based on helical CT images. In: LNCS 3565: Proc. Information Processing in Medical Imaging (IPMI) 2005. (2005) 664–676
6. Mendoca, P.R.S., Bhotika, R., Sirohey, S.A., Turner, W.D., Miller, J.V., Avila, R.S.: Model-based analysis of local shape for lesion detection in CT scans. In: LNCS 3749: Proc. Medical Image Computing and Computer-Assisted Intervention (MICCAI) 2005. (2005) 688–695
7. Zhang, X., Stockel, J., Wolf, M., Cathier, P.: A new method for spherical object detection and its application to computer aided detection of pulmonary nodules in CT images. In: LNCS 4791: Proc. Medical Image Computing and Computer-Assisted Intervention (MICCAI) 2007. (2007) 842–849

8. Armato, S.G., McLennan, G., McNitt-Gray, M.F., Meyer, C.R., Yankelevitz, D., Aberle, D.R., Henschke, C.I., Hoffman, E.A., Kazerooni, E.A., MacMahon, H., Reeves, A.P., Croft, B.Y., Clarke, L.P.: Lung image database consortium: Developing a resource for the medical imaging research community. Radiology **232** (2004) 739–748

9. Raffy, P., Fetita, C., Reigelman, C., Preteux, F., Wood, S., Grenier, P.: Evaluation of computer-aided detection (CAD) performance using mathematically simulated lung nodules. In: Proc. Computer Assisted Radiology and Surgery 2004. (2004) 935–940

10. Zhang, X., Olcott, E., Raffy, P., Yu, N., Chui, H.: Simulating solid lung nodules in MDCT images for CAD evaluation: modeling, validation, and applications. In: Proc. SPIE Medical Imaging 2007. Volume 6514. (2007) 0Z01–0Z08

11. Rexilius, J., Hahn, H.K., Schlter, M., Bourquain, H., Peitgen, H.O.: Evaluation of accuracy in MS lesion volumetry using realistic lesion phantoms. Academic Radiology **12** (2005) 17–24

12. Bookstein, F.L.: Principal warps: thin-plate splines and the decomposition of deformations. IEEE Transactions on Pattern Anal. and Mach. Intell. **11** (1989) 567–585

13. Danielsson, P.E.: Euclidean distance mapping. Computer Graphics and Image Processing **14** (1980) 227–248

14. Razifar, P., Sandstrm, M., Schnieder, H., Lngstrm, B., Maripuu, E., Bengtsson, E., Bergstrm, M.: Noise correlation in PET, CT, SPECT and PET/CT data evaluated using autocorrelation function: a phantom study on data, reconstructed using FBP and OSEM. BMC Med Imaging **5** (2005) 5

15. Angel, E.: Interactive computer graphics: a top-down approach using OpenGL. Addison-Wesley (2003)

Respiratory Motion Modeling and Estimation

Tobias Klinder[1,2], Cristian Lorenz[2], and
Jörn Ostermann[1]

[1] Institut für Informationsverarbeitung, Leibniz University of Hannover, Germany,
klinder@tnt.uni-hannover.de,
[2] Philips Research Europe - Hamburg, Sector Medical Imaging Systems, Germany [*]

Abstract. In order to cope with the problems caused by breathing motion, it would be beneficial for many applications to incorporate prior knowledge of respiratory motion. In this paper, we present the extraction, modeling, and prediction of respiratory motion based on inhale-exhale pairs of CT images. Intra- and inter-patient motion models of the lungs are built and adapted to unseen data by the use of sparse motion indicators. The created models are thereby represented as a linear model by applying Principal Component Analysis (PCA) on the covariance of motion vectors of corresponding landmarks. For model adaptation, diaphragm and rib-cage are investigated as model stimulators and compared to a systematical selection of landmarks holding most of the model's variability. While the diaphragm motion correlates well with the breathing motion achieving an average estimation error of 3.0 mm for the intra- and 4.1 mm for the inter-patient models in average, the predictability of breathing based on the rib-cage motion is significantly worse. Using both diaphragm and rib-cage as stimulators, we obtained an average estimation error of 2.8 mm for the intra- and 3.7 mm for the inter-patient models improving prediction.

1 Introduction

Respiratory motion is a key issue in radiation therapy, tumor ablation, and other treatments of the thorax and upper abdomen [1] but also for data acquisition. Since breathing motion causes a significant organ movement and deformation, an accurate knowledge of the localisation of the object in focus is difficult to obtain. However, a precise prediction of the structures of interest would be highly desirable for many applications, e.g., for dose reduction of healthy tissue during radiotherapy treatment. One approach to reduce the uncertainties caused by breathing is to use prior knowledge of the respiratory motion as, e.g., in the form of breathing models. In contrast to biomechanical models [2, 3] that aim a physically-based modeling of the lung, our goal is to build a general breathing model from an ensemble of motion fields extracted, e.g., from 4D-CT or 4D-MR, and individualize this general model using sparse motion indicators to obtain a dense motion field for the organ under investigation.

Little work has been presented for building motion models based on extracted motion fields from 4D image data while not considering tissue properties. Admittedly, in many cases the generated models were patient-specific, e.g., [4, 5], which means that in a clinical scenario, motion information can only be incorporated if images of the patient are already available or additionally acquired. In contrast to patient-specific approaches, Sundaram et al. [6] created a dynamic model of average lung deformation also registering between subjects. However, the method did not address the clinically relevant case of adapting a learned general model to a certain patient. Motion model adaptation of an inter-subject model by the use of sparse motion information has been recently presented for liver deformation to predict the drift of the exhalation position of corresponding points inside the liver [7].

In this paper, we address the problem of extracting, modeling and estimating breathing motion based on inhale-exhale pairs of CT images. We do not only deal with intra-patient model building as well as its adaptation to estimate breathing motion on different days throughout treatment, but also build inter-subject models to predict patient-specific motion learned from a general model.

Section 2.1 introduces the available image data. Based on the motion field extraction introduced in Sect. 2.2, motion model building is explained in Sect. 2.3. The adaptation of our motion models to unseen data based on sparse motion indicators is presented in Sect. 2.4. Systematic selection of regions providing sparse motion information is explained in Sect. 2.5. Finally, performance of our motion models is presented in Sect. 3.

2   Methods

2.1   Image Data

Inhale and exhale thorax CT images of seven patients all suffering from lung cancer were available over several weeks of treatment during radiotherapy. For each patient, image data of up to seven weeks plus two weeks of preliminary examination was acquired. Thus, this image data captures the variability in breathing motion throughout different weeks of treatment. Due to the fact that all images where acquired at breath hold, motion artifacts caused by breathing were not present. All images had an in-plane resolution of 0.85-0.97 mm and a slice thickness of 2.5 mm. Since the images were acquired for radiotherapy planning where the focus was set on the trajectory of the tumor, not all cases show the lungs entirely. These cases were removed from further investigation.

2.2   Motion Field Extraction

Lung motion fields are derived from inhale-exhale images using a surface-based tracking technique where the surface is represented as a triangulated mesh. By propagating a topologically identical patient-specific lung surface mesh from inhale to exhale, anatomical point correspondences are assumed to be preserved.

The trajectories of corresponding points of the adapted meshes thus provide a sparse motion field. A continuous description of the sparse motion field is finally obtained by interpolation using thin-plate-splines [8]. For the extraction of motion fields from 4D-CT data an abundant amount of other methods exist. In [9], we have recently compared surface-based tracking to three other common techniques. Mesh propagation extracts motion fields that provide an accuracy to manually set landmarks of approximately the voxel size. In a qualitative analysis, the motion fields show plausible characteristics also similar to the other methods. Most prominent advantages of surface-tracking are its ability to cope with discontinuities in motion fields, its low computational cost, and the fact that it directly provides a segmentation in all phases.

Patient-specific lung surface models for motion field extraction are obtained by selecting one chosen reference inhale image per patient and apply a triangulation of the thresholded image. These meshes cover the outer border of the lung lobes and also the surfaces of the bronchial and pulmonary vessel tree as well as the tumor surfaces [8]. By adapting the generated patient-specific reference lung mesh to inhale and exhale images over all weeks of treatment, vertex correspondences are preserved for all images of the same patient.

In addition to the lung motion field, the movement of the rib-cage is extracted. For that purpose, a general rib-cage surface model [10] is automatically positioned in all inhale images, adapted, and finally propagated to the corresponding exhale image. Since all ribs and vertebrae are labeled separately, motion fields can be easily assigned to each individual structure.

For both lung and rib-cage, mesh adaptation is performed using an iterative shape-constrained deformable surface model method [11]. In each iteration, the concordance of model and object boundary in the image is optimized. An appropriate parameter setting for lung surface mesh propagation can be found in [8]. Figure 1 shows the adapted surface meshes of lung and rib-cage in one data set.

2.3 Motion Modeling

The proposed motion models capture the variability in motion of a given learning set by applying principal component analysis (PCA) on the covariance matrix of the motion fields. Since the key issue for model building is to establish corresponding landmarks, this aspect will be explained separately for intra- and inter-patient motion models in detail below.

Suppose we have a set of M **motion fields** denoted as $\mathbf{v}_1, \ldots, \mathbf{v}_M$. Each $\mathbf{v}_i$ contains the components of N **motion vectors** $\mathbf{a}_j$ defined at N corresponding landmark positions $\mathbf{v}_i = [a_{1i_x}, a_{1i_y}, a_{1i_z}, \ldots, a_{Ni_x}, a_{Ni_y}, a_{Ni_z}]^T$. After aligning our training sets, averaging all vectors, and applying PCA on the covariance matrix, we obtain a mean motion field $\bar{\mathbf{v}}$ and its principal modes of variation ϕ_k covered in the matrix $\mathbf{\Phi}$. Thus, we can express a given motion field $\mathbf{v}_i$ as

$$\mathbf{v}_i = \bar{\mathbf{v}} + \mathbf{\Phi}\mathbf{b}_i \qquad (1)$$

where $\mathbf{b}_i$ is the weight vector.
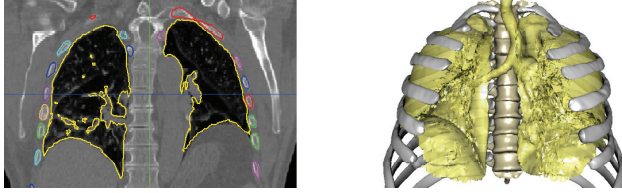
Fig. 1. Extracted lung surface mesh and adapted rib-cage shown in one slice and as surface rendering. Coloring in image slice indicates labeling of individual structures.

Intra Patient Since a topologically identical mesh is adapted to all images of the same patient as explained in Sect. 2.2, corresponding vertices are assumed to provide the correspondences of our intra-patient motion model.

Inter Patient Motion field extraction was based on patient specific surface meshes. Thus, vertex correspondences between meshes of different patients were not given. In order to establish inter-subject correspondences, a model of the outer surfaces of the lung [12] is additionally adapted to all data sets. By adapting a topologically identical surface model, we again assume anatomical point correspondences to be preserved. However, in order to not only establish correspondences on the outer surface of the lung but also in the inside, we moreover define a cartesian grid inside the lung mesh. A patient-specific grid inside the lung is obtained by applying a thin-plate spline deformation field calculated from corresponding points of individualized surface and lung model to the grid points. Figure 2 illustrates the definition of inter-subject correspondences.

In case of the rib-cage, corresponding landmarks are obtained for both intra- and inter-patient models from corresponding mesh vertices since the same model was adapted to all patients.

(a) Reference grid of lung model   (b) Patient individualized grid
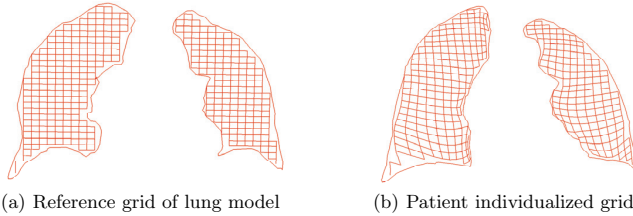
Fig. 2. Establishing inter-patient correspondences. Reference cartesian is deformed by calculated thin-plate spline deformation field which uses corresponding vertices of the outer lung surfaces. Examplarily, warping is shown for two-dimensional contour.

2.4 Motion Model Adaptation

The task of estimating the patient's motion field $\mathbf{v}$ under the assumption of sparse motion indicators expressed as $\mathbf{v}_b$ can be modeled as a conditional distribution $p(\mathbf{v}_a|\mathbf{v}_b)$ where $\mathbf{v}_a$ are the motion vectors of the vertices to be estimated. We are interested in the maximum likelihood estimation of the conditional distribution under the condition that the motion of a small set of landmarks is known. For that purpose, we partition the motion field $\mathbf{v}$ into to two disjoint subsets $\mathbf{v}_a$ and $\mathbf{v}_b$ resulting in

$$\mathbf{v} = \begin{pmatrix} \mathbf{v}_a \\ \mathbf{v}_b \end{pmatrix} \tag{2}$$

and equivalent partitions for the mean vector $\bar{\mathbf{v}}$ and the covariance matrix $\boldsymbol{\Sigma}$

$$\bar{\mathbf{v}} = \begin{pmatrix} \bar{\mathbf{v}}_a \\ \bar{\mathbf{v}}_b \end{pmatrix} \qquad \boldsymbol{\Sigma} = \begin{pmatrix} \boldsymbol{\Sigma}_{aa} & \boldsymbol{\Sigma}_{ab} \\ \boldsymbol{\Sigma}_{ba} & \boldsymbol{\Sigma}_{bb} \end{pmatrix} . \tag{3}$$

The maximum likelihood estimate $\bar{\mathbf{v}}_{a|b}$ of the conditional distribution given $\mathbf{v}_b$ can be calculated as [13]

$$\bar{\mathbf{v}}_{a|b} = \bar{\mathbf{v}}_a + \boldsymbol{\Sigma}_{ab}\boldsymbol{\Sigma}_{bb}^{-1}(\mathbf{v}_b - \bar{\mathbf{v}}_b) . \tag{4}$$

Note that usually, $\boldsymbol{\Sigma}_{bb}^{-1}$ is not invertible owing to multi-collinearity in the landmark positions and unreliable due to chance covariance in a limited training set. Therefore, we apply some regularization. In this case, we perform a ridge regression [14] by replacing $\boldsymbol{\Sigma}_{bb}^{-1}$ with $\tilde{\boldsymbol{\Sigma}}_{bb}^{-1} = \boldsymbol{\Sigma}_{bb}^{-1} + \gamma\mathbf{I}$, where γ is a positive and typically small constant.

2.5 Landmark Selection

For many clinical applications, an indication of optimal positions for sparse motion indicators is of special interest, e.g., in case of MRI navigator images. Assume that we were able to place motion indicators at arbitrary positions inside the lung neglecting any practical limitations. Then the question arises what will be the optimal choice for indicator positions? We tackle the problem in an iterative procedure using the properties of our created motion model. In order to find positions with most predictive power, we try out any motion vector as motion predictor and select the corresponding landmark position that belongs to the motion vector which reduces the variance of the model more than others. After calculating the modes invariant with respect to that chosen motion vector, we rerun the selection.

For the pointwise selection, we follow the formulation given in [15], where the variability of the weight vectors $\mathbf{b}_i$ before and after creating the invariant modes are compared.

Assuming the same partioning into two disjoint subsets as derived for the mean and covariance in Eq. 3 also for the corresponding matrix $\boldsymbol{\Phi}$ containing

the eigenvectors, we obtain:

$$\boldsymbol{\Sigma} = \boldsymbol{\Phi}\mathbf{D}\boldsymbol{\Phi}^T = \begin{pmatrix} \boldsymbol{\Phi}_a\mathbf{D}\boldsymbol{\Phi}_a^T & \boldsymbol{\Phi}_a\mathbf{D}\boldsymbol{\Phi}_b^T \\ \boldsymbol{\Phi}_b\mathbf{D}\boldsymbol{\Phi}_a^T & \boldsymbol{\Phi}_b\mathbf{D}\boldsymbol{\Phi}_b^T \end{pmatrix} \quad \text{with} \quad \boldsymbol{\Phi} = \begin{pmatrix} \boldsymbol{\Phi}_a \\ \boldsymbol{\Phi}_b \end{pmatrix} . \tag{5}$$

The diagonal matrix of eigenvalues is denoted as $\mathbf{D}$.

Following [15], we express the invariant mode $\hat{\mathbf{b}}_i$ by

$$\hat{\mathbf{b}}_i = \mathbf{b}_i - \mathbf{R}_b\boldsymbol{\Phi}_b\mathbf{b}_i = (\mathbf{I} - \mathbf{R}_b\boldsymbol{\Phi}_b)\mathbf{b}_i \,, \tag{6}$$

with $\mathbf{R}_b = \mathbf{D}\boldsymbol{\Phi}_b^T[\boldsymbol{\Phi}_b\mathbf{D}\boldsymbol{\Phi}_b^T]^{-1}$.

Comparing the variability of the weight vectors before and after creating the invariant mode with respect to the stimulator expressed as $\mathbf{v}_b$, the points with maximal predictive power can be found. Note that in this case $\mathbf{v}_b$ contains the motion vector corresponding to one landmark position. When applying this procedure iteratively, the k best predictors can be found.

3   Results

As shown in [9], the motion fields extracted as in Sect. 2.2 provide a landmark accuracy of about the voxel size. Thus, for the following evaluation of our motion models, we assume the extracted motion fields as our ground truth. While this enables us to evaluate the performance of our models at a few thousand discrete positions meaning the mesh vertices, it has to be noted that there might be a slight bias compared to landmarks set by experts.

For evaluation of our motion models, we are interested in the possibility to predict inhale-exhale motion fields of the lungs. We compare the prediction of anatomical related regions to systematic landmark selection from Sect. 2.5. In each case prediction is performed as described in Sect. 2.4. As anatomical related regions, we focus on the main breathing motors which are diaphragm and rib-cage. The diaphragm motion is extracted by manually labelling the dome of the surfaces close to the diaphragm on all lung models. As stimulators for model adaptation as described in Sect. 2.4, we then use all motion vectors belonging to the vertices of the labelled surfaces. In case of the rib-cage, we focus on the rib motion since there is almost no spine motion.

Intra patient motion models are evaluated in a leave-one-out study. In each case, one week is chosen and a motion model is built out of all remaining weeks and adapted to the 'unseen' motion field. The estimated motion field is then compared to the respective extracted motion field and the error at mesh vertices is evaluated. The results when using diaphragm and rib-cage as model stimulators are given in Tab. 1. By comparing the prediction power of the respective anatomical region to the mean motion model $\bar{\mathbf{v}}$, it can be seen that model stimulation using the diaphragm denoted as PCA_{DIA} significantly improves the estimation while the rib-cage (PCA_{RC}) even worsens the prediction. However, using both diaphragm and rib motion (PCA_{RCD}) as model stimulators yields almost the optimal parameter fit with an error of 2.8 mm in average. The best

possible prediction that can be achieved with our model is given for comparison by projecting the true motion field in the PCA space. Figure 3 shows the prediction in the sagittal view for two selected cases.

Inter patient motion models are evaluated in a similar leave-one-out study. Motion models of all patients besides the one under consideration are built and adapted to all weeks of the corresponding patient with results given in Tab. 1. Although we observed significant differences in breathing patterns between patients, there have to be similarities in the respective motion fields which can be seen from the fact that the mean model already compensates about 40 % of the breathing motion. In case of the inter-patient model, we took into account 20 eigenvectors that correspond to the largest eigenvalues covering about 98% of the variance. Again, the prediction using the diaphragm is much better than using the rib-cage with 4.1 mm in average compared to 7.4 mm.

When applying landmark selection from Sect. 2.5 on both intra- and inter-patient models, the predictive power of the first N-landmarks was investigated. In each case, we additionally took all neighboring vertices into account to be less sensitive to the exact landmark position. Having motion information only at the positions of the first $N=3$ selected landmarks including their neighbors yields an error of 4.0 (5.8) mm in average for intra-(inter-) patient models and for $N=5$ and $N=7$ an error of 3.5 (5.1) mm and 3.0 (4.8) mm, respectively. A typical distribution for $N=5$ is given in Fig. 4. For many cases, it could be observed that the first selected landmark was located close to the diaphragm while the second one in the anterior part of the corresponding other upper lung.
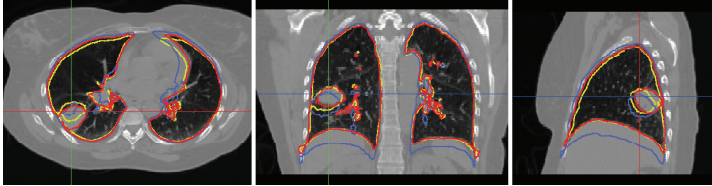
Fig. 3. Result of motion estimation shown in three orthogonal slices for Patient 3 using the inter-subject motion model and assuming the motion of the dome of the diaphragm to be known. Inhale contour is shown blue, red exhale, and yellow estimation.

4 Conclusion

Breathing motion is a complicated factor in several applications working on the thorax or upper abdomen. In many cases, the treatment would benefit from prior knowledge of the organ deformation and location. Although latest image devices are able to acquire dynamic images covering the respiratory motion, e.g.,

Table 1. Result of intra- and inter-patient motion models. Approximation of inhale-exhale motion by model in leave-one-out test. All values are given in millimeters. Performance of mean motion model as well as PCA model with optimal parameter fit indicates model quality. Prediction error is given when using different model stimulators. Mean motion of entire lung and respective regions presented for comparison ($\mathbf{m}$, $\mathbf{m}_{DIA}$, $\mathbf{m}_{RC}$). Last row gives mean value of all rows as absolute value and relative to mean lung motion in percent. For details see text.

Intra Patient Model								
	$\bar{\mathbf{v}}$	PCA	PCA_{DIA}	PCA_{RC}	PCA_{RCD}	$\bar{\mathbf{m}}$	$\bar{\mathbf{m}}_{DIA}$	$\bar{\mathbf{m}}_{RC}$
Pat. 1	4.1 ± 1.2	2.1 ± 0.7	2.3 ± 0.7	4.8 ± 1.9	2.2 ± 0.7	10.9 ± 1.7	15.4	3.7
Pat. 2	4.5 ± 1.4	2.8 ± 0.9	4.1 ± 1.3	4.9 ± 1.6	3.2 ± 0.9	12.6 ± 2.4	23.1	9.7
Pat. 3	2.2 ± 0.5	1.4 ± 0.3	1.5 ± 0.4	2.0 ± 0.5	1.5 ± 0.4	6.9 ± 1.5	12.0	2.4
Pat. 4	4.6 ± 1.9	1.7 ± 0.5	1.9 ± 0.6	3.3 ± 1.5	2.0 ± 0.4	10.0 ± 4.1	20.9	3.4
Pat. 5	4.7 ± 1.5	3.3 ± 0.7	4.0 ± 0.7	4.9 ± 1.3	3.9 ± 0.6	7.3 ± 2.2	11.1	4.6
Pat. 6	4.2 ± 0.6	3.8 ± 0.8	4.2 ± 1.2	4.8 ± 0.5	3.9 ± 1.0	10.8 ± 2.1	20.2	3.2
Pat. 7	3.2 ± 0.7	3.0 ± 0.8	3.3 ± 1.0	3.4 ± 1.4	3.1 ± 1.1	10.0 ± 1.4	18.7	4.1
mean	3.9 (39.8)	2.6 (26.4)	3.0 (30.6)	4.0 (40.1)	2.8 (28.9)	9.8	17.3	4.4

Inter Patient Model								
	$\bar{\mathbf{v}}$	PCA	PCA_{DIA}	PCA_{RC}	PCA_{RCD}	$\bar{\mathbf{m}}$	$\bar{\mathbf{m}}_{DIA}$	$\bar{\mathbf{m}}_{RC}$
Pat. 1	6.3 ± 1.2	3.2 ± 0.5	4.1 ± 0.8	7.0 ± 1.6	4.1 ± 0.8	10.9 ± 1.7	15.4	3.7
Pat. 2	8.6 ± 1.8	4.3 ± 0.6	6.1 ± 1.0	11.0 ± 3.0	5.0 ± 0.9	12.6 ± 2.4	23.1	9.7
Pat. 3	3.6 ± 0.4	1.8 ± 0.1	2.7 ± 0.3	5.0 ± 1.0	2.1 ± 0.3	6.9 ± 1.5	12.0	2.4
Pat. 4	6.7 ± 1.8	2.8 ± 0.7	3.9 ± 1.0	6.6 ± 2.4	3.6 ± 1.0	10.0 ± 4.1	20.9	3.4
Pat. 5	5.9 ± 1.6	2.5 ± 0.4	3.1 ± 0.4	7.4 ± 0.8	3.2 ± 0.5	7.3 ± 2.2	11.1	4.6
Pat. 6	8.0 ± 1.2	4.0 ± 1.0	4.8 ± 1.0	9.0 ± 2.3	4.6 ± 1.0	10.8 ± 2.1	20.2	3.2
Pat. 7	4.3 ± 1.0	2.9 ± 0.4	4.1 ± 0.5	6.0 ± 0.7	3.6 ± 0.6	10.0 ± 1.4	18.7	4.1
mean	6.2 (63.3)	3.1 (31.6)	4.1 (41.9)	7.4 (75.8)	3.7 (37.7)	9.8	17.3	4.4

4D-CT, there is not always a multiphase breathing gated examination available. Thus, in this paper, we focused on motion model creation from inhale-exhale pairs of CT data sets and adaptation to 'unseen' data. Due to the fact that the images were acquired at breath-hold, no image artifacts caused by breathing were present. From the extracted motion fields, patient-specific models but also inter-subject models have been built and compared. For model stimulation, different sparse motion indicators have been investigated. If the diaphragm motion is known, we achieved a prediction error of 3.0 (4.1) mm for the intra (inter)-patient model covered. Using sparse motion information obtained from the rib-cage is thereby less appropriate. With a systematic selection of landmarks, most important regions for providing sparse motion information have been found. Although there are several attempts for patient-specific modeling, it is to the best of our knowledge the first time that an inter-subject breathing model based on statistical properties of extract motion fields has been built and adapted to patient data.

Since we worked so far on a small size of training data, improvement can probably be expected when enlarging the amount of patient data. Future studies will be carried out on multiphase data taking the dynamic properties of respiration into account.

(a) First Invariant Mode   (b) Selected Points (Front)   (c) Selected Points (Top)
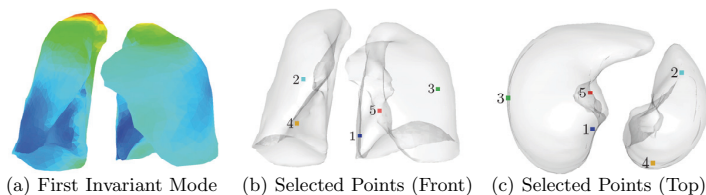
Fig. 4. Systematic point selection. Potential predictive power of each landmark position displayed on mesh surface (a). Ratio of variance of coefficients $\hat{\mathbf{b}}_i$ and $\mathbf{b}_i$ from Eq. 6 color coded from blue (small) to red (high). Selected first five points color coded from first (blue) to fifth (red) shown in (b) and (c). First selected point typically close to diaphragm while second in corresponding other lung.

References

1. Balter, J.M., ten Haken, R.K., Lawrence, T.S., Lam, K.L., Robertson, J.M.: Uncertainties in CT-based radiation therapy treatment planning associated with patient breathing. International Journal on Radiation Oncology, Biology, Physics **36**(1) (1996) 167–174
2. Villard, P.F., Beuve, M., Shariat, B., Baudet, V., Jaillet, F.: Simulation of lung behaviour with finite elements: Influence of bio-mechanical parameters. In: Proc. MEDIVIS. (2005) 9–14
3. Werner, R., Ehrhardt, J., Schmidt, R., Handels, H.: Modeling respiratory lung motion: a biophysical approach using finite element methods. In: Proc. SPIE. Volume 6916. (2008) 69160N1–11
4. McClelland, J.R., Blackall, J.M., Tarte, S., et al.: A continuous 4D motion model from multiple respiratory cycles for use in lung radiotherapy. Medical Physics **33**(9) (2006) 3348–3358
5. Zhang, Q., Pevsner, A., Hertanto, A., Hu, Y.C., Rosenzweig, K.E., Ling, C.C., Mageras, G.S.: A patient-specific respiratory model of anatomical motion for radiation treatment planning. Medical Physics **34**(12) (2007) 4772–4782
6. Sundaram, T.A., Avants, B.B., Gee, J.C.: A dynamic model of average lung deformation using capacity-based reparameterization and shape averaging of lung mr images. In: Proc. MICCAI. Volume 3217. (2004) 1000–1007
7. von Siebenthal, M., Szkely, G., Lomax, A., Cattin, P.: Inter-subject modelling of liver deformation during radiation therapy. In: Proc. MICCAI. Volume 4791. (2007) 659–666

8. Klinder, T., Lorenz, C., von Berg, J., Renisch, S., Blaffert, T., Ostermann, J.:
   4DCT image-based lung motion field extraction and analysis. In: Proc. SPIE
   Medical Imaging. Volume 6914. (2008) 69141L1–11
9. Vik, T., Kabus, S., von Berg, J., Ens, K., Dries, S., Klinder, T., Lorenz, C.: Val-
   idation and comparison of registration methods for free-breathing 4d lung-ct. In:
   Proc. SPIE Medical Imaging. Volume 6914. (2008) 69142P1–10
10. Klinder, T., Lorenz, C., von Berg, J., Dries, S., Blow, T., Ostermann, J.: Auto-
    mated model-based rib cage segmentation and labeling in CT images. In: Proc.
    MICCAI. Volume 4792. (2007) 195–203
11. Weese, J., Kaus, M., Lorenz, C., et al.: Shape constrained deformable models for
    3D medical image segmentation. In: Proc. IPMI. (2001) 380–387
12. Blaffert, T., Barschdorf, H., von Berg, J., Dries, S., et al.: Lung lobe modeling and
    segmentation with individualized surface meshes. In: Proc. SPIE Medical Imaging.
    Volume 6914. (2008) 69141I1–10
13. Bishop, C.M.: Pattern Recognition and Machine Learning. Springer (2006)
14. Hoerl, A., Kennard, R.: Ridge regression: Biased estimation for nonorthogonal
    problems. Technometrics 12(1) (1970) 5567
15. Hug, J., Brechbühler, C., Szekely, G.: Model-based initialisation for segmentation.
    In: Proc. ECCV. Volume 1843. (2000) 290–306

Registration-based lung tissue mechanics assessment during tidal breathing

Kai Ding[1], Kunlin Cao[2], Shalmali V. Bodas[1], Gary E. Christensen[2]
Eric A. Hoffman[3,1], and Joseph M. Reinhardt[1]

[1] Department of Biomedical Engineering
[2] Department of Electrical and Computer Engineering
[3] Department of Radiology
The University of Iowa, Iowa City, IA 52242

{kai-ding, kunlin-cao, shalmalividyadhar-bodas, gary-christensen,
eric-hoffman,joe-reinhardt}@uiowa.edu

Abstract. Lung tissue expansion and contraction can be assessed by acquiring multiple 3D CT images at different lung volumes. Static "breath-hold" imaging has been shown to produce tissue deformation estimates that match well with other measures of lung function. However, dynamic imaging protocols that image the breathing lung may produce more physiologically meaningful estimates of lung function. We use non-linear image registration to match retrospectively reconstructed respiratory-gated lung "dynamic" CT volumes acquired during tidal breathing. We compare the lung expansion and contraction estimates from the dynamic acquisitions to regional ventilation assessed by xenon-enhanced CT. The Jacobian-based lung volume change estimate shows a good agreement with the xenon-CT (average $r^2 = 0.85$) at image pair acquired at 50% and 75% of the inspiration duration.

1   Introduction

The lungs expand and contract during the respiratory cycle. Lung tissue mechanics depends on the material properties of the lung parenchyma and the mechanical inter-relationships between the lungs, diaphragm, and other parts of the respiratory system. Pulmonary diseases, such as fibrosis and emphysema, can change the tissue material properties of lung parenchyma and the associated lung function.

Lung tissue expansion and contraction can be assessed by acquiring multiple 3D CT images at different lung volumes. Static "breath-hold" imaging has been shown to produce tissue expansion estimates that match well with other measures of lung expansion [1]. However, static imaging is not able to observe kinematic effects, and dynamic imaging protocols that image the breathing lung may produce more physiologically meaningful estimates of lung function.

Various efforts have been made to non-invasively assess lung function. Guerrero et al. have used optical-flow registration to compute lung ventilation from 4D CT [2,3]. Gee et al. have used non-rigid registration to study pulmonary

kinematics [4] using magnetic resonance (MR) imaging. Christensen et al. have
used image registration to match images across cine-CT sequences and esti-
mate rates of local tissue expansion and contraction [5]. While they were able to
show that their accumulated measurement matched well with the global mea-
surement, they were not able to compare the registration-based measurements
to local measures of regional tissue ventilation. Others have used hyperpolar-
ized gas MR imaging of the lung to assess lung function and to demonstrate
pathophysiological changes [6]. Xenon-enhanced CT (Xe-CT) has been used to
measure regional ventilation by observing the gas wash-in and wash-out rate on
serial CT images [7–9]. All of these approaches have some disadvantages, includ-
ing limited spatial coverage, limited spatial resolution, high cost, or the need for
special equipment or personnel.

We have previously shown a good correlation (linear regression, average
$r^2 = 0.73$) between specific ventilation measured from Xe-CT and specific vol-
ume change assessed by image registration [1]. By applying 3D registration to
CT images of the lung acquired at different levels of inflation (static breath-
hold imaging), we estimate the regional lung expansion from the displacement
field calculated from the image registration. We compared these lung expan-
sion estimates to Xe-CT derived measures of regional ventilation to validate our
measurements and to establish their physiological significance.

However, it is reasonable to think that the static breath-hold scenario does
not fully depict the behavior and function of the moving, breathing lung. Dy-
namic imaging, where images are acquired across the respiratory cycle, may
provide better estimates of lung mechanics, and may more accurately reflect
the behavior of the breathing lung. In this paper we compare registration-based
estimates of lung mechanics derived from dynamic imaging protocols to Xe-CT
estimates of lung ventilation and evaluate the accuracy of the registration applied
in this study.

2   Materials and Methods

2.1   Data Acquisition

Appropriate animal ethics approval was obtained for these protocols from the
University of Iowa Animal Care and Use Committee and the study adhered
to NIH guidelines for animal experimentation. Four adult sheep were used for
experiments. The sheep were anesthetized using intravenous pentobarbital and
mechanically ventilated during experiments. The dynamic respiratory-gated CT
images are acquired with the animals in the supine position using the dynamic
imaging protocol with a pitch of 0.1, slice collimation of 0.6 mm, rotation time of
0.5 sec, slice thickness of 0.75 mm, slice increment of 0.5 mm, 120 kV, 400 mAs,
and kernel B30f. Images are reconstructed retrospectively at 0, 25, 50, 75, and
100% of the inspiration duration and 75, 50 and 25% of the expiration duration
(herein denoted as the $T0$, $T1$, $T2$, $T3$, $T4$, $T5$, $T6$, and $T7$ images). Twelve
contiguous axial locations and approximately 40 breaths for Xe-CT studies are
selected from the whole lung volumetric scan performed near end-expiration.

Images are acquired with the scanner set in ventilation triggering mode with 80 keV energy (for higher Xe signal enhancement), 160 mAs tube current, a 360° rotation, a 0.33 sec scan time, and 2.4 mm slice thickness. Respiratory gating is achieved by replacing the standard ECG gating signal with a trigger signal from a LabView program. Both of the two types of images are acquired with a matrix of 512 by 512 and without moving the animal between scans, so after acquisition the data sets are in rigid alignment.

2.2 Image Registration

Inverse consistent linear elastic image registration was applied to register phase change image pairs [10]. The registration estimates the inverse consistency error between the forward and reverse transformation, so it provides more accurate correspondences between two images compared to independent forward and reverse transformations. The registration minimizes the cost function defined as:

$$C = \sigma[C_{SIM}(I_0 \circ h, I_1) + C_{SIM}(I_1 \circ g, I_0)] + \tag{1}$$

$$\chi[C_{ICC}(u, \tilde{w}) + C_{ICC}(w, \tilde{u})] +$$

$$\rho[C_{REG}(u) + C_{REG}(w)],$$

where the forward transformation h is used to deform the image I_0 into the shape of the image I_1, and the reverse transformation g is used to deform the shape of I_1 into that of I_0. The deformed template and target images are denoted by $(I_0 \circ h)$ and $(I_1 \circ g)$, respectively. The vector displacement function $\boldsymbol{u}(x, y, z)$ that maps image I_0 to image I_1 is used to calculate the local lung expansion using the Jacobian determinant $J(x, y, z)$ defined as:

$$J(x, y, z) = \begin{vmatrix} 1 + \frac{\partial u_x(x,y,z)}{\partial x} & \frac{\partial u_x(x,y,z)}{\partial y} & \frac{\partial u_x(x,y,z)}{\partial z} \\ \frac{\partial u_y(x,y,z)}{\partial x} & 1 + \frac{\partial u_y(x,y,z)}{\partial y} & \frac{\partial u_y(x,y,z)}{\partial z} \\ \frac{\partial u_z(x,y,z)}{\partial x} & \frac{\partial u_z(x,y,z)}{\partial y} & 1 + \frac{\partial u_z(x,y,z)}{\partial z} \end{vmatrix},$$

where $u_x(x, y, z)$ is the x component of $\boldsymbol{u}(x, y, z)$, $u_y(x, y, z)$ is the y component of $\boldsymbol{u}(x, y, z)$, and $u_z(x, y, z)$ is the z component of $\boldsymbol{u}(x, y, z)$.

The C_{SIM} term of the cost function defines the symmetric intensity similarity. The C_{ICC} term is the inverse consistency constraint or inverse consistency error cost and is minimized when the forward and reverse transformations are inverses of each other. The C_{REG} term is used to regularize the forward and reverse displacement fields. The functions $u, w, \tilde{u}, \tilde{w}$ are voxel displacement fields and are related to the forward and reverse transformations by the equations: $h(x) = x + u(x)$, $g(x) = x + w(x)$, $h^{-1}(x) = x + \tilde{u}(x)$, $g^{-1}(x) = x + \tilde{w}(x)$. The constants σ, χ and ρ are used to enforce/balance the constraints. In our registrations, we set the weighting constants $\sigma = 1$, $\chi = 600$, and $\rho = 0.00125$. The parameters were made on the basis of pilot experiments, previous work and experience.

The Jacobian measures the differential expansion at position (x, y, z) in the image I_1. If the Jacobian is greater than one, there is local tissue expansion; if the Jacobian is less than one, there is local tissue contraction (Figure 1).

In our study, we register $T0$ to $T1$, $T1$ to $T2$, ..., and $T6$ to $T7$ in the dynamic respiratory-gated image sequence and thus acquire seven transformation pairs.
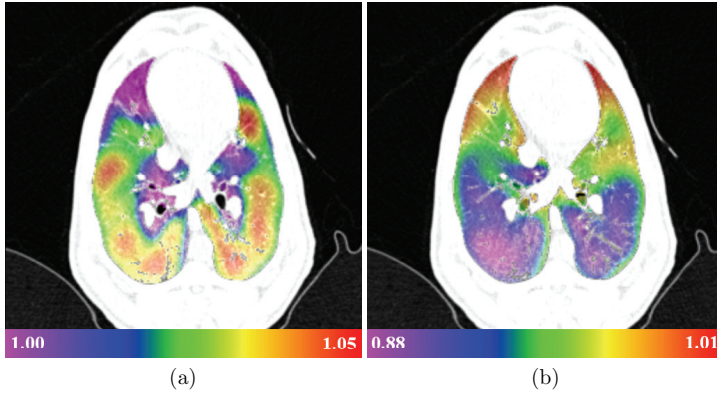
Fig. 1. Color-coded maps showing (a) the Jacobian of the image registration transformation (unitless) for approximately the same anatomic slice computed from the $T0 - T1$ inspiration image pair and (b) the $T4 - T5$ expiration image pair. Note that the color scales are different for (a) and (b). Red regions on the inspiration image (a) are regions that have high expansion while dark blue regions on the expiration image (b) have high contraction.

2.3 Assessment of Image Registration Accuracy

For each dynamic respiratory-gated image sequence, 20 anatomic landmarks were matched across all eight image $T0, ..., T7$. The selected landmarks were recognizable branchpoints of the vascular and airway branches. For each landmark, the actual landmark position was compared to the registration-derived estimate of landmark position and the error was calculated.

2.4 Jacobian-Based Lung Expansion Compared to Lung Ventilation

Xe-CT estimates of specific ventilation (sV) are computed using the "time-series image analysis" (TSIA) software described in [9]. To compare the Jacobian values with the sV, we must identify corresponding regions in the two images. The Xe-CT has only twelve slices of axial coverage and the data sets are acquired in rigid

alignment as described in Section 2.1, so we register the twelve-slice Xe-CT data
to the $T0$ whole-volume dynamic respiratory-gated CT data using rigid affine
registration as shown in Figure 2. We subdivide the Xe-CT data into 30 slabs
along the y (ventral–dorsal) axis. We track the deformation of each slab across
the sequence of volume images (i.e., from $T0$ to $T1$ to $T2$, etc.) and compare the
average Jacobian within each slab to the corresponding average sV measurement
in the Xe-CT images.

(a)   (b)   (c)
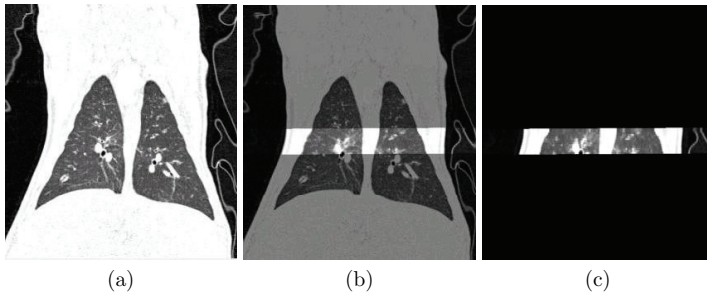
Fig. 2. Example of the result of affine registration between Xe-CT data and
dynamic respiratory-gated CT data. (a) $T0$ whole-volume dynamic respiratory-
gated CT data. (b) Fused image. (c) Deformed first breath of the Xe-CT data.

3   Results

3.1   Registration Accuracy

Figure 3a-b shows the projection of the manually-selected landmark locations
onto coronal and sagittal slices for one animal. Figure 4 shows the registration
accuracy as assessed by predicting the motion of the 20 manually-defined land-
marks across seven phase change pairs. Overall the registration accuracy is on
the order of 1 mm, or about 2 voxels.

3.2   Lung Expansion and Xe-CT Estimates of sV

Figure 5a shows the average Jacobian vs. lung height for all phase change pairs.
Figure 5b shows the average sV vs. lung height calculated from the Xenon-CT
study. Figure 6 shows average Jacobian vs. average sV at the $T0$ to $T1$ inspiration
phase change pair and $T4$ to $T5$ expiration phase change pair. The figure gives
the equation of the linear regression line with r^2 values and 95% confidence for
the linear fits between average sV and the average Jacobian. Figure 7 shows the
correlation coefficients r^2 from the linear regression of average Jacobian and sV
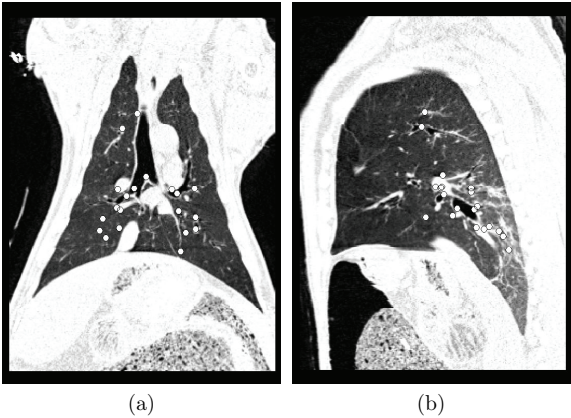
Fig. 3. Manually-selected landmark locations projected onto (a) a coronal slice and (b) a sagittal slice for one animal at T0 breathing phase.
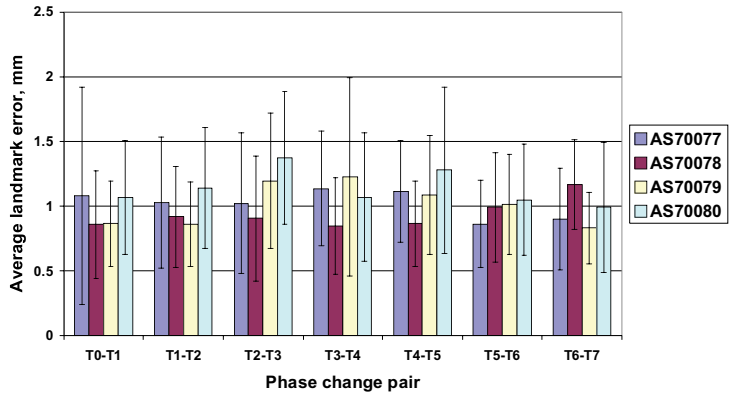
Fig. 4. Registration accuracy by mean ± standard deviation of landmark errors for each phase change pair and for each animal.

for each phase change pair and each animal. The phase change pair $T2$ to $T3$ shows the highest average correlation $r^2 = 0.85$ among all phase change pairs.
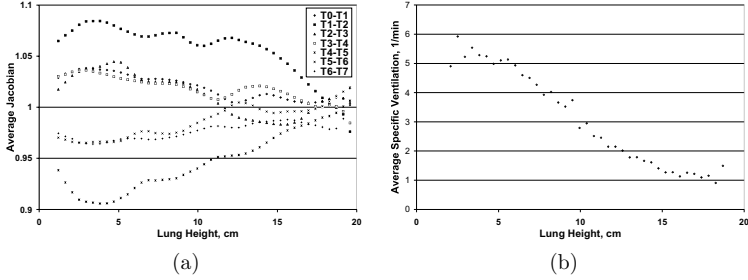
Fig. 5. Example of Jacobian and sV measurements vs. lung height for one animal. (a) Average Jacobian values for all phase change pairs and (b) average sV vs. lung height. Lung height is 0 cm is the most dorsal position and positive heights move toward the ventral direction.
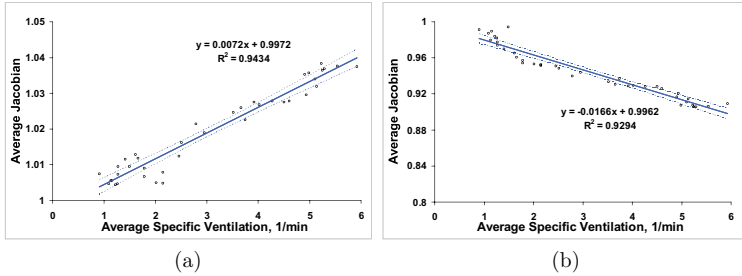
Fig. 6. Examples of scatter plot of average sV and average Jacobian for one animal with linear regression with 95% confidence interval for (a) $T0$ to $T1$ phase change pair and (b) $T4$ to $T5$ phase change pair.

4   Discussion and Summary

We have calculated estimates of lung expansion from the Jacobian of the registration deformation field during tidal breathing for respiratory phase change
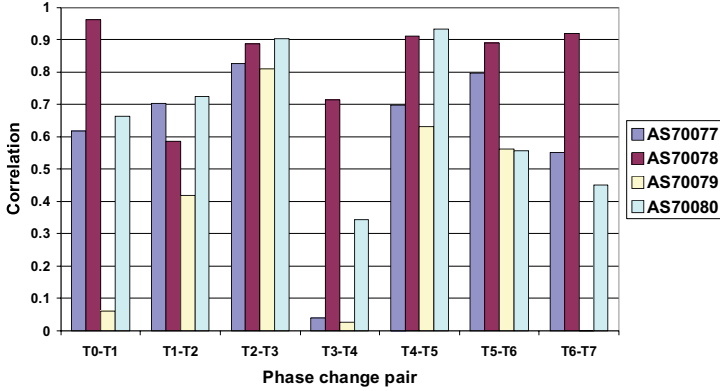
Fig. 7. Correlation coefficients r^2 from the linear regression of average Jacobian and sV for each phase change pair and for each animal.

pairs. The Jacobian values were compared to the Xe-CT based measures of specific ventilation. The correlation to the Xe-CT sV is higher for the Jacobian calculated from the dynamic respiratory-gated images ($r^2 = 0.84$) than for the static breath-hold images ($r^2 = 0.74$) as we reported [1]. The linear regression relationship in Figure 7 shows a wide range between animals and phases. We suspect it is caused by the natural anatomical difference between animals. Further more, since we used time based respiratory gating, the lung expand and contract differently between each phase change pair. For example, the lungs do not expand as much in $T3$ to $T4$ phase change pair as in $T2$ to $T3$ phase change pair. It would be interesting to determine if pressure based respiratory gating method will give better result.

The average registration error was about 1 mm in the phase change image pairs. Since the landmarks in these experiments were manually defined by picking anatomic features on a computer display, some component of this error is likely attributable to human error.

Since the Xe-CT data is collected over several breaths during tidal breathing, it is reasonable to expect that the Jacobian calculated from the dynamic respiratory-gated volume images would more closely reflect the ventilation patterns measured by the Xe-CT. The moving, breathing lung has mass, inertia, and hysteresis, and the true dynamics of the respiratory system are probably better revealed using images acquired across free breathing. Figure 8 shows the different arrival phases for different regions of the lung when largest expansion and largest contraction occur. Most of the lung regions will have the largest expansion at the middle phase ($T1$ to $T2$ or 25% to 50% of the inspiration du-

ration) and the largest contraction at the beginning phase ($T4$ to $T5$ or 100% to 75% of the expiration duration). It shows the lung does not expand or contract uniformly along the phases.

| T0-T1 | T1-T2 | T2-T3 | T3-T4 | T4-T5 | T5-T6 | T6-T7 |
| (a) | | | | (b) | | |
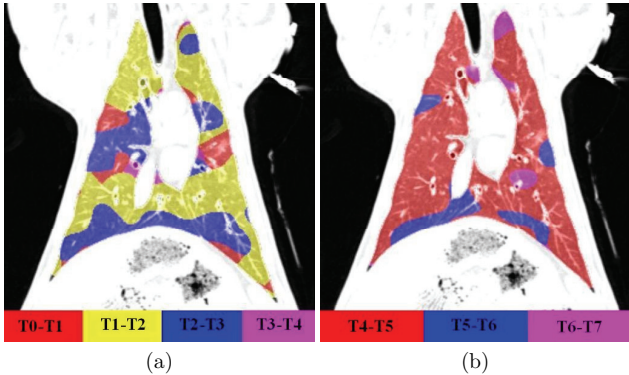
Fig. 8. (a) Color coded image showing the (a) the phase point when largest expansion occur and (b) the largest contraction occur. Note the color scales are different for (a) and (b).

Additional work is needed to validate these methods before translation to use in humans. For imaging humans during normal respiration, the dynamic imaging may pose some challenges. The retrospectively reconstructed respiratory-gated images require regular and repeatable breathing patterns, so the subject must be trained to breath properly before images are acquired. In addition, since image data is gathered over several breaths and then reconstructed at different phases, the dynamic imaging will deliver more radiation dose than a single pair of breath-hold scans. The development of low dose imaging and/or *prospective* respiratory-gating may be able to reduce the radiation exposure. With the validation of the best correlated phase change pair during tidal breathing, the subject will only need to be scanned at two certain phases and the radiation dose will be significantly reduced.

If these methods can eventually be extended to humans, they would provide important new tools for studying the lung. Xe-CT requires the use of expensive xenon gas and the associated hardware to control delivery of the gas and harvest the gas from expired air for recycling. In addition, it is known that xenon gas has a strong anesthetic effect that must be carefully monitored. Finally, Xe-CT imaging protocols require high temporal resolution imaging, so typically axial coverage is limited to just a few slices at a time. However, if a registration-based analysis of images acquired during tidal breathing could be registered across

respiratory phase, high-resolution maps of lung expansion could be produced for
the entire lung with low cost and dose.

5   Acknowledgments

References

1. Reinhardt, J.M., Ding, K., Cao, K., Christensen, G.E., Hoffman, E.A., Bodas, S.V.:
   Registration-based estimates of local lung tissue expansion compared to xenon-ct
   measures of specific ventilation. Medical Image Analysis (In press)
2. Guerrero, T., Sanders, K., Noyola-Martinez, J., Castillo, E., Zhang, Y., Tapia, R.,
   Guerra, R., Borghero, Y., Komaki, R.: Quantification of regional ventilation from
   treatment planning CT. Int. J. Radiation Oncology Biol. Phys. **62**(3) (Jul 1 2005)
   630–634
3. Guerrero, T., Sanders, K., Castillo, E., Zhang, Y., Bidaut, L., Komaki, T.P.R.:
   Dynamic ventilation imaging from four-dimensional computed tomography. Phys
   Med Biol. **51**(4) (Feb. 21 2006) 777–791
4. Gee, J., Sundaram, T., Hasegawa, I., Uematsu, H., Hatabu, H.: Characterization of
   regional pulmonary mechanics from serial magnetic resonance imaging data. Acad.
   Radiol. **10** (2003) 1147–1152
5. Christensen, G.E., Song, J.H., Lu, W., Naqa, I.E., Low, D.A.: Tracking lung tissue
   motion and expansion/compression with inverse consistent image registration and
   spirometry. Med Physics **34**(6) (June 2007) 2155–2165
6. Moller, H.E., Chen, X.J., Saam, B., Hagspiel, K.D., Johnson, G.A., Altes, T.A.,
   de Lange, E.E., Kauczor, H.U.: MRI of the lungs using hyperpolarized noble gases.
   Magnetic Resonance in Medicine **47**(6) (2002) 1029–1051
7. Marcucci, C., Nyhan, D., Simon, B.A.: Distribution of pulmonary ventilation
   using Xe-enhanced computed tomography in prone and supine dogs. J. Applied
   Physiology **90**(2) (2001) 421–430
8. Tajik, J.K., Chon, D., Won, C.H., Tran, B.Q., Hoffman, E.A.: Subsecond multi-
   section CT of regional pulmonary ventilation. Acad. Radiol. **9** (2002) 130–146
9. Chon, D., Simon, B.A., Beck, K.C., Shikata, H., Saba, O.I., Won, C., Hoffman,
   E.A.: Differences in regional wash-in and wash-out time constants for xenon-CT
   ventilation studies. Respiratory Physiology & Neurobiology **148** (2005) 65–83
10. Christensen, G., Johnson, H.: Consistent image registration. IEEE Transactions
    on Medical Imaging **20**(7) (July 2001) 568–582

Lung ventilation estimation based on 4D-CT imaging

Sven Kabus[1], Jens von Berg[1], Tokihiro Yamamoto[2], Roland Opfer[1], and
Paul J. Keall[2]

[1] Philips Research Europe – Hamburg, Germany
`sven.kabus@philips.com` *
[2] Department of Radiation Oncology, Stanford University School of Medicine,
Stanford, CA, USA

Abstract. The estimation of lung ventilation would allow to prevent
high functional lung regions from radiation injuries during radiation ther-
apy of lung cancer. As 4D-CT images are a standard procedure for radi-
ation therapy planning, the usage of these images for the additional lung
ventilation estimation would be advantageous over other imaging meth-
ods specifically performed for ventilation estimation only. This would re-
quire both a registration of images of different respiratory phases yielding
deformation vector fields and a suitable metric of ventilation estimation
based on the deformation vector fields. We apply two different image
registration methods and two different metrics to a set of 4D-CT images
and compare them with each other and with a global reference measure
based on independent lung volume measures from image segmentation.

1 Introduction

In radiation therapy planning of lung tumors it would be advantageous to con-
sider spatial information of lung function (e.g., ventilation) to prevent high func-
tional lung regions from radiation injury and achieve better quality of life. The
current standard of care for ventilation assessment is nuclear medicine (NM)
imaging [1, 2]. Recent techniques based on oxygen-enhanced magnetic resonance
(MR) imaging utilize hyperpolarized noble gases (e.g., ^{129}Xe) [3, 4]. Images ac-
quired with NM techniques suffer from a low spatial resolution. Also, for both of
NM and MR techniques, there are issues including a long scan time, high costs,
and low availability in radiotherapy departments. More recently, several groups
[5–8] have employed four-dimensional (4D) CT images [9, 10] for ventilation as-
sessment. Four-dimensional CT images for treatment planning can be used for
ventilation assessment as well, thus not requiring any additional imaging ses-
sions. Therefore, the 4D-CT-based approach potentially has several advantages
over NM or MR technique as it is free, fast, available, and of high spatial reso-
lution. However, the accuracy is dependent on image registration algorithms as

well as of the metrics for ventilation. Several investigators demonstrated varia-
tions in registration results between algorithms [7, 11, 12]. In the previous work,
we showed apparent discrepancies in deformation vector fields (DVFs) and ven-
tilation images calculated by different registration algorithms [8]. There have
been two different metrics used for ventilation assessment based on the DVFs.
Guerrero et al. have employed difference in Hounsfield units (HU) as a metric of
ventilation, and showed good agreements between global measures of ventilation
(i.e., tidal volume) for patient data calculated using this method and those of
manual measurements [6]. Reinhardt et al. have used the Jacobian determinant
and demonstrated good agreements between local ventilation for sheep data de-
termined by this approach and Xe-CT [5].

The contribution of the current work is to compare two different image reg-
istration methods with each other, both applied to 4D-CT and to also apply
both of the reported metrics to estimate ventilation from the resulting DVFs
on a set of four human 4D-CT scans. In addition, we compare our ventilation
metrics with the relative change in lung volume for each phase transition taken
from independent lung segmentation of each CT volume. Jacobian and relative
volume change should be equal by definition – however, from a technical point
of view this is not trivial and worth to show. Moreover, it can be considered as
a necessary condition for a reliable ventilation estimation.

2   Methods

2.1   Data Sets and Image Acquisition

The 4D-CT scans were performed on the GE Discovery ST multislice PET/CT
scanner (GE Medical Systems, Waukesha, WI) in cine mode at Stanford Cancer
Center. Case 1 of the four examined cases was a patient with lung cancer, the
other cases were breast cancer patients not showing lung metastases. During the
CT scan, patient respiratory traces were acquired using the Varian RPM system
(Varian Medical Systems, Palo Alto, CA), with the marker block placed on the
upper abdomen. Scan parameters were set as follows: 0.5 s gantry rotation, 0.45 s
cine interval, and 2.5 mm slice thickness. Each image reconstruction took 360 deg
of data. The projection images were retrospectively sorted into ten respiratory
phase-based bins of 3D CT image data (i.e., from 0% to 90% phase at 10%
intervals). Figure 1 shows coronal views of two exemplary phases of Case 3 and
all ten phase-to-phase subtraction images.

2.2   Registration Schemes

Given two reconstructed phases (one named reference phase, the other template
phase in the following), image registration tries to find a deformation vector field
$u : \mathbb{R}^3 \to \mathbb{R}^3$ such, that the displaced template phase is similar to the reference
phase. In this work, two fundamentally different registration schemes have been
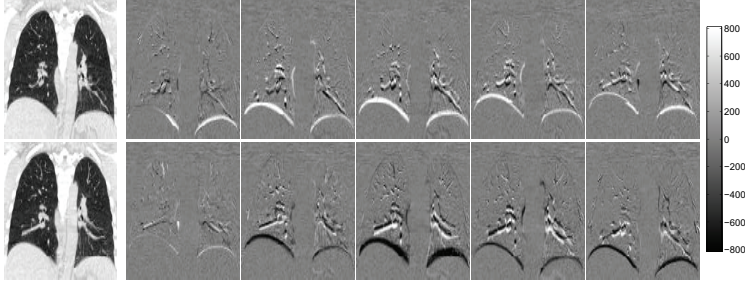
Fig. 1. Coronal views of reconstructed phases 30% and 80% (first column) and phase-to-phase subtraction images (remaining columns): row-wise are displayed $10\% - 0\%$, $20\% - 10\%$, ..., $0\% - 90\%$. Motion of diaphragm and vessel tree, particularly in the lower regions of the lung, is clearly visible. The change in sign for the diaphragm region occurs when changing from exhalation (0% to 60%) to inhalation (60% to 0%).

employed to estimate u. Whereas the first scheme matches surfaces (lung wall, vessel tree) followed by a thin-plate-spline interpolation to achieve the dense DVF, the second scheme is volumetric by itself.

2.3 Surface-based Registration

In the first, surface-based method we automatically determine the lung surface by a Hounsfield threshold at -650 HU and a marching cube triangulation in one phase. This surface not only covers the outer lung border but also the inner structures separating the parenchyma from the larger lung vessels. The number of triangles of this iso-surface ranges from 38,000 to 87,000. The method of deformable surface models [13] is now applied on this iso-surface of the lungs to adapt it to the second phase by minimizing the energy term $E = E_{ext} + \alpha E_{int}$. The external energy E_{ext} drives the mesh towards the surface points obtained in a surface detection step. The internal energy E_{int} restricts the flexibility by penalizing differences from the shape model. A number of such minimization steps is iteratively performed on the mesh. Details on the candidate point selection and on the calculation of the external energy can be found in [14]. The internal energy

$$E_{int} = \sum_j \sum_{k \in N(j)} ((\mathbf{v}_j - \mathbf{v}_k) - s\mathbf{R}(\hat{\mathbf{v}}_j - \hat{\mathbf{v}}_k))^2 \tag{1}$$

preserves shape similarity of all mesh vertices $\mathbf{v}_i$ to the model vertices $\hat{\mathbf{v}}_i$ from the initial iso-surface. $N(j)$ is the set of neighbors of vertex j. The neighboring vertices are those connected by a single triangle edge. The scaling factor s and the rotational matrix $\mathbf{R}$ are determined by a closed-form point-based registration method based on a singular value decomposition prior to calculation of (1) to

allow a similarity transformation (rigid transformation plus isotropic scaling) without effecting the internal energy.

By use of all vertex deformation vectors of the deformable surface, we construct a coarse displacement field for the lungs. Interpolation by thin-plate-splines on a subset of these vertices is used to create a dense field from it. The subset is obtained by replacing all vertices having a neighbor closer than 10 mm. Further details on this shape-tracking method, the impact of its parameters, and its computational performance are given in [14].

2.4 Volume-based Registration

The volumetric registration tries to find a DVF such, that the displaced template phase minimizes both a certain similarity measure $\mathcal{D}$ and a regularizing term $\mathcal{S}$. By adding a regularizing term, the registration problem is well-posed. For $\mathcal{D}$ we choose the popular sum of squared differences while for $\mathcal{S}$ an elastic regularizer [15] based on the Navier-Lamé equation is employed. The elastic regularizer assumes that the underlying images can be characterized as an elastic and compressible material. Its properties are modeled by the so-called Lamé constants λ, μ.

Based on calculus of variations we arrive at a system of non-linear partial differential equations to be solved,

$$\mu\Delta u + (\mu + \lambda)\nabla \cdot \nabla u = \nabla T_u(R - T_u) \, , \qquad (2)$$

with T_u and R corresponding to the displaced template phase and the reference phase, respectively.

For the discretization of (2) finite differences in conjunction with Neumann boundary conditions have been chosen. The resulting system of linear equations consists on one hand of a sparse, symmetric and highly structured matrix arising from the regularizer and, on the other hand, of a so-called force vector corresponding to the similarity measure. By nature, the larger the contrast of misaligned image structures is, the larger the modulus of the force vector is. Therefore, in CT images bone structures get typically perfectly matched whereas soft tissue may be not aligned. This holds particularly for the lung-rib interface with on the one side the parenchyma following the breathing motion and on the other side the ribs staying in place or even moving in opposite direction. To circumvent mis-alignment of parenchymal structures we added a simple masking of the force vector. For every voxel with Hounsfield value above 0 HU in the reference phase, the force vector is set to zero for this voxel position. This results in a lung deformation which is not influenced by mis-alignment of the rib cage.

The corrected linear system of equations is then linearized and iteratively solved by a conjugate gradient scheme. The whole registration method is embedded into a multi-resolution setting (typical image pyramid has a resolution of $512 \times 512 \times 136$ at level 0, $256 \times 256 \times 136$ at level 1, $128 \times 128 \times 68$ at level 2 etc.; registration is executed on levels 4 to 1) and preceeded by an affine pre-registration.

3 Evaluation

For the datasets considered in this work no ground-truth such as nuclear medicine data or annotated landmark positions was available. But both registration methods have been previously validated with a method reported in [8]. For that validation similar CT scans with known annotated landmark positions were used. Unlike the current study, ten biphasic thorax CT data sets (max inhale and max exhale reconstructions only) were used. Two independent experts set corresponding landmark pairs (18–20 pairs) at vessel- and bronchi-bifurcations in all images. The landmark positions at one state were transformed using the DVFs obtained from the two registration methods and compared to the reference landmarks. For the validation the average Euclidean distance before and after registration have been compared. Given an average error of 5.99 mm ± 3.97 mm before registration, both methods showed a reduced error in all testcases (2.50 mm ± 2.16 mm for the surface-based method and 2.28 mm ± 1.87 mm for the volumetric method). Taking into account the relatively large difference between max inhale and max exhale state in that previous study, we expect a more accurate result and smaller registration errors for the consecutive registration in 10% steps given in the 4D-CT datasets investigated in this work.

To analyze the DVFs restricted to the lung, segmentation of the lung is required. For this purpose we use the segmentation scheme reported in [16]. The segmentation provides a binary mask of the lung but with the vessel tree excluded. To evaluate the segmentation result, we computed the total lung mass at each respiratory phase by integrating the density $\rho = (HU + 1000)/1000$ over the segmented lung. Then, the lung mass at each phase is subtracted from their mean (mean lung mass is 593 g, 558 g, 580 g, 628 g for patients 1 to 4, respectively). From these absolute differences, for patient 1 to 4 a mean error of 8.2 g ± 6.6 g, 11.0 g ± 4.5 g, 14.5 g ± 6.2 g, 6.5 g ± 3.7 g, respectively, result. Each DVF (restricted to those parenchyma voxels having a Hounsfield value above −1000 HU) is analyzed using the two different metrics:

1. **Jacobian analysis.** Given two phases, the corresponding DVF u maps each position in the reference phase onto the corresponding position in the template phase. Thus, the determinant of the Jacobian of $x + u(x)$ represents the local volume change at position x between reference phase and template phase,

$$\Delta V_{\text{Jac}}(x) := \det(\nabla(x + u(x))) - 1. \qquad (3)$$

   A value of zero indicates local volume preservation whereas a positive (negative) value corresponds to local expansion (contraction).

2. **Hounsfield analysis.** Given two phases and the corresponding DVF u, a local change in lung volume can be estimated from the relative difference of corresponding Hounsfield values (cf. [6]). By denoting a Hounsfield value at position x in phase i as $HU_i(x)$ the metric is defined by

$$\Delta V_{\text{HU}}(x) := 1000 \frac{HU_i(x + u(x)) - HU_{i+1}(x)}{HU_{i+1}(x)(HU_i(x + u(x)) + 1000)}. \qquad (4)$$

Again, a positive (negative) value indicates local expansion (contraction). Note that the images have been smoothed with a Gaussian kernel before computing ΔV_{HU}.

Moreover, the lung segmentation [16] of each phase was used to compute the lung volume and from that the relative change in lung volume for each phase transition. This we took as a global reference value for our volume change estimates.

4   Results

Without any further parameter tuning, both methods have been successfully applied to the four patient cases with ten phases each. The quality of registration has been assessed by visually inspecting each residual image (i.e., the subtraction image obtained after registration). Almost all registrations indicate an optimal match of lung wall and diaphragm. In addition, in most cases the vessel-tree has been correctly aligned. The results for method 2 are slightly better than for method 1 which can be explained by the volumetric nature of method 2 (see Figure 2 for an exemplary view). For outer-lung regions such as spine or rib cage the residual images from both methods indicate a worse alignment compared to the subtraction images obtained before registration (compare with Figure 1). This is based in method 1 on the extraction of the lung surface only and in method 2 on the threshold applied on the force vector.

According to Section 3 the resulting DVFs are analyzed using two different metrics to estimate the local change in lung volume (see second and third column of Figure 2 for exemplary views). By using the lung segmentation of each phase, the values ΔV_{Jac} and ΔV_{HU} are integrated within the lung (displayed in Figure 3). Dependent on the patient case these relative volume changes lie between -7% for exhalation and $+11\%$ for inhalation.

Finally, the volume change estimates from both registration methods and both analysis metrics are compared with our global reference segmentation (depicted red in Figure 3). A quantitative comparison is provided by Table 1. Here, for each registration method and each analysis metric the absolute difference between estimated volume change and segmentation-based volume change is given in percentage. Overall, the Jacobian analysis yields an error in volume change estimation of 0.8% and 0.4% for method 1 and 2, respectively, while the estimation error of the Hounsfield analysis is larger (1.4% and 0.9%).

5   Discussion and Conclusion

We have applied two fully independent image registration methods to 4D-CT lung scans. The resulting DVFs for each phase-to-phase transition are analyzed with two different metrics in order to estimate local lung ventilation. An overall inspection reveals similar estimates of Jacobian and Hounsfield analysis.
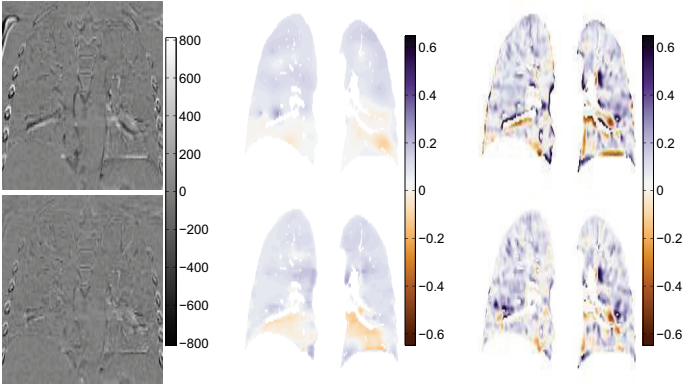
Fig. 2. For the same coronal view as in Figure 1, residual (left), Jacobian analysis (center), and Hounsfield analysis (right) are investigated after registration of 90% onto 80% phase with method 1 (top) and method 2 (bottom). Note that the two metrics are analyzed and visualized within the lung only – white regions belong to either the vessel tree or are outside the lung.

For a quantitative inspection we integrated these estimates over the lung and compared this estimated change of lung volume to that arising from an independent lung segmentation. Our results show that for both registration methods the volume change estimated by Jacobian analysis agrees well with the segmentation-based volume change (mean deviation of 0.8% and 0.4%). The Hounsfield analysis as the second investigated metric indicates a less optimal result (mean deviation of 1.4% and 0.9%). Since this metric is sensitive to local changes in the DVF but also to imaging- or reconstruction-related artifacts, a concluding rating is difficult. For a deeper insight, beside the global comparisons a point-based comparison between Jacobian and Hounsfield analysis seems to be worthwhile.

Table 1. Mean and standard deviation of absolute difference (in %) between estimated volume change and segmentation-based volume change.

Case	Jacobian analysis		Hounsfield analysis	
	method 1	method 2	method 1	method 2
1	0.84 ± 0.75	0.54 ± 0.48	1.75 ± 1.57	1.26 ± 1.59
2	0.74 ± 0.48	0.47 ± 0.42	1.29 ± 1.55	1.08 ± 1.04
3	1.09 ± 0.61	0.26 ± 0.29	1.59 ± 1.09	0.70 ± 0.94
4	0.50 ± 0.29	0.35 ± 0.25	0.91 ± 0.72	0.54 ± 0.50
mean	0.79 ± 0.53	0.41 ± 0.36	1.38 ± 1.23	0.89 ± 1.02
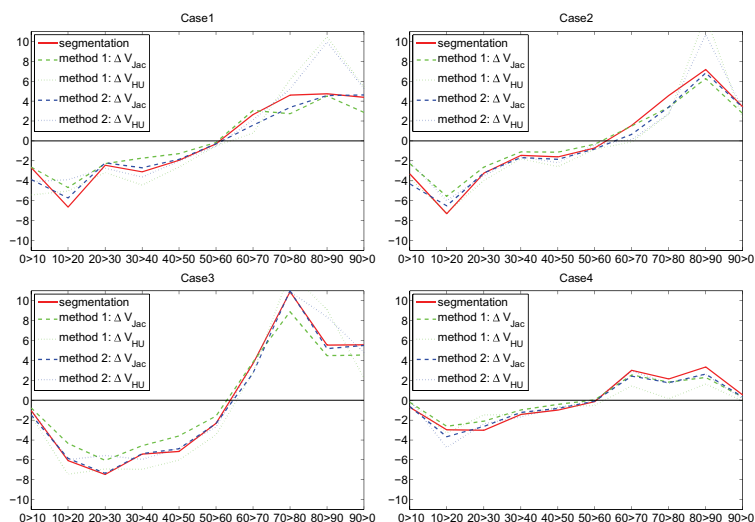
Fig. 3. Volume changes (in %) for each patient case and for each phase-to-phase transition obtained by either Jacobian analysis (dashed) or HU subtraction analysis (dotted) for method 1 (green) and method 2 (blue). For comparison, the volume change derived by a lung segmentation is depicted (red).

References

1. Suga, K.: Technical and analytical advances in pulmonary ventilation SPECT with Xenon-133 gas and Tc-99m-Technegas. Ann Nucl Med **16**(5) (2002) 303–10
2. Harris, R.S., Schuster, D.P.: Visualizing lung function with positron emission tomography. J Appl Physiol **102**(1) (2007) 448–58
3. Kauczor, H.U., Hanke, A., Van Beek, E.J.: Assessment of lung ventilation by MR imaging: current status and future perspectives. Eur Radiol **12**(8) (2002) 1962–70
4. Fain, S.B., Korosec, F.R., Holmes, J.H., O'Halloran, R., Sorkness, R.L., Grist, T.M.: Functional lung imaging using hyperpolarized gas MRI. J Magn Reson Imaging **25**(5) (2007) 910–23
5. Reinhardt, J.M., Christensen, G.E., Hoffman, E.A., Ding, K., Cao, K.: Registration-derived estimates of local lung expansion as surrogates for regional ventilation. Inf Process Med Imaging **20** (2007) 763–74
6. Guerrero, T., Sanders, K., Castillo, E., Zhang, Y., Bidaut, L., Pan, T., Komaki, R.: Dynamic ventilation imaging from four-dimensional computed tomography. Phys Med Biol **51**(4) (2006) 777–91
7. Yaremko, B.P., Guerrero, T.M., Noyola-Martinez, J., Guerra, R., Lege, D.G., Nguyen, L.T., Balter, P.A., Cox, J.D., Komaki, R.: Reduction of normal lung irradiation in locally advanced non-small-cell lung cancer patients, using ventila-

tion images for functional avoidance. Int J Radiat Oncol Biol Phys **68**(2) (2007) 562–71

8. Vik, T., Kabus, S., von Berg, J., Ens, K., Dries, S., Klinder, T., Lorenz, C.: Validation and comparison of registration methods for free-breathing 4D lung CT. In: Proc. of SPIE Medical Imaging. Volume 6914., SPIE (2008) 69142P–1–69142P–10

9. Keall, P.J., Starkschall, G., Shukla, H., Forster, K.M., Ortiz, V., Stevens, C.W., Vedam, S.S., George, R., Guerrero, T., Mohan, R.: Acquiring 4D thoracic CT scans using a multislice helical method. Phys Med Biol **49**(10) (2004) 2053–67

10. Rietzel, E., Pan, T., Chen, G.T.: Four-dimensional computed tomography: image formation and clinical protocol. Med Phys **32**(4) (2005) 874–89

11. Kashani, R., Hub, M., Kessler, M.L., Balter, J.M.: Technical note: a physical phantom for assessment of accuracy of deformable alignment algorithms. Med Phys **34**(7) (2007) 2785–8

12. Brock, K.K.: A multi-institution deformable registration accuracy study. Int J Radiat Oncol Biol Phys **69**(3) (2007) S44

13. Weese, J., Kaus, M., Lorenz, C., Lobregt, S., Truyen, R., Pekar, V.: Shape constrained deformable models for 3D medical image segmentation. In: Proc. of IPMI, Springer–Verlag (2001) 380–387

14. von Berg, J., Barschdorf, H., Blaffert, T., Kabus, S., Lorenz, C.: Surface based cardiac and respiratory motion extraction for pulmonary structures from multi-phase CT. In: Proc. of SPIE Medical Imaging. Volume 6511., SPIE (2007) 65110Y

15. Broit, C.: Optimal Registration of Deformed Images. PhD thesis, University of Pennsylvania (1981)

16. Wiemker, R., Opfer, R., Bülow, T., Rogalla, P., Steinberg, A., Dharaiya, E., Subramanyan, K.: Towards computer aided emphysema quantification on ultra-low-dose CT: Reproducibility of ventrodorsal gravity effect measurement and correction. In: Proc. of SPIE Medical Imaging. Volume 6514., SPIE (2007) 65140O

Estimating Respiratory Motion from Cone-Beam Projections

Jef Vandemeulebroucke[1,2,3], Patrick Clarysse[1], Jan Kybic[3], and David Sarrut[1,2]

[1] Université de Lyon, Creatis-LRMN Laboratory, INSA-Lyon, France
[2] Université de Lyon, Centre Léon Bérard, 28 rue Laennec, 69353 Lyon, France
[3] Center For Machine Perception, Czech Technical University, Prague, Czech Republic

Abstract. Respiratory motion introduces uncertainties when planning and delivering radiotherapy treatment to lung cancer patients. Cone-beam projections potentially constitute a valuable source of motion information that could serve for motion compensated reconstruction and to learn the relationship between internal motion and respiratory correlated signals. We propose a method for respiratory motion estimation directly from cone-beam projections by including prior knowledge about the patient's breathing motion. The method requires that a four-dimensional computed tomography is available from which a patient specific model is constructed. Each cone-beam projection is compared to cone-beam projection views of the model and motion estimation is accomplished by optimizing the model parameters with respect to a similarity measure. Experiments on simulated data show satisfying results. Experiments on real cone-beam projections are currently being undertaken in order to confirm these observations.

1 Introduction

External beam radiotherapy is the primary treatment modality for patients with non-operable lung cancer. Respiratory motion introduces uncertainties during imaging, treatment planning and treatment delivery [1]. To reduce respiratory motion induced image artifacts, respiratory-correlated acquisition techniques have been developed resulting in four-dimensional computed tomography images (4DCT). These images provide additional information about tumor and organ-at-risk position and trajectory that can be incorporated in the treatment planning process. To account for respiratory motion additional margins are considered during treatment planning and treatment delivery methods such as breath-hold and gating methods have been developed. These methods often use an external respiratory correlated signal such as lung air flow or abdominal height as a surrogate for tumor motion during the actual treatment. Changes of the respiratory motion over the duration of the treatment have been reported [2] and may influence the relationship between tumor motion and surrogate. Recent developments have made cone-beam CT (CBCT) mounted on the linear accelerator [3] available, which makes it possible to acquire patient images in treatment

position just prior to treatment. As with conventional CT, three-dimensional
(3D) CBCT is heavily influenced by respiratory motion. Respiratory correlated
CBCT [4] reduces respiratory motion artifacts. However, seeing only a subset of
the CB projections is used to reconstruct each 3D CBCT image, the resulting
image quality is lower. Motion compensated CB reconstruction techniques [5],
[6] have been proposed enabling the use of all the acquired projections. These
methods however require knowledge of the motion present during CB acquisition.

CB projections are potentially a valuable source of motion information that
could be used for motion compensated reconstruction and to study the rela-
tionship between an external surrogate signal and tumor motion. Zijp *et al.* [7]
proposed a fast and robust method to extract the breathing phase from a se-
quence of CB projections of the thorax. The method produces a breathing signal
from which the phase can be derived and has been successfully applied for respi-
ratory correlated CB reconstruction [4]. Zeng *et al.* [8] proposed a method for 3D
motion estimation from a sequence of CB projections. A B-spline deformation
model is used to deform a reference CT volume to match the motion observed in
the CB projections. The applied deformations are optimized by computing the
similarity between the CB projection views of the deformed volume and the CB
projection sequence. Optimization of the numerous parameters of the B-spline
deformation model was regularized by introducing spatial and temporal motion
roughness penalties and an aperiodicity penalty for the estimated breathing mo-
tion. Results on simulated data were encouraging, demonstrating the feasibility
of the approach.

We propose a method for estimating 3D respiratory motion from a CB projec-
tion sequence by incorporating prior knowledge about the patient's respiratory
motion. The proposed method requires that a 4DCT image of the patient is
available at CB acquisition time. From this image the respiratory motion is es-
timated using deformable registration. The motion estimation is incorporated
into a patient specific motion model with two parameters: the breathing phase
and amplitude. For each CB projection, the 3D motion estimation comes down
to finding the model parameters for which the modeled CB projection view best
matches the CB projection with respect to a similarity measure. In the next sec-
tion the images used and the construction of the patient specific motion model
are discussed. Next the proposed cost function and optimization scheme are de-
tailed. Section 3 contains the experiments performed to test the proposed method
and the results for these experiments. Section 4 focuses on known limitations and
issues related to the chosen approach and outlines future work.

2   Materials and Methods

2.1   4DCT and Deformable Registration

The images used for this study are part of a freely available data set [9]. It
consists of a 4DCT image of the thorax composed of 10 3DCT images represent-
ing different phases of the respiratory cycle.The data set is completed by 400
anatomical landmarks (40 in each 3DCT image) identified by medical experts.

Each of the 3DCT images will be referred to as a *phase* of the 4DCT. Deformations between the phases were estimated by deformably registering them to one particular phase, arbitrarily chosen to be the end-inhalation phase. Registration was performed using the demons algorithm [10], of which its effectiveness on CT images of the thorax was verified in earlier work by our group [11]. The accuracy of the registrations was assessed using the landmarks provided with the data set. The phase to phase displacements of 40 anatomical points, as estimated by medical experts, were compared to the displacements given by the registration results. The average registration error was thus estimated to be 1.2mm with a standard deviation of 0.4mm. The maximum misalignment for these landmarks was found to be 2.6mm.

2.2   Patient Model

Using the registration results obtained above, we can construct a patient specific motion model. The proposed model is composed of a reference image I_{ref} and a deformation model T. I_{ref} should be *warped* using the deformation model T to obtain a modeled breathing state S. There are at least two ways of warping an image with a deformation field: you can obtain the target through either backward or forward mapping of the source. Suppose we have a source volume I and a target volume J. In addition suppose we have estimated the displacements of voxels of I to their corresponding positions in J and inversely. We can obtain J through backwards mapping of the source voxels:

$$J(\mathbf{x}) = I(\mathbf{x} + D_{J \to I}(\mathbf{x})) . \tag{1}$$

where $D_{J \to I}(\mathbf{x})$ represents the displacement necessary to map a voxel with position $\mathbf{x}$ in the target space to its corresponding position in the source space. J is thus obtained by fetching for each voxel position of J its corresponding (interpolated) value in I. The second possibility is through forward mapping:

$$J(\mathbf{x} + D_{I \to J}(\mathbf{x})) = I(\mathbf{x}) . \tag{2}$$

where this time $D_{I \to J}(\mathbf{x})$ represents the displacement necessary to map a voxel with position $\mathbf{x}$ in the source space to its corresponding position in the target space. J is obtained by adding for each voxel position of I, a contribution to the neighboring voxels of its corresponding position in J. If not explicitly taken into account, forward mapping can lead to holes in the target image, i.e. voxels for which no contribution was added. Backwards mapping is usually preferred as it allows for a more efficient implementation.

The reference image for the proposed model is the *mean-position image* (MPI), which we define as the image in which all structures appear at their time-weighted mean position. This concept was also used by Wolthaus *et al.* [12] for the definition of the *mid-ventilation phase*. Note that due to hysteresis of the respiratory motion, the mean position of a moving structure does not necessarily lie on the trajectory of that structure. It was obtained in the following way. For each voxel of the end-inhalation used as reference image for the registrations

described in Section 2.1, we calculated the mean of the deformation vectors mapping this voxel to its positions in the other phases (including a zero vector to account for the position of the voxel in the end-inhalation phase itself). The end-inhalation phase was then deformed using this mean deformation field, through backwards mapping. This requires the inversion of the mean deformation field. The Figure 1 shows a graphical representation of the MPI.

The second element in the patient model is the deformation model T allowing I_{ref} to be warped to a different breathing state. As we want to incorporate the phase to phase deformations observed in the 4DCT image, T should at least be able to reach these states. The deformations between the MPI and the phases of the 4DCT were again estimated through deformable registration. It was preferred to re-estimated the deformations to avoid accumulating errors by composing the previously obtained deformation fields. The resulting deformation fields were combined to form a 4D vector image, the fourth dimension being the breathing phase ϑ. This vector image was recursively filtered to obtain a 4D continuous cubic B-spline representation [13], denoted by $T(\mathbf{x}, \vartheta)$. In addition to modelling the phase to phase deformations, a second model parameter α was introduced to allow inter- and intracyle variations of the deformations. α can be interpreted as an instantaneous amplitude, linearly scaling the displacements of all voxels given by T for a certain value of ϑ.
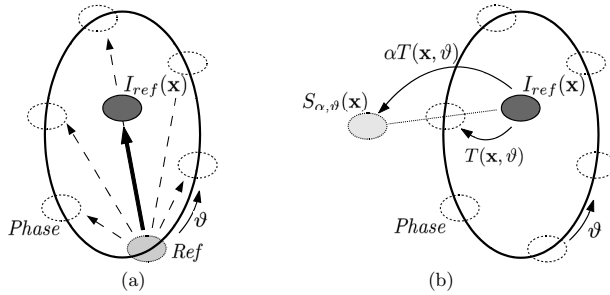
Fig. 1. (a) Schematic 2D representation of the construction of the MPI. A reference phase of the 4DCT (*Ref*) is registered to all other phases. For each voxel the estimated displacements are averaged (*bold arrow*) and after inversion used to obtain $I_{ref}(\mathbf{x})$ through backwards mapping of the reference phase. (b) Schematic 2D representation of the patient model. An anatomical point of I_{ref} is shown at its time-weighted mean position $\mathbf{x}$. Its corresponding position in the phases of the 4DCT was estimated through deformable registration. Interpolating these positions yields the closed contour represented in bold, which can be interpreted as an estimated trajectory. $T(\mathbf{x}, \vartheta)$ will map $\mathbf{x}$ on this trajectory, whereas $\alpha T(\mathbf{x}, \vartheta)$ can map $\mathbf{x}$ to any breathing state $S_{\vartheta, \alpha}(\mathbf{x})$ in the plane of the ellipse.

Backwards mapping I_{ref} using T would require to use the deformations defined
from the space of the phases of the 4DCT (the target volumes) to I_{ref} (the source
volume). In this case interpolating these deformation fields with a cubic spline
would be interpolating the end point of vectors which have different starting
positions. Similarly, when scaling the deformations by α, one should scale the
deformations defined from I_{ref} to the phases of the 4DCT. Tests were performed
performed with both backwards and forward mapping and showed very little
difference to the resulting target volumes and especially to their CB projection
view. This is due to the fact that the deformation fields vary quite smoothly in
space. In this work forward mapping was used when warping I_{ref} to S. With
$T(\mathbf{x}, \vartheta)$ representing the deformation given by the cubic spline for a phase ϑ on
position $\mathbf{x}$, we can represent a model breathing state S through forward mapping
as:

$$S(\mathbf{x} + \alpha T(\mathbf{x}, \vartheta)) = I_{ref}(\mathbf{x}) . \tag{3}$$

We will use $S_{\vartheta,\alpha}$ to represent the image that is the result of deforming I_{ref}
through forward mapping using the deformations as given by $\alpha T(\mathbf{x}, \vartheta)$. The
phase parameter $\vartheta \in [0, 1[$, goes through all breathing phases from the end-
exhale phase to end-inhale and back to end-exhale when rising from 0 to 1. For
the parameters at values $\alpha = 1$ and $\vartheta = 0, 0.1, 0.2, ...$ the modeled breathing
states correspond to the phases of the 4DCT. For $\alpha = 0$, I_{ref} is found. The
right panel of Figure 1 shows a schematic representation of the proposed patient
model.

2.3   Cost Function and Optimization

In order to compare the 3D modeled breathing state $S_{\vartheta,\alpha}$ to a CB projection p_ϕ
taken from a projection angle ϕ, we calculate its CB projection view. Let $\mathcal{A}_\phi$ de-
note the ideal CB projection operator from a projection angle ϕ. The projection
operator $\mathcal{A}_\phi$, was made to simulate the geometry of the Elekta Synergy. Figure 2
shows the modeled CB projection view of I_{ref}. For comparison we also show a
CB projection of the same patient acquired on the Elekta Synergy. The hori-
zontal object present in the CB projection is a reinforcement of the treatment
table. For each CB projection can define the optimization problem as follows:

$$(\hat{\vartheta}, \hat{\alpha})_\phi = arg \max_{\vartheta,\alpha}(\mathcal{F}(p_\phi, \mathcal{A}_\phi S_{\vartheta,\alpha})) . \tag{4}$$

where $\hat{\vartheta}$ and $\hat{\alpha}$ represent the estimated model parameters, $\mathcal{F}(., .)$ is a similarity
measure for which we assumed that higher values correspond to higher similarity.
The optimization is performed in the space of the model parameters with respect
to the similarity measure between the real projection and the modeled projection
view. We used mutual information as similarity measure [14]. The cost function
was handed to a Powell optimization strategy [15]. Each subsequent optimization
was initialized with the model parameters found for the previous CB projection.

(a)   (b)
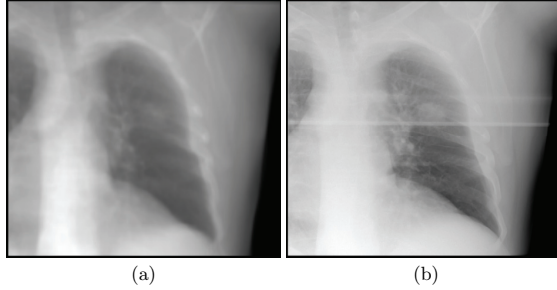
Fig. 2. (a) CB projection view of I_{ref} and (b) CB projection of the same patient acquired on the Elekta Synergy.

3   Experiments and Results

Validation of the motion estimation on real CB sequences is difficult as no ground truth is available. In this work experiments were performed on a simulated sequence of CB projections. We randomly generated a continuous phase signal with a varying period using the statistical properties reported for the breathing period by George *et al.* [16]. This resulted in a breathing period that varied between 2.2s and 5.6s. The phase was not limited to linear functions of time, allowing for intracycle breath rate variations. The continuous amplitude signal α was randomly generated using a lognormal random number generator, and was limited to slow variations in time and values close to one. The thus obtained amplitude signal varied between 0.88 and 1.21. These parameter values were fed to the patient model and the CB projection views were calculated every $0.6°$, twice per second and starting from the right lateral side of the patient. This is in good correspondence with actual CB acquisition parameters. We ran the proposed method for the first 150 projections, which amounts to 30 seconds of scanning time over an angle of $90°$. The optimization for the first projection was initialized with $\vartheta = 0$ and $\alpha = 1$ which was close to the generated values $\vartheta = 0.02$, $\alpha = 1.1$. It is our opinion that including a second parameter to account for the inter- and intracycle variations of amplitude extents the grasp of the method and makes the estimation of the phase more robust. To verify this we ran the optimization a second time on the same generated CB projection sequence but now fixing the amplitude of the patient model to one. The phase estimated using this method will be noted as $\check{\vartheta}$. Results obtained through this method will be referred to as estimated using a *phase based model* (PhM), as opposed to using the *phase-amplitude based model* (PhAM). The generated phase signal and amplitude signal of the simulated CB sequence and their estimated values are shown in Figure 3 in function of the projection number. The estimation errors are summarized for the whole CB sequence in Table 1. We also assessed the 3D misalignment for the landmarks discussed in Section 2.1.

Their 3D positions in the breathing states used to generate the CB sequence
are compared to their estimated positions. Figure 4 shows the misalignment in
function of the projection number while Table 2 summarizes these measures for
the whole CB sequence. Figure 5 shows an example of a generated CB projec-
tion for which difference images where calculated with the CB projection view of
I_{ref} and with the estimated CB projections. For this projection (with projection
number 138), estimation errors were around their average value. When running
the optimization with the PhAM, results are accurate for the entire sequence,
yielding acceptable maximum errors and excellent mean errors for parameter
value. Figure 4 shows that the maximum misalignment of up to 16mm was ef-
fectively reduced to below 1mm for this experiment. When using the PhM, the
phase estimate showed bad correspondence. With respect to the 3D misalign-
ment, the estimated sequence hardly improved the mean misalignment of the
landmarks with respect to before the optimization.

Table 1. Summary of the estimation errors for the parameter values. The table con-
tains the mean absolute estimation error (ME), its standard deviation (SD) and the
maximum absolute estimation error ($MaxE$). $\hat{\vartheta}$ and $\hat{\alpha}$ represent the estimates of the
model parameters when using a PhAM, $\check{\vartheta}$ is the phase estimate when PhM.

Parameters	ME	SD	MaxE
$\hat{\vartheta}$	0.0024	0.0265	0.0350
$\hat{\alpha}$	0.0040	0.0001	0.0605
$\check{\vartheta}$	0.1109	0.2074	0.4800

Table 2. The 3D position of the landmarks is compared to their position in: I_{ref}
(*before*), the estimated sequence when using a PhM and the estimated sequence when
using a PhAM. The table contains the mean of the misalignment ($MeanM$), its standard
deviation (SD) and the maximum misalignment($MaxM$).

Method	MeanM (mm)	SD (mm)	MaxE (mm)
before	3.8544	2.1760	16.7856
PhM	3.5033	1.7747	12.6477
PhAM	0.0184	0.0366	0.9596

4   Discussion and Conclusions

Motion estimation using a phase-amplitude based model performed well on this
simulated data. It was expected that the estimates when using a PhM would be
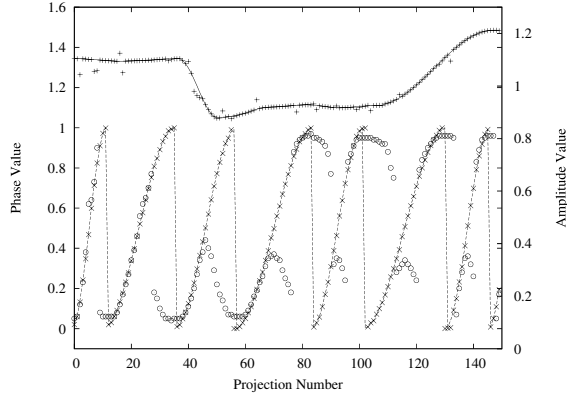
Fig. 3. The randomly generated amplitude signal α (*full line*) and phase signal ϑ (*dashed line*) used to generate the CB projection sequence, and the estimated amplitude $\hat{\alpha}$ (*plus marks*) and phase $\hat{\vartheta}$ (*cross marks*) when using a PhAM. The phase $\check{\vartheta}$ that was obtained using a PhM is shown in the same graph (*circles*).
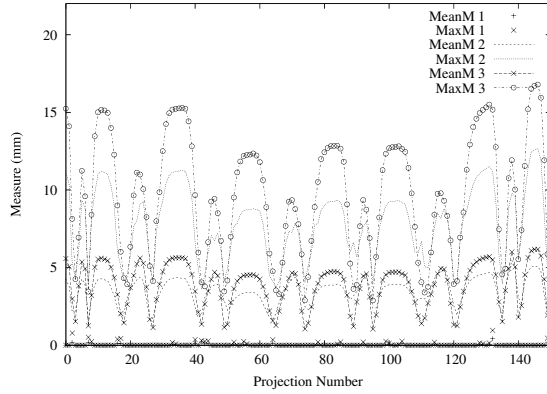
Fig. 4. The 3D position of the landmarks is compared to the their position in the estimated CB projection. The mean and maximum misalignment with the I_{ref} or before optimization (*MeanM 3, MaxM 3*), with the estimated sequence when using a PhM (*MeanM 2, MaxM 2*), with the estimated sequence when using a PhAM (*MeanM 1, MaxM 1*)
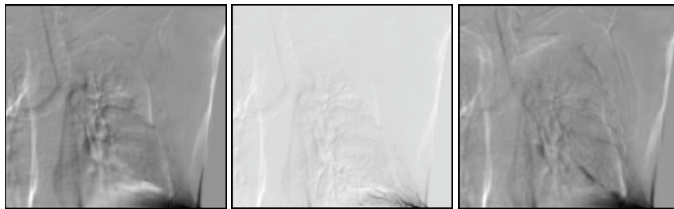
Fig. 5. Difference images between the generated CB projection and the: (*left*) CB projection view of I_{ref}, (*middle*) modeled CB projection view after optimization using the PhAM, (*right*) modeled CB projection view after optimization using the PhM. For the middle image the windowing was altered the emphasize the present differences.

less accurate since the CB sequence was generated using a PhAM. The phase estimates when using a PhM were however surprisingly bad. They indicate that the phase estimation is highly sensitive to even relatively small changes in amplitude. We are currently working on expanding the method for real CB projection sequences. One major issue is validation as it is hard to quantify the accuracy of the motion estimation in the absence of a ground truth. Some other remarks regarding real CB sequences can be made. In order to create CB projection views from the patient model, one must first be able to accurately reproduce the CB geometry and place the patient model in the same physical location as the patient. This can be done by rigidly registering the reconstructed 3D CBCT volume to for example a phase of the 4DCT. Although this CT-CBCT registration should not pose a problem, preliminary tests have shown that the method is highly sensitive to misalignment of the volumes, becoming unpredictable even for small misalignments. In the experiments performed, the similarity measure used showed reliable identification of the optimal parameters on the simulated data, and proved relatively easy to optimize. However several elements could undermine its efficiency when confronted with real CB data. For example, the borders of the treatment table will appear in a large part of the CB projections. Scatter and detector noise, not taken into account by the ideal projection operator, might turn the similarity measure less reliable and harder to optimize. Finally it is possible that changes of the patient anatomy and breathing motion over the course of the treatment are too large to remain adequately represented by the proposed patient model. An example of such changes are baseline shifts (variations of the mean tumor position), which have been reported for lung tumors [2]. The proposed patient model might not handle well large baseline shifts, but can in that case be expanded so that they are explicitly taken into account.

References

1. Keall, P.J., Mageras, G.S., Balter, J.M., Emery, R.S., Forster, K.M., Jiang, S.B., Kapatoes, J.M., Low, D.A., Murphy, M.J., Murray, B.R., Ramsey, C.R., van Herk,

M.B., Vedam, S.S., Wong, J.W., Yorke, E.: The management of respiratory motion in radiation oncology report of AAPM task group 76. Med Phys **33**(10) (Oct 2006) 3874–3900

2. Sonke, J.J., Lebesque, J., van Herk, M.: Variability of four-dimensional computed tomography patient models. Int J Radiat Oncol Biol Phys (Nov 2007)

3. Jaffray, D.A., Siewerdsen, J., Wong, J.W.: Flat-panel cone-beam computed tomography for image-guided radiation therapy. Int J Radiat Oncol Biol Phys **53**(5) (Aug 2003) 1337–1349

4. Sonke, J., Zijp, L., Remeijer, P., van Herk, M.: Respiratory correlated cone beam CT. Med Phys **32**(4) (Apr 2005) 1176–1186

5. Li, T., Koong, A., Xing, L.: Enhanced 4D cone-beam CT with inter-phase motion model. Med Phys **34**(9) (Sep 2007) 3688–3695

6. Rit, S., Wolthaus, J., van Herk, M., Sonke, J.J.: On-the-fly motion-compensated cone-beam CT using an a priori motion model. In: Medical Image Computing and Computer-Assisted Intervention (MICCAI). Volume 5241., New York, USA (2008) 729–736

7. Zijp, L., Sonke, J., van Herk, M.: Extraction of the respiratory signal from sequential thorax cone-beam X-ray images. In: International Conference on the Use of Computers in Radiation Therapy. (May 2004)

8. Zeng, R., Fessler, J.A., Balter, J.M.: Estimating 3-D respiratory motion from orbiting views by tomographic image registration. IEEE Trans Med Imaging **26**(2) (Feb 2007) 153–163

9. Vandemeulebroucke, J., Sarrut, D., Clarysse, P.: The POPI-model, a point-validated pixel-based breathing thorax model. In: XVth International Conference on the Use of Computers in Radiation Therapy, ICCR. (2007)

10. Thirion, J.P.: Image matching as a diffusion process: an analogy with maxwell's demons. Med Image Anal **2**(3) (Sep 1998) 243–260

11. Sarrut, D., Boldea, V., Ayadi, M., Badel, J., Ginestet, C., Clippe, S., Carrie, C.: Nonrigid registration method to assess reproducibility of breath-holding with ABC in lung cancer. Int J Radiat Oncol Biol Phys **61**(2) (Feb 2005) 594–607

12. Wolthaus, J.W.H., van Herk, M., Muller, S.H., Belderbos, J.S.A., Lebesque, J.V., de Bois, J.A., Rossi, M.M.G., Damen, E.M.F.: Fusion of respiration-correlated pet and ct scans: correlated lung tumour motion in anatomical and functional scans. Phys Med Biol **50**(7) (Apr 2005) 1569–1583

13. Unser, M.: Splines: A perfect fit for signal and image processing. IEEE Signal Processing Magazine **16**(6) (November 1999) 22–38 IEEE Signal Processing Society's 2000 magazine award.

14. Maes, F., Collignon, A., Vandermeulen, D., Marchal, G., Suetens, P.: Multimodality image registration by maximization of mutual information. IEEE Trans Med Imaging **16**(2) (Apr 1997) 187–198

15. Press, W.H., Flannery, B.P., Teukolsky, S.A., Vetterling., W.T.: Numerical Recipes in C. Cambridge University Press, second edition (1992)

16. George, R., Vedam, S.S., Chung, T.D., Ramakrishnan, V., Keall, P.J.: The application of the sinusoidal model to lung cancer patient respiratory motion. Med Phys **32**(9) (Sep 2005) 2850–2861

Automated Quantitative Assessment of Lung Fissure Integrity on CT

Matthew Brown, Robert Ochs, Fereidoun Abtin, Arash Ordookhani, Megan Brown, Hyun Kim, Greg Shaw, Daniel Chong, Jonathan Goldin

Department of Radiological Sciences, David Geffen School of Medicine at UCLA, Los Angeles, CA, USA

M Brown – mbrown@mednet.ucla.edu; R Ochs – rochs@ucla.edu; F Abtin – fabtin@mednet.ucla.edu; A Ordookhani – aordookhani@mednet.ucla.edu; M Brown – meganbrown@ mednet.ucla.edu; G Shaw – gshaw@ucla.edu; D Chong – dchong@mednet.ucla.edu; J Goldin – Jgoldin@mednet.ucla.edu

Abstract. This study investigates an unsupervised machine learning approach for quantitative analysis of pulmonary fissure completeness. The analysis of pulmonary fissures has traditionally been a subjective task, relying on visual assessment. Fissure analysis is becoming an important consideration as clinical studies suggest that fissural completeness significantly correlates with success of new endobronchial valve therapies for emphysema. Using an unsupervised clustering approach and 600 fissure samples, three clusters emerged: fissure, non-fissure, and bronchovascular bundle. The performance of the system was then evaluated using 84 randomly selected fissure images. The test cases were also independently contoured by two observers to form a reference standard for fissural completeness. Analysis of the results showed that there was no statistically significant difference between the CAD system and the human observers in calculating fissure percentage completeness (T-Test $P > 0.05$). Pair-wise comparisons of CAD-reader and reader-reader assessment of fissure completeness showed comparable levels of agreement $\geq 77\%$.

Keywords: CAD, Fissure, Emphysema,

1   Introduction

Quantitative Image Analysis (QIA) is an important component in the process of patient selection and treatment targeting in new emphysema treatments involving minimally invasive lung volume reduction. One of the new treatments utilizes one-way endobronchial valves placed in the segmental airways to exclude and deflate an emphysematous lobe without the need for surgery, thereby allowing other (healthier) lobes to further expand and improve lung function [1].

The degree of fissure integrity (i.e., completeness) is emerging as a potential predictor of treatment efficacy (deflation of emphysematous lobe) [2]. An incomplete pulmonary fissure, as shown in Figure 1, indicates the potential for collateral ventilation between adjacent lobes, circumventing complete lobar occlusion and lobar collapse [3,4]. Thus analysis of fissure completeness may play a key role in

identifying patients who would benefit from the therapy and determining the target
lobe for the endobronchial valve treatment.

   Determination of fissural completeness using HRCT imaging has previously been
carried out by visual inspection and subjective grading into three categories:
complete, incomplete, and absent [2,5,6]. This is a tedious and difficult task requiring
review of a large number of images under multi-planar reformats, which has
motivated the development of an automated quantitative system to assess fissure
integrity. A variety of methods have been introduced for fissure detection [7-13], but
none have been used to quantify fissural completeness or proven successful in an
emphysematous lung [14]. In emphysema patients, the lung parenchyma is
inconsistent in structure due to the presence of enlarged air-sacs known as bullae
which lead to fissures that are irregular and indistinct on imaging.

   An automated quantitative system could be of tremendous benefit in making
assessment of fissure integrity more discriminative, reproducible, and broadly
applicable. The aim of this pilot study is to demonstrate the feasibility of automated
quantitative assessment of fissure integrity from thin-section CT in emphysema
subjects with abnormal and incomplete pulmonary fissures.

2   Methods

2.1   Image Data Collection

The CT images used in this study were selected at random from a cohort of 486
emphysema subjects from a research database (see cohort demographics in Table 1).
Images were acquired with the following imaging parameters: 120 kVp, 140 to 300
mAs, and a pitch ranging from 0.984 to 1.5. Images were reconstructed with slice
thicknesses ranging from 1 to 3 mm and using standard reconstruction algorithms
(e.g., GE STD, Siemens B30f, Philips B, and Toshiba FC10 filters).

   For this pilot study we focused our initial analyses on the left major fissure. The
research database contains 12,391 technologist-drawn and radiologist approved 2D
left major fissure contours from 486 subjects. From these contours samples were
randomly selected to form independent training and test sets.

Table 1. Demographics of 486 emphysema subjects in a research database.

	Absolute (Mean ± SD)	Range
Age (y)	63 ± 7	$41 - 76$
PFT TLC (L)	7.62 ± 1.48	$3.4 - 12.14$
PFT RV (L)	4.83 ± 1.21	$0.86 - 9.22$
CT TLC (L)	6.99 ± 1.38	$3.98 - 10.45$
CT RV (L)	5.12 ± 1.26	$2.29 - 8.99$
Voxels below -910 HU (%)	56.53 ± 10.30	$30.45 - 81.27$

2.2   Image Features

The Hessian matrix of second order partial derivatives (Eq. 1) was used to characterize the variations in intensity about a point. In our initial implementation the Hessian matrix was computed in two-dimensions, however, the technique can be generalized to three-dimensions [13]. The computation has been described in greater detail by other authors [7,15]. The eigenvalues of the Hessian matrix, κ_1, κ_2, are computed and ranked according to their absolute value $|\kappa_1| > |\kappa_2|$.   Prior to computation of the Hessian, Gaussian smoothing was applied to the image with a kernel with standard deviation $\sigma = 1.0$ mm.

$$Hf(x,y) = \begin{bmatrix} f_{xx} & f_{yx} \\ f_{xy} & f_{yy} \end{bmatrix}.$$

(1)

Fissures are modeled as faint, plate-like structures (or lines in 2D images) due to their thin surface and partial volume averaging. κ_1 should correspond to the gradient change normal to the fissural line (along the first eigenvector). The other orthogonal eigenvalue (κ_2) should correspond to the second eigenvector in the direction of the fissure and should be closer to zero.   Based on this assumption we compute a "plateness image" using Equation 2 (see Figure 1).

$$I = 1 - \left| \frac{\kappa_2}{\kappa_1} \right|$$, where the ratio κ_2 / κ_1 is expected to be high for fissures.

(2)
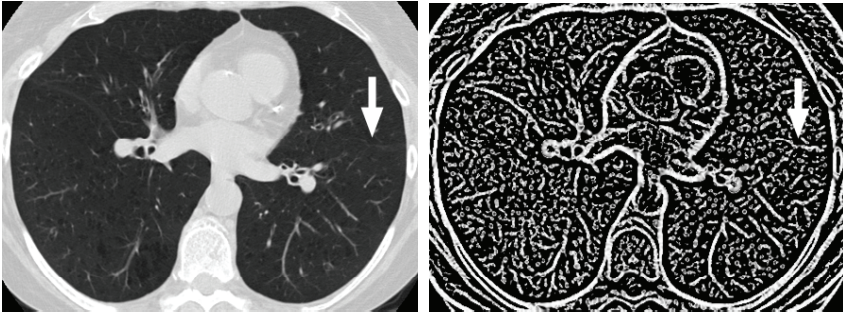
Figure 1. Original (on left) and plateness (on right) images with an incomplete left major fissure indicated by arrows. The fissures appear as faint white lines in the original image and solid lines in the plateness image. Since there are blood vessels with similar appearance the left major fissure is marked with an arrow.

Three features are computed for each voxel along the fissure path within a 32 x 32 patch around the voxel from the original CT and plateness images as shown in Table 2. The fissure path is based on a manually drawn contour as described in section 2.1.

Table 2. Features computed along fissure path in a 32 x 32 patch around the voxel of interest in the original CT and plateness (gradient-based) images.

Image	Feature	Range	Expected values
Original CT	Median gray level	0 - 255	higher if voxel is a fissure
Plateness	20th centile of histogram	0 - 255	higher if voxel is a fissure
Plateness	Max run length of voxels with intensity > 100	0 - 32	longer if voxel is a fissure

2.3 Classification Model

The classification model was machine-learned using unsupervised k-means clustering. The number of clusters was set as K=3 based on the premise that three classes are expected: fissure, non-fissure, bronchovascular bundle (near the hilum).

The classifier was trained using 600 feature samples from among the 12,391 fissure contours in our research database. Each sample (computed from a 32x32 patch) was taken from a different fissure contour (i.e., different CT image) to minimize dependence between the samples and bias in the classifier. Sample patches surrounding the learned cluster centers are shown in Figure 2.

Using the cluster centers, a minimum-distance classifier was implemented to classify each voxel along the fissural path. For each voxel along the path, the three features are computed in the 32x32 neighborhood and input to the classifier. The fissural completeness is then calculated as the percentage of pixels along the path classified as fissure.

          (a)   (b)   (c)
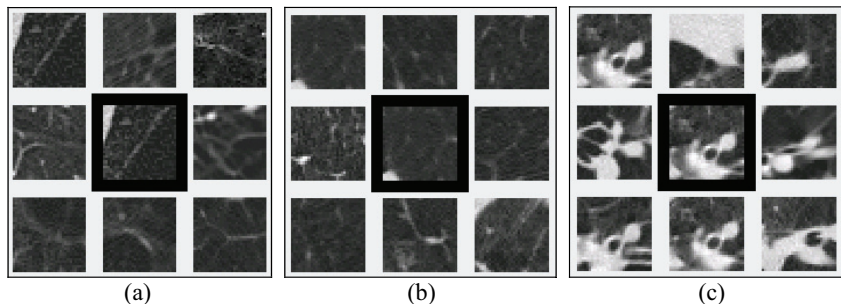
Figure 2. Nine training samples nearest to the learned cluster centers for (a) fissures, (b) non-fissures, and (c) bronchovascular bundles.

2.4   Experimental Testing

100 HRCT slices were selected at random from our research database. Two human observers independently reviewed all 100 slices. They were shown the endpoints of the fissure marked initially by the technologist (which was also used as the input fissure path to the CAD system). If either one of the observers in the current

experiment disagreed with the endpoints of these previous contours the case was rejected. 16 cases were rejected during this process, resulting in 84 CT images which the observers and CAD system independently analyzed.

The human observers were trained on lung anatomy, software application, and sensitivity calibration. We did not use radiologists for this pilot study due to the previous determination that well-trained observers produce similar inter-observer variability as licensed radiologists [16].

The observers assessed the entire fissural path within each CT image, they marked portions of the fissure they considered to be intact and left blank the parts of the fissure they considered to be absent. They also marked portions of the path passing through vessel bundles (near the hilum) which were excluded from completeness calculations. The CAD system was also applied to each case and automatically classified voxels along the fissural path into the same three categories. For each observer the fissure integrity was then assessed by calculating the percent-completeness of the fissure:

$$Completeness = \frac{N_{Fiss}}{N_{Fiss} + N_{NonFiss}} \times 100 \ . \qquad (3)$$

where N_{Fiss} = number of voxels classified as fissure and $N_{NonFiss}$ = number of voxels classified as non-fissure.

The completeness percentage values were then compared pair-wise between Reader 1, Reader 2, and CAD. The percentages were also converted into a binary decision of complete or incomplete for each observation. Based on previous studies involving visual assessment, a completeness percentage of $\geq 90\%$ was considered as a complete fissure for potential treatment planning [16]. Bland-Altman plots (Figure 5) were used to quantify statistically significant agreements between treatments. Paired T tests were also performed to further analyze the agreement between treatments (See Section 3).

3. Results

Figures 4a-c are scatter plots showing the pair-wise comparisons between Reader 1, Reader 2, and CAD. Figure 4a compares the two readers, and Figures 4b and 4c compare CAD against each of the readers. Dashed lines indicate 90% fissural completeness. Points above/right of these lines (in the upper right quadrant) represent agreement on complete fissures, while points in the lower left represent agreement on incomplete fissures. The upper left and lower right quadrants represent disagreement between the observers. The numbers, N, shown in each quadrant are the number of pair-wise observations in each category. For example, in Figure 4a there were 55 fissures where both readers agreed that the fissure was complete, 18 where they agreed it was incomplete and 6+5=11 where they disagreed. Similar agreement/disagreement counts are shown for Figures 4b and 4c involving CAD vs Reader.

Bland-Altman plots were also used to determine the pair-wise agreement between fissure completeness scores. All three plots contain a cluster of points with zero difference at 100% completeness. In each of the Bland-Altman plots, 6 readings were significantly different (with 95% confidence), leaving a 78-reading agreement for each comparison. The paired T tests resulted in p values of 0.3401, 0.0863, and 0.3876 for the comparisons between Reader 1 and Reader 2, Reader 1 and CAD, and Reader 2 and CAD respectively.

Table 3 summarizes the agreement between the three observers. The first row indicates that all three agree on completeness/incompleteness in 72.6% of the fissures. In 14.3% of the fissure the two readers agree but CAD does not. In the remaining cases CAD agrees with one of the readers. Figure 3 shows examples of fissures with different levels of fissure completeness and observer agreement.

Table 3. Agreement counts between Reader 1, Reader 2, and the CAD system.

Agreement On Fissural Completeness			
Readers that agree	≥90%	<90%	Total
R1, R2, and CAD	49	12	61 (72.6%)
Only R1 and R2	6	6	12 (14.3%)
Only R1 and CAD	5	2	7 (8.3%)
Only R2 and CAD	3	1	4 (4.8%)

(a)   (b)   (c)
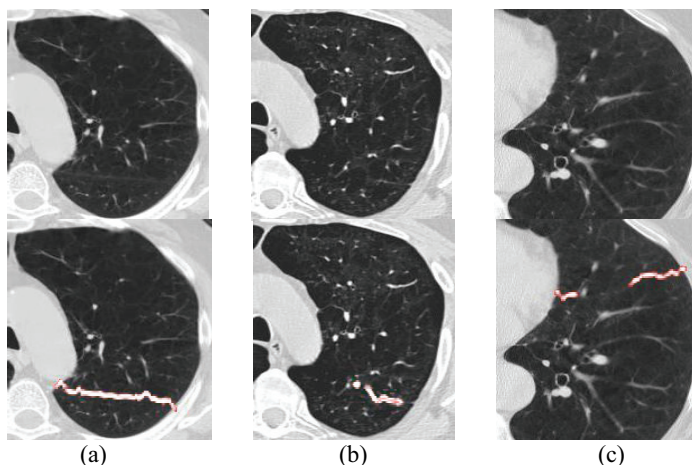
Figure 3. Original CT images and CAD detected fissures for cases where (a) both readers and CAD classified the fissure as complete, (b) both readers and CAD classified the fissure as incomplete, and (c) one reader and CAD classified it as incomplete and the other reader classified it as complete.
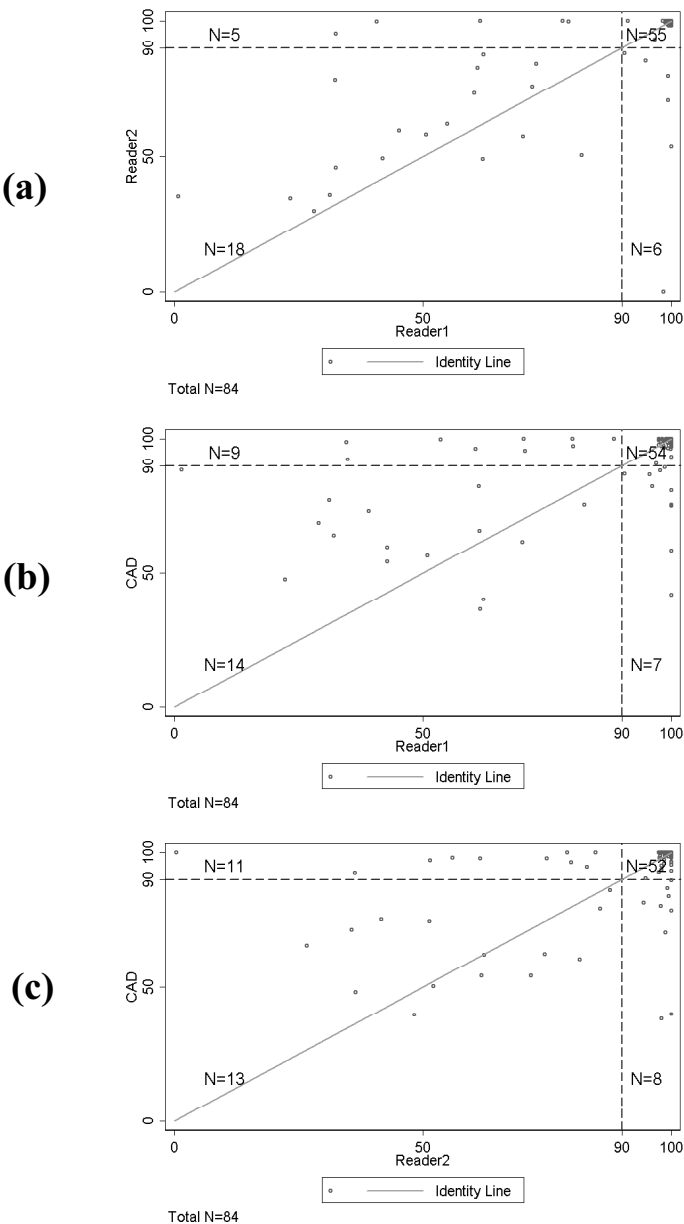
Figure 4. Scatter plot of fissure percentage completeness for (a) Reader 2 vs Reader 1, (b) CAD vs Reader 1, and (c) CAD vs Reader 2.
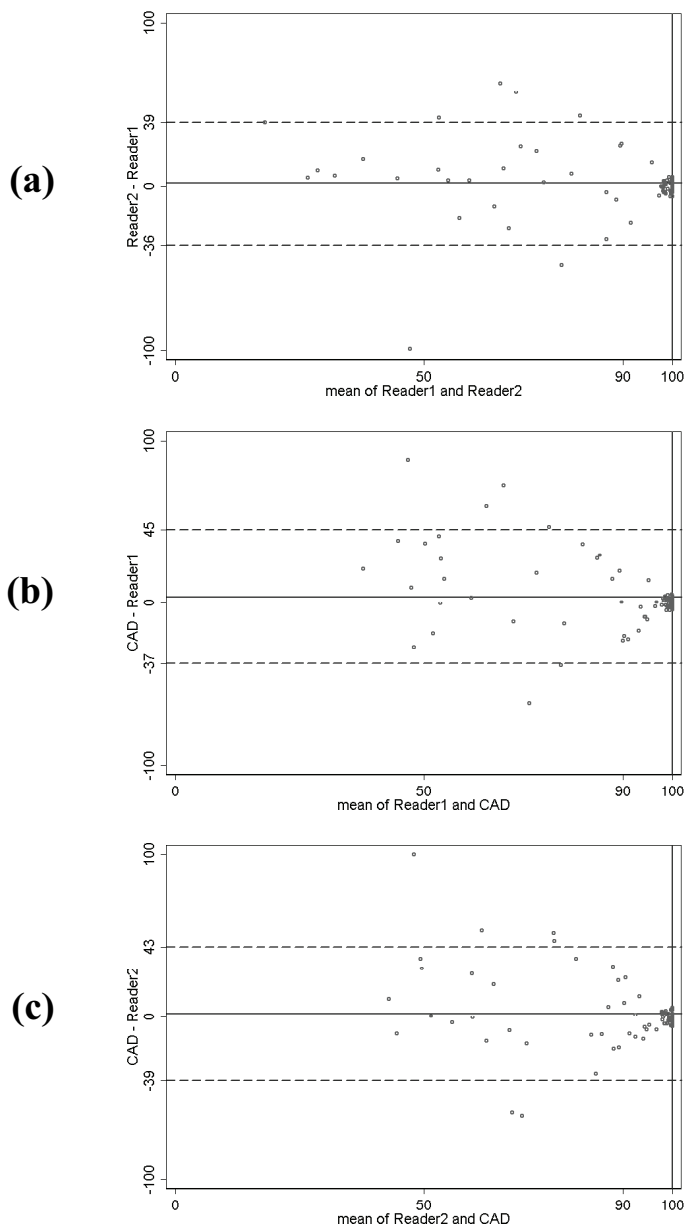
Figure 5. Bland-Altman plot of fissure completeness percentages for (a) Readers 1 and 2, (b) CAD and Reader 1, and (c) CAD and Reader 2.

4   Discussion

There was good pair-wise agreement for the percentage of fissural completeness, with no statistically significant differences within any of the three comparisons. When applying the 90% threshold for completeness CAD had good agreement with each of the readers (81% and 77% respectively - see Figures 4b,c), although the agreement between the readers themselves was higher (88% - see Figure 4a). This performance is particularly encouraging since it involved subjects with emphysema and incomplete fissures, rather than clearly defined normal fissures as has typically been the case in previous work on fissure detection. The difficulty of the fissure classification problem in emphysema subjects is indicated by the agreement levels between the readers. The scatter plot in Figure 4a shows that even trained readers have quite different completeness scores in some (difficult) cases.

We developed our prototype system for the left major fissure since its appearance seems more consistent than the right major or minor fissures. The completed system will include fissures from the right lung which may necessitate a larger sample size for classifier training. In this pilot study we used images with a range of CT technical factors, and further systematic investigation on their influence on features is needed. We will also continue investigation of supervised machine-learning approaches with expert-labeled voxels [13].

In future studies we expect improved agreement. For this pilot study we intentionally focused on a small number of features that were selected a priori and a simple classification method. Thus we find the results to be very encouraging and expect the approach to generalize well to larger data sets. Also, we were using only one CT image per assessment, both in terms of CAD feature calculation and for reader review. This was done for speed/simplicity of calculation and observer review in this initial investigation of the feasibility of using CAD to assess abnormal fissures. Previous fissure detection approaches have used multiple slices (three dimensions) and lung vasculature data [7-13] to guide fissure detection. We will extend the features to 3D and the reader evaluation to multiple slices and expect to achieve better performance and greater agreement. We will also investigate the feature space further with a labeled training set to determine whether a non-linear classifier is appropriate.

By maintaining a high accuracy for detection of complete and almost complete fissures we aim to maximize the number of successful outcomes for the endobronchial valve treatment [1,2].

5   Conclusion

The CAD prototype system showed good agreement with human observers in computing the percentage of fissure completeness and in classifying the fissures as complete or incomplete. The pair-wise agreement between CAD and each reader was comparable to that between the readers themselves. This work is one of the first efforts to specifically detect abnormal/incomplete fissures in subjects with emphysema and has important clinical applications in targeting of endobronchial valve treatments. The results are very encouraging for a challenging problem and we

expect the extension of the features to three-dimensions will yield further improvements in performance.

References

1. Strange, C., Herth, F., Kovitz, K.L., McLennan, G., Ernst, A., Goldin, J., Noppen, M., Criner, G.J., Sciurba, F.C.: Design of the Endobronchial Valve for Emphysema Palliation Trial (VENT): a non-surgical method of lung volume reduction. BMC Pulm Med 2007; 7(1):10 (2007)

2. Abtin, F., Goldin, G.G., Strange, C., Criner, G.J., Marquette, C., Sciurba, F.C., Brown, M.S., Rao, A., Kim, H., Roback, D., Jiang, A., Pais, R., Irani, Z., Herth, F.J., Ernst, A., Kovitz, F.L., McLennan, G.: The Influence of Fissural Anatomy on the Treatment Outcome of Patients with Emphysema. Posterboard presentation #34;A755 at ATS at Toronto (2008)

3. Scanlon, T.S., Benumof, J.L.: Demonstration of interlobar collateral ventilation. J Appl Physiol 46:658--661 (1979)

4. Hogg, J.C., Macklem, PT, Thurlbeck, W.M.: The resistance of collateral channels in excised human lungs. J Clin Invest. 48:421--431 (1969)

5. Proto AV, Ball JB Jr. Computed tomography of the major and minor fissures. Am J Roentgenol 140:439–448 (1983)

6. Ariyurek, O.M., Karabulut, N., Yelgec, N.S., Gulsun, M.: Anatomy of the minor fissure: assessment with high-resolution CT and classification. Eur Radiol. 12:175-180 (2002)

7. Sato, Y., Nakajima, S., Shiraga, N., Atsumi, H., Yoshida, S., Koller, T., Gerig, G., and Kikinis, R.: Three-dimensional multi-scale line filter for segmentation and visualization of curvilinear structures in medical images. Med. Image Anal. 2(2):143--168 (1998)

8. Zhang L., Hoffman, E.A., Reinhardt, J.M.: Lung lobe segmentation by graph search with 3-D shape constraints, Proc. SPIE. 4321:204--215 (2001)

9. L. Zhang, et al., Atlas-driven lung lobe segmentation in volumetric X-ray CT images, Proc. SPIE. 5031:306--315 (2003)

10. Wang, J., Betke, M., Ko, J.P., Shape-based curve growing model and adaptive regularization for pulmonary fissure segmentation in CT, in Lecture Notes in Computer Science. Berlin, Germany: Springer-Verlag, MICCAI. 3216: 541-548 (2003)

11. Saita S, Yasutomo M, Kubo M, Kawata Y, Niki N, Eguchi K, Ohmatsu H, Kakinuma R, Kaneko M, Kusumoto M, Moriyama N, Sasagawa M. An extraction algorithm of pulmonary fissures from mutli-slice CT image. Proc SPIE. 5370:1590--1597 (2004)

12. van Rikxoort EM, van Ginneken B, Klik M, Prokop M. Supervised enhancement filters: application to fissure detection in chest CT scans. IEEE Trans Med Imaging. 2008 Jan;27(1):1-10.

13. Ochs R., Goldin J., Abtin F., Kim H.J., Brown K., Batra P., Roback D., McNitt-Gray M., M. Brown. Automated classification of lung bronchovascular anatomy in CT using AdaBoost. Medical Image Analysis Volume 11, Issue 3, June 2007, Pages 315-324

14. Sluimer, I., Schilham, A., Prokop, M., van Ginneken, B.: Computer analysis of computed tomography scans of the lung: a survey. IEEE Trans. Med. Imaging 4, 385--405 (2006)

15. Krissian, K., Malandain, G., Ayache, N., Vaillant, R., Trousset, Y. Model Based Detection of Tubular Structures in 3D Images. Computer Vision Image Understanding 80, 130-171 (2000)

16. Abtin, F., Goldin, G.G., Brown, M., Rao, A., Kim, H., Jiang, A., Ahmad, S., Roback, D., Pais, R., Irani, Z.: Variation in Fiossure Anatomy in Emphysema Patients Treated with Endobronchial Valve. Posterboard presentation #33;A754 at ATS at Toronto (2008)

Optimal Graph-Based Segmentation of 3D Pulmonary Airway and Vascular Trees Across Bifurcations⋆

Xiaomin Liu[1], Danny Z. Chen[1], Xiaodong Wu[2], and Milan Sonka[2]

[1] Department of Computer Science & Engineering, University of Notre Dame,
xliu9@nd.edu
[2] Department of Electrical & Computer Engineering and Department of Radiation Oncology, The University of Iowa

Abstract. Segmenting airway and vascular trees in CT volume images plays a fundamental role in pulmonary image analysis. However, accurate assessment of complete tree morphology is difficult due to their complex nature. In this paper, we extend an optimal graph search based technique to identifying tree-structured airways and lung vessels with one or more interrelated surfaces. Based on a pre-segmentation that preserves the object topologies, our approach utilizes the 3D medial axes to re-sample the volume image and construct a geometric graph. By designing appropriate cost functions, the segmentation of both airways and vessels is performed across tree bifurcations in a single optimization process for the entire tree. Segmentation results of double surfaces for airways and single surface for vascular trees are presented.

1 Introduction

The airway and blood vessels are two major components of the human lung. Quantitative assessment of both the airway and vascular trees provides important information for functional understanding of pulmonary anatomy and objective measures of lung diseases. Due to the large image sizes and highly branching structure, it is tedious and time-consuming to manually locate individual tree branches and draw their contours on 2D slices. Furthermore, 2D manual tracing and image analysis methods may not be as effective and reliable, since they lack the ability to incorporate 3D information. Hence, developing automated and accurate 3D segmentation methods for lung images is a critical task in pulmonary image analysis and computer aided diagnosis. Further, airway trees and vessel trees may have multiple interrelated layers of surfaces for segmentation, some of which are extremely hard to detect individually. Several techniques have been proposed to segment the tubular lung objects ([1, 2]), but they cannot guarantee global optimality.
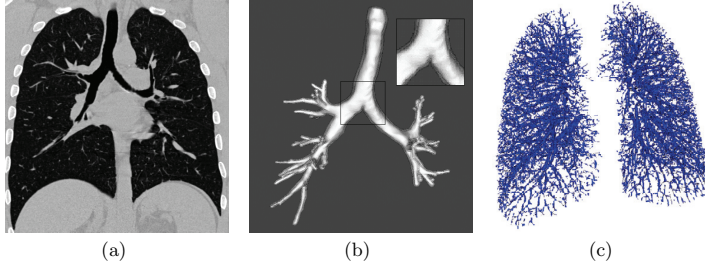
Fig. 1. Illustrating 3D airway and vessel segmentation: (a) The original lung CT image; (b) the rendered airway inner and outer walls with cross bifurcation segmentation; (c) the rendered vascular tree.

Wu and Chen [3] reported an optimal graph search based algorithm, which was extended to multiple surface segmentation by Li *et al.* [4]. These graph search based schemes transform the image segmentation problem to computing a minimum-cost closed set in a derived vertex-weighted graph, and obtain optimal segmentation. The methods have been successfully applied to non-branching airway segmentation [4] and MR arterial wall segmentation [4], but it did not directly extend to segmenting objects with a tubular and tree-like topology such as airways or blood vessels (see Fig. 1). The methods in [3, 4] are only suitable for the objects that have a relatively simple topology (e.g., cylindrical or spherical). In these cases, a 3D geometric graph can be built either by unfolding the sought surfaces to terrain-like surfaces or resampling the image along the normal surface directions within a narrow band. However, these "simple" methods for building 3D graphs are not directly applicable to objects with complicated structures since they may cause severe interferences among the resampled vectors and may fail to intersect (or capture) the sought surfaces.

In this paper, we present a technique for segmenting multiple interrelated layers of surfaces for airway and vessel trees. Specifically, to address the above-mentioned drawbacks, we propose a new scheme for building the 3D graphs for segmenting tubular and tree-structured objects. To overcome the difficulty of segmenting branching structures, we use medial axes and surface dilation to guide and produce an effective image resampling. Our graph search on the resampled images uses task-specific cost functions for airway and vascular trees. Consequently, we obtain segmentation results of multiple interrelated layers of surfaces for airway and vessel trees with significantly improved quality.

2 Method

The graph search based algorithms [3, 4] solve the segmentation problem by transforming it to finding a minimum-cost closed set in a directed vertex-weighted

graph, which is solvable in polynomial time. Our graph search based image segmentation approach consists of the following four major steps:

(1) Pre-segmentation. A pre-segmentation is needed to provide the basic information about the object's global structure. It is not necessary for the pre-segmentation to be locally accurate. However, it is crucial to preserve the topology of the target object. If the pre-segmentation does not yield a mesh, we also need to transform the volumetric result into a mesh representation.

(2) Image resampling. Based on the outcome of the pre-segmentation, the image is resampled on each vertex of the initial surface mesh, resulting in a set of vectors (called *columns*) of voxels. In this paper, the medial axes are applied to determine the directions and lengths of the resampling columns.

(3) Graph construction. Each voxel in the columns is considered as a node in the graph. There are three types of edges, representing the relations of voxels within the same surface or between different surfaces. A cost is assigned to each node which reflects the certain property of the sought surfaces.

(4) Graph search. Finally, we apply a minimum s-t cut algorithm [3, 4] to the resulting graph to simultaneously search for multiple interrelated surfaces.

For segmenting tubular and tree-structured objects, the most nontrivial task is to build a vertex-weighted geometric graph to model the volumetric image. When constructing this graph model, we need to carefully resample the volumetric image so that the following two constraints are satisfied: 1) All the sought surfaces must be captured by the graph; 2) the relations among the voxel columns should be consistent with the surface topology specified by the preliminary mesh from the pre-segmentation, meaning that interferences among different voxel columns cannot be allowed (more on this later).

2.1 Pre-segmentation

The algorithm we use for the pre-segmentation of pulmonary vascular trees is based on a hybrid method of the tube enhancement filtering and traversal approaches [5]. First, the tube enhancement based on the cylindrical shape model using an eigenvalue of the Hessian matrix serves as a filter to extract vessels. Then, a traversal step detects the change of signs of those eigenvalues to improve the vessel's connectivity. Finally, objects with many branch points are selected to distinguish between vascular trees and noise components. Airway trees are pre-segmented using commercially available Pulmonary Workstation PW+ software (VIDA Diagnostics, Oakdale, IA). Once a labeled image is generated by the pre-segmentation, it is transformed into a triangulated mesh using the marching cube algorithm.

2.2 Image resampling based on medial axes

To segment an optimal surface in the image using the corresponding preliminary meshed surface, our approach needs to perform a resampling of the image for
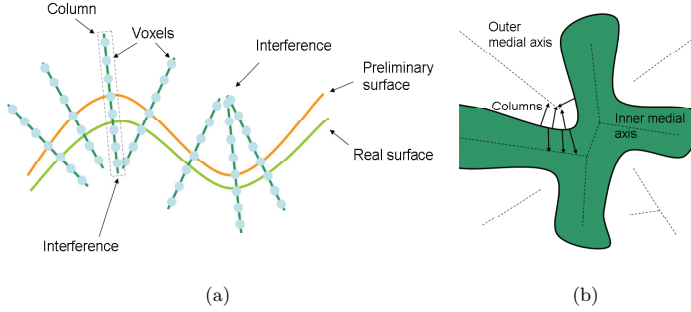
Fig. 2. Illustrating the image resampling: (a) The interferences caused by inappropriate column lengths, (b) a 2D example of image resampling based on medial axes.

every surface vertex along the normal direction of the meshed surface at that vertex, resulting in a column of voxels for each vertex. In this process, we seek to avoid two "bad" situations: (1) The length of a column is too long, so that it interferes with (intersects) other columns; (2) the length of a column is too short, so that it fails to capture enough information about the real surface. To avoid possible interferences among the resampled columns, we need to determine the proper directions and lengths for the columns. Intuitively, the normal direction at each mesh vertex is the direction without any bias or prior information about the location of the real surface.

A medial axis of the preliminary surface is a set of points each of which has at least two nearest points on the surface [6]. At each mesh vertex, the medial axis determines the maximum distance that a column can be extended along the normal direction without any interference with other columns. Although the exact computation of the medial axis is possible in principle, it is complicated to implement due to significant algebraic difficulties [6] – approximate solution can be obtained using computational geometry. Suppose the vertices of the mesh form a point set S. We can compute the Voronoi diagram and the dual Delaunay triangulation of S in 3D [7]. The medial axis is computed by using the poles in the Delaunay triangulation, which are selected from the centers of the big Delaunay balls adjacent to vertices. Then we assign a pole to each of the mesh vertices by selecting the largest pole among the vertex's k-nearest neighbors, in order to reduce the impact of possible noise on the surface [7]. Next, the lengths of the columns are obtained by computing the distances from the mesh vertices to their corresponding medial axis points (on both the inner and outer medial axes of the preliminary surface, Fig. 2).

A sought surface may contain very sharp angles at its branches. In such situations, a medial axis could be very close to the branching portions of the surface and consequently the columns computed based on the medial axis could

be quite short, giving very little flexibility to the graph search algorithm. To avoid this problem, we first grow (dilate) the preliminary surface by a certain distance, and then compute the columns based on the grown surface and its medial axes. The distance we used in the experiments is 3 voxels, which should be depending on the size of the object. With an appropriate value for the growth distance, we can obtain a set of columns that extend across the sought surfaces.

2.3 Graph search based segmentation

A graph $G = (V, E)$ has a set V of nodes (or vertices) that are connected by edges in E. In this paper, the nodes in V of G are voxels in the resampled volumetric image, which are organized by the columns. Each column of nodes is associated with a vertex of the preliminary mesh, and is sampled along the normal direction at that vertex. We assign edges to connect neighboring nodes in G and to ensure the geometric constraints (e.g., the *smoothness constraint* and the *separation constraints* of the surfaces [4]). Generally, the graph G is constructed in a similar manner as in [4].

Airway tree segmentation The detection of airway outer wall is difficult since the outer surface is often surrounded by other adjacent tissue with similar gray scale intensities in CT images. Instead of segmenting the airways section by section and gluing the branches together afterwards, we consider the airway tree as a whole and search for both the optimal inner and outer walls simultaneously.

For each sought surface, we construct a graph that is designed to contain that surface. At each mesh vertex v, there are two columns, corresponding to the inner and outer surfaces, respectively. Denote these two columns at vertex v by $Col_i(v) = \{n_i(v, 0), \ldots, n_i(v, K-1)\}$ and $Col_o(v) = \{n_o(v, 0), \ldots, n_o(v, K-1)\}$. Within every column, say, $Col_i(v)$, each node $n_i(v, k)$ is connected by a directed edge to $n_i(v, k-1)$ for $k \geq 1$. Between each pair of adjacent columns, a set of edges is assigned to ensure the *smoothness constraint* [4] within the surface. Let v_1 and v_2 be two adjacent vertices on the mesh, and suppose $n_i(v_1, k_1)$ is connected to $n_i(v_2, k_2)$ by an inter-column edge. Then the smoothness constraint Δ enforces that:

$$-\Delta \leq k_1 - k_2 \leq \Delta \qquad (1)$$

With the smoothness constraint, we avoid any dramatic change of the neighboring voxels on the same surface, which consequently results in smooth surfaces of the sought medical objects. In the case of double surface detection, another set of edges, called inter-surface edges, is added to impose the *surface separation constraint* [4] between the two surfaces. The inter-surface edges are assigned between vertices $n_i(v, k_i)$ and $n_o(v, k_o)$ for all $v \in V$ so that the following separation constraint is satisfied:

$$\delta^l \leq k_i - k_o \leq \delta^u \qquad (2)$$
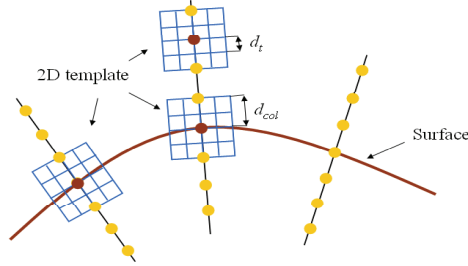
Fig. 3. Illustrating the cost function constructed at each node of the graph.

where δ^l (resp., δ^u) specifies the smallest (resp., largest) allowed distance between the two sought surfaces. (Assume that $\delta^l, \delta^u > 0$.) The separation constraint ensures that two sought surfaces are not too far away and also not too close to each other (e.g., they may not be allowed to intersect).

After the graph is constructed with the above three types of edges between the nodes in the columns, we need to assign a cost function to each of these nodes. The cost function must reflect the possibility for a voxel (node) to belong to a certain surface. For airway wall detection, the two surfaces differ from each other in the direction of intensity change. Since the airway lumen is darker than the airway wall, the intensity increases from low to high at the inner border. Conversely, the intensity decreases from high to low at the outer border when only parenchymal tissue is adjacent. However, this intensity change for the outer border shall also hold when non-parenchymal surrounding tissue is present. Since the airway wall is not completely connected to the surrounding tissue, there ought to be a little gap outside the wall that represents a "lower" intensity.

The cost function we use for airway segmentation is a combination of the first and second derivative edge detectors and is based on the cost function proposed in [8]. This is due to the property that the two edge detectors tend to yield the maximum magnitude on one or the other side of the true edge, causing certain over-estimate or under-estimate of the airway wall position. Thus, a weighted sum of the first and second derivatives works better for the accurate border location. Ideally, a 3D edge detector should be applied in our situation because of the 3D nature of the image. However, a test of these two edge detectors shows that the 3D edge detector performs similarly as the 2D one but with a significant increase of the computational overhead. In order to achieve a higher efficiency, the 2D edge detector is adopted to build the cost function at each node v

$$Cost_{outer}(v) = \omega \cdot I_t(v) * M_{Sobel_1} + (1 - \omega) \cdot I_t(v) * M_{Marr} \qquad (3)$$

$$Cost_{inner}(v) = \omega \cdot I_t(v) * M_{Sobel_2} + (1 - \omega) \cdot I_t(v) * M_{Marr} \qquad (4)$$

where M_{Sobel_1} and M_{Sobel_2} represent two inversely oriented 2D Sobel templates, and M_{Marr} represents the 2D Marr template (all in 5×5 size). $I_t(v)$ is a 5×5 gray-scale image template resampled from the 3D image around the node v. As we did for the resampling of columns, the template is also sampled along the corresponding vertex's normal direction (as illustrated in Fig. 3). Thus, the 2D image template is centered at the node v and is on a plane perpendicular to the surface. The voxels are sampled at the interval size d_t such that

$$d_t = d_{col}/2 \tag{5}$$

where d_{col} is the interval between neighboring voxels in a resampled column.

Vascular trees segmentation For the vascular tree segmentation, the graph is constructed in the similar manner although there is only one surface to be detected (the only one visible on CT image). Hence, only the intra-column and inter-column edges are needed to build the graph. The cost function used for vascular tree segmentation is the magnitude of gradient computed from the Gaussian smoothed images.

3 Validation of Pulmonary Image Segmentation

The validation of airway and vascular tree segmentation has been performed on a set of 6 CT scans of lung CT images. The sizes of the images vary from $512 \times 512 \times 562$ to $512 \times 512 \times 671$, voxel size $0.68 \times 0.68 \times 0.6\text{mm}^3$, Siemens Sensation 64-slice MDCT. After an initial surface of the airway inner wall is extracted (Section 2.1), our graph search approach succeeded in extracting the outer wall as well as optimizing the location of the inner wall in all 6 cases. An example of the cross bifurcation segmentation is shown in Fig. 4. To segment the vascular trees, each connected blood vessel component is filtered out and labeled with a different number. In comparison with the pre-segmentation results (which were considered final till now), the graph search captures the wall surfaces more accurately especially across bifurcations. (See Fig. 5.) While the initial results show the approach is promising for simultaneously segmenting single- and multiple-surfaces of pulmonary airway and vascular trees, we understand that more quantitative evaluation needs to be done in our following work.

4 Discussion and conclusion

In this paper, we extended the optimal graph search based approaches introduced in [3, 4] to segmenting airway and vascular trees in 3D pulmonary CT images. By using medial axes to guide the resampling, the 3D image is sampled properly based on the preliminary segmentation. Even though both the airways and vessels have complicated tree structures (airways with multiple interrelated surfaces), our proposed resampling scheme is able to extract sufficient information from the image data and provide it to the graph search algorithm to identify the
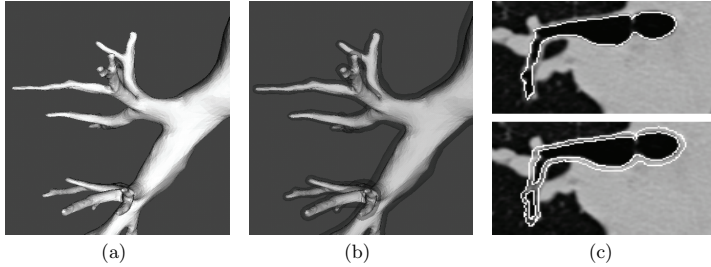
Fig. 4. Airway segmentation: (a) The inner wall (result of pre-segmentation); (b) double surfaces after graph search; (c) the comparison of preliminary result and graph search result in 2D slices. Note that the inner as well as outer surfaces are smooth and three-dimensionally accurate across the airway branching.
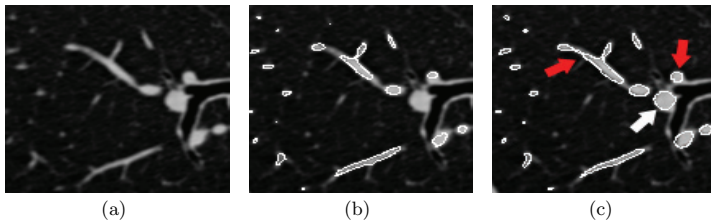
Fig. 5. Vascular tree segmentation results: (a) The original image; (b) the presegmentation; (c) the graph search result showing improved delineation (red arrows). Note the ability to correctly detect the surface even if the preliminary segmentation fails locally – as long as the preliminary segmentation is in the vicinity of the desired surface (white arrow).

optimal surfaces. By applying cost functions with directional information, our algorithm succeeds in detecting both the inner and outer surfaces of the airway walls as well as the vascular wall surfaces across bifurcations. As shown by the examples, the segmentation considerably improved the results of the preliminary segmentation.

References

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

2. Boldak, C., Rolland, Y., Toumoulin, C.: An improved model-based vessel track-
ing algorithm with application to computed tomography angiography. Journal of
Biocybernetics and Biomedical Engineering **23** (2003) 41–63
3. Wu, X., Chen, D.Z.: Optimal net surface problems with applications. In: Widmayer,
P., Triguero, F., Morales, R., Hennessy, M., Eidenbenz, S., Conejo, R. (eds.) ICALP
2002. LNCS. (2002) 1029–1042
4. Li, K., Wu, X., Chen, D.Z., Sonka, M.: Optimal surface segmentation in volumetric
images - A graph-theoretic approach. IEEE Transactions on Pattern Analysis and
Machine Intelligence **28** (2006) 119–134
5. Shikata, H., Hoffman, E.A., Sonka, M.: Automated segmentation of pulmonary
vascular tree from 3D CT images. In Amini, A.A., Manduca, A., eds.: Proc. of SPIE
Medical Imaging 2004: Physiology, Function, and Structure from Medical Images.
Volume 5369. (2004) 107–116
6. Attali, D., Boissonnat, J.D., Edelsbrunner, H.: Stability and computation of medial
axes: A state of the art report. In T. Möller, B.H., Russell, B., eds.: Mathemati-
cal Foundations of Scientific Visualization, Computer Graphics, and Massive Data
Exploration. Springer-Verlag, Mathematics and Visualization (2007)
7. Dey, T.K., Sun, J.: Normal and feature estimations from noisy point clouds. Tech-
nical Report OSU-CISRC-7/50-TR50 (2005)
8. Sonka, M., Reddy, G.K., Winniford, M.D., Collins, S.M.: Adaptive approach to
accurate analysis of small-diameter vessels in cineangiograms. IEEE Trans. Med.
Imag. **16** (1997) 87–95

Vessel-guided airway segmentation based on voxel classification

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

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

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

1 Introduction

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

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

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

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

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

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

2   Vessel-guided airway segmentation

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

2.1   Airway Appearance Model

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

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

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

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

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

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

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

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

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

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

where A is the airway class, $K_A(\boldsymbol{x})$ is the number of nearest neighbors around $\boldsymbol{x}$ belonging to the airway class, obtained from a total of K nearest neighbors.
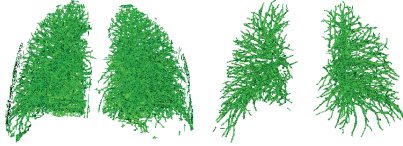
Fig. 1. Vessel segmentation with thresholding only on intensities (left) and on both intensities and tubeness measures (right).

2.2 Vessel Orientation

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

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

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

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

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

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

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

2.3 Orientation Similarity Measure

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

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

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

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

2.4   Segmentation Framework

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

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

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

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

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

2.5   Optimal Threshold Selection

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

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

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

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

3   Results on 20 Low-dose CT Images

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

3.1   Parameter Settings

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

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

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

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

3.2   Results

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

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

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

4   Discussion and Conclusions

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

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

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

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

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

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

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

References

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

ment. Radiology **235**(3) (Jun 2005) 1055–1064

2. Kiraly, A.P., Higgins, W.E., Hoffman, E.A., McLennan, G., Reinhardt, J.M.: 3D human airway segmentation for virtual bronchoscopy. In: SPIE Medical Imaging 2002: Physiology and Function from Multidimensional Images. Volume 4683. (April 2002) 16–29

3. Schlathölter, T., Lorenz, C., Carlsen, I.C., Renisch, S., Deschamps, T.: Simultaneous segmentation and tree reconstruction of the airways for virtual bronchoscopy. Volume 4684., SPIE (2002) 103–113

4. Kitasaka, T., Mori, K., Suenaga, Y., Hasegawa, J., Toriwaki, J.: A method for segmenting bronchial trees from 3D chest X-ray CT images. In: MICCAI (2). (2003) 603–610

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

6. Lo, P., de Bruijne, M.: Voxel classification based airway tree segmentation. In: Medical Imaging 2008: Image Processing. Volume 6914., SPIE (2008) 69141K

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

8. Graham, M.W., Gibbs, J.D., Higgins, W.E.: Robust system for human airway-tree segmentation. In: Medical Imaging 2008: Image Processing. Volume 6914., SPIE (2008) 69141J

9. Sonka, M., Park, W., Hoffman, E.: Rule-based detection of intrathoracic airway trees. IEEE Transactions on Medical Imaging **15**(3) (1996) 314–326

10. Weickert, J., Ishikawa, S., Imiya, A.: On the history of Gaussian scale-space axiomatics. In Sporring, J., Nielsen, M., Florack, L., Johansen, P., eds.: Gaussian Scale-Space Theory. Kluwer Academic Publishers, Dordrecht, The Netherlands (1997) 45–59

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

12. Hu, S., Hoffman, E., Reinhardt, J.: Automatic lung segmentation for accurate quantitation of volumetric X-ray CT images. Medical Imaging, IEEE Transactions on **20**(6) (June 2001) 490–498

13. Bülow, T., Wiemker, R., Blaffert, T., Lorenz, C., Renisch, S.: Automatic extraction of the pulmonary artery tree from multi-slice CT data. Volume 5746., SPIE (2005) 730–740

14. Wang, T., Basu, A.: A note on 'A fully parallel 3D thinning algorithm and its applications'. Pattern Recognition Letters **28**(4) (March 2007) 501–506

15. Arya, S., Mount, D.M., Netanyahu, N.S., Silverman, R., Wu, A.Y.: An optimal algorithm for approximate nearest neighbor searching fixed dimensions. J. ACM **45**(6) (1998) 891–923
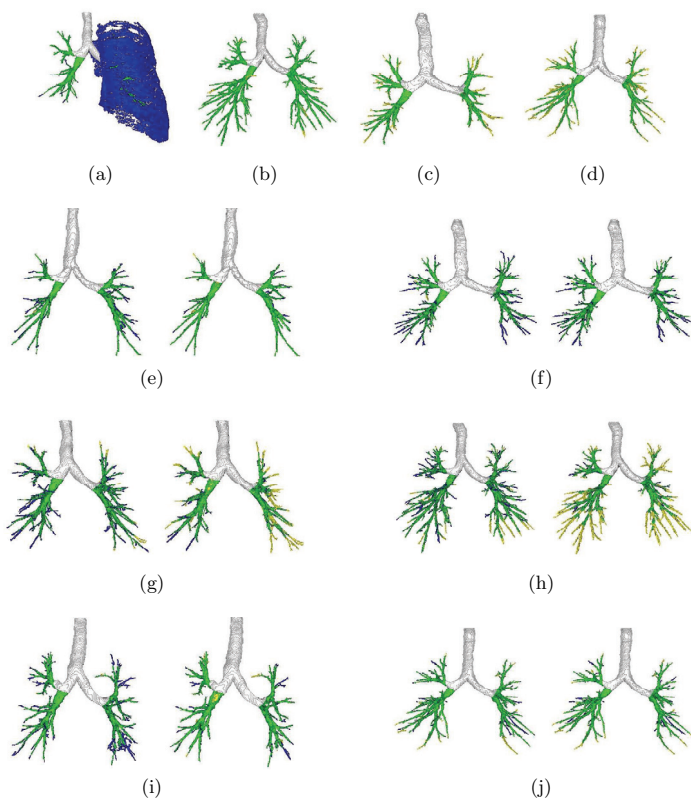
Fig. 2. Surface renderings of the segmentation results. Results from region growing with intensity are given in (a)-(d). Comparison of segmentation results using airway probability with (left) and without (right) orientation similarity measure are given in (e)-(j). Results from the following pairs of figures are from the same test image: (a) and (i), (b) and (j), (c) and (f), (d) and (h). The pre-segmented trachea and main bronchus are shown in white, true positives are show in green, false positives are shown in blue and false negatives in yellow, all with respect to the surrogate ground truth. (Refer to the electronic version for colours.)

Poster Presentations

Lung Tissue Analysis Using Isotropic Polyharmonic B–Spline Wavelets

Adrien Depeursinge[1], Dimitri Van De Ville[2], Michael Unser[2], and
Henning Müller[1,3]

[1] Service of Medical Informatics, Geneva University Hospitals and University of
Geneva, CH, `adrien.depeursinge@sim.hcuge.ch`,
[2] Biomedical Imaging Group, École Polytechnique Fédérale de Lausanne, CH,
[3] Business Information Systems, University of Applied Sciences Sierre, CH.

Abstract. A texture classification system is described, based on isotropic polyharmonic B–spline wavelets that identify lung tissue patterns from high–resolution computed tomography (HRCT) images of patients affected with interstitial lung diseases (ILD). Along with several desirable properties for isotropic texture analysis, the nonseparable transform with a quincunx subsambling scheme allows a mean of 94.3% of correct matches among six lung tissue classes. A comparison with a classical dyadic transform suggests that the isotropic quincunx transform is preferable for lung tissue analysis. This is part of work on a tool for integrating visual and clinical features as diagnostic aid for emergency radiology.

1   Introduction

The interpretation of high–resolution computed tomography (HRCT) images of the chest showing patterns associated with interstitial lung diseases (ILDs) is time–consuming and requires high clinical expertise due to rare cases and a large number of different diseases. The diagnosis of ILD is established from the interpretation of several clinical parameters of the patient in addition to radiological findings [1]. The most common imaging procedure used is the chest x–ray because of its low cost and weak radiation exposure. However, chest x–rays appear as normal in a large portion of diseases and are often unspecific where HRCT of the chest contains essential visual data for the characterization of lung tissue patterns associated with ILDs [2]. HRCT produces three–dimensional (3D) images of the pulmonary volumes, avoids the superposition of anatomic structures, and is well suited for the assessment of lung tissue texture. However, the increase of data volume compared to the chest x–rays makes the interpretation task more complex. The high spatial resolution generates a large variety of lung tissue patterns, which induces confusion of diverse pathologic lung tissues. In addition, the radiologist has to go through the whole stack of slices, which can result in interpretation errors by omission [3]. In emergency radiology, radiologists have recourse to a large diversity of imaging modalities such as conventional

projection radiography, computed tomography (CT), magnetic resonance imaging (MRI), functional imaging (fMRI, PET), and ultrasound applied to different organs such as the brain, colon, breast, chest, liver, kidney and the vascular and skeletal systems. They have to provide a first radiological report with ideas on the diagnosis quickly. Automatic detection and categorization of pathologic lung tissue patterns can help the radiologists to cope with the complexity and challenges of interpreting HRCT [4, 5]. The suspicious (abnormal) patterns in the new, non–interpreted HRCT are highlighted to the radiologist with a proposed tentative diagnostic [6]. The radiologist has to consider the system as a second opinion for providing a differential diagnosis.

The taxonomy used by radiologists to interpret patterns in HRCT images often relates to texture properties, which suggest that texture analysis is relevant for the characterization of ILD which is typically diffuse. Texture analysis in digital image processing has been an active research domain over more than thirty years. In [7], texture in digital images is defined as nonfigurative and cellularly organized areas of pixels. Early examples of texture features are the autocorrelation function, textural edginess, measurements derived from mathematical morphology, run–length and gray–level co–occurence matrices, the latter being the most popular of the lot [7, 8]. Unfortunately, building co–occurence matrices from HRCT images where grey–levels are corresponding to Hounsfield Units (H.U.) with values from –1000 H.U. (air) to 1500 H.U. (high density bones) is unrealistic because the number of possible co-occurrences is simply too large to be stored and estimated reliably. Complementary to the characterization of spatial dependencies, the distribution of grey–levels values can be studied through statistical measures of grey–level histograms. The Fourier transform has also been proposed for texture analysis, based on the property that some image patterns (especially, periodic ones) are well described in terms of sinusoidal components [8, 9]. However, the latter is not appropriate for segmentation because the Fourier transform is global.

1.1   Dyadic Versus Quincunx Wavelet Frames for Texture Analysis

The multiresolution analysis provided by the wavelet transform (WT) is an attractive solution for texture analysis. The signal under investigation is decomposed onto a set of wavelet functions at various scales; i.e., representing details at different resolutions. For example, in the discrete version with dyadic subsampling the analyzed image is iteratively filtered and subsampled by a factor of 2 in each dimension. The WT is particularly well suited to compactly represent piecewise–smooth signals, which partly explains its success in biomedical imaging applications [10]. Wavelet bases are desirable to perform compact data representations, as they allow orthogonal decompositions. They have fast implementations, but have the main drawback to lack translation invariance. Discrete wavelet frames (DWF), on the other hand, are redundant and offer more flexibility for image analysis. DWF are truly shift–invariant and can be obtained from a wavelet basis by removing the subsampling stage of the algorithm and upsampling the filters instead. When compared to the WT, the DWF tends to decrease

the variability of the estimated texture features thereby improving classification performance [11].

Although widely deployed, the separable dyadic form of DWF has two major drawbacks to perform texture analysis. On one hand, the scale–progression is large as images are downsampled by a factor of 2 (in each dimension) between two decomposition levels. Relevant information might be padded out when having major energy contained in a narrow subband located between two successive levels of the dyadic transform. Subtle changes in the scale of lung tissue patterns (i.e. micronodules versus macronodules) might be neglected by the dyadic scale–progression [12]. On the other hand, separability allows computational efficiency because wavelet coefficients within each subband can be obtained by successive one–dimensional (1D) convolutions along the columns and the rows of the image. Unfortunately, this process tends to favor the vertical and the horizontal directions, and produces a so–called "diagonal" wavelet component, which does not have a straightforward directional interpretation. Under the assumption that lung tissue patterns in axial slices of HRCT data do not have privileged directionalities, neither horizontal, vertical, nor diagonal, the separable transform is not appropriate for their analysis.

In Section 2.2, we propose the use of isotropic polyharmonic B–spline wavelets together with a fine scale–progression (equivalent factor of $\sqrt{2}$ based on the quincunx subsampling scheme). Beside scale–progression and isotropy we demonstrate that isotropic polyharmonic B–splines have many other desirable properties for lung tissue analysis.

1.2   Non–Separable Wavelets for Biomedical Texture Classification

During the last twenty years, the WT has been utilized widely in biomedical applications, as well as for the characterization of textures of biomedical tissues [10]. The increased spatial resolution of modern imaging techniques allows for assessment from anatomical structures to textures of tissues. More recently, nonseparable wavelet transforms have been used for detection of pathologic tissues with no a priori privileged directionalities in several imaging modalities.

In [13], the quincunx wavelet transform is used for the characterization of liver tissue in noisy ultrasonic B–scan images. Compared to the classical WT, the nonseparable transform allows for an increased classification performance. However, the use of the compacted pyramidal representation of the subbands is not appropriate to obtain shift–invariant features for classification.

The redundant quincunx wavelet transform along with support vector machines (SVM) were used in [14] to classify 5 lung tissue patterns associated with ILDs and achieved 94.3% of global accuracy. Nevertheless, the classification task is slightly biased since the training set contains an equal number of instances for each of the five patterns, which is usually not the case in clinical practice.

In [12], grey–level histograms with discrete wavelet frame features were evaluated using a k–nearest neighbor classifier. In this paper, we use isotropic polyharmonic B–splines as scaling functions to implement a two–dimensional (2D) redundant quincunx wavelet transform in order to characterize 6 types of lung

tissue in HRCT data with optimized SVMs. Lung tissue texture classification
using co-occurence matrices, Gabor filters and Tamura texture features was in-
vestigated in [15]. The classification of regions of interest (ROIs) delineated by
the user consitutes the intial steps towards automatic detection of abnormal lung
tissue patterns in the whole HRCT volume.
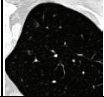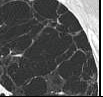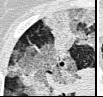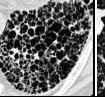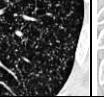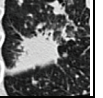
2   Methods

2.1   Dataset

The dataset used is part of an internal multimedia database of ILD cases con-
taining HRCT images with annotated ROIs created in the *Talisman* project[1].
843 ROIs from healthy and five pathologic lung tissue patterns are selected for
training and testing the classifiers selecting classes with sufficiently high repre-
sentation (see Table 1).

The wavelet frame decompositions with dyadic and quincunx subsampling
are implemented in Java [11, 16] as well as optimization of SVMs. The basic
implementation of the SVMs is taken from the open source Java library *Weka*[2].

Table 1. Visual aspect and distribution of the ROIs per class of lung tissue pattern.

visual aspect						
class	healthy	emphysema	ground glass	fibrosis	micronodules	macronodules
# of ROIs	113	93	148	312	155	22
# of patients	11	6	14	28	5	5

2.2   Isotropic Polyharmonic B–Spline Wavelets

As mentioned in Section 1.1, isotropic analysis is preferable for lung texture
characterization. The Laplacian operator plays an important role in image pro-
cessing and is clearly isotropic. Indeed, $\Delta = \frac{\partial^2}{\partial x_1^2} + \frac{\partial^2}{\partial x_2^2}$, is rotationally invariant.
The polyharmonic B–spline wavelets implement a multiscale smoothed version
of the Laplacian [16]. This wavelet, at the first decomposition level, can be char-
acterized as

$$\psi_\gamma(\mathbf{D}^{-1}\mathbf{x}) = \Delta^{\frac{\gamma}{2}} \{\phi\}(\mathbf{x}), \tag{1}$$

---

[1] TALISMAN: Texture Analysis of Lung ImageS for Medical diagnostic AssistaNce,
   http://www.sim.hcuge.ch/medgift/01_Talisman_EN.htm
[2] http://www.cs.waikato.ac.nz/ml/weka/

where ϕ is an appropriate smoothing (low–pass) function and $\mathbf{D} = [1\ 1; 1\ -1]$ is the quincunx subsampling matrix. The so-called *order* γ tunes the iterate of the Laplacian operator (comparable to the traditional vanishing moments). Large values of γ reduce the energy of the wavelet coefficients but increase the ringing effect [17].

2.3   Lung Tissue Classification: Dyadic Versus Quincunx

The quincunx scale–progression is finer compared to dyadic decomposition, with an equivalent downsampling factor of $\sqrt{2}$ instead of 2. Compared to the dyadic separable case, quincunx subsampling generates only one wavelet subband per decomposition level (versus three for dyadic). This leads to a direct and easy interpretation of the subbands; the small number of subbands also breeds small features spaces, which are preferable for classification.

In summary, using isotropic polyharmonic B–splines as scaling functions have the following desirable properties for lung tissue analysis:

– rotational invariance
– fine scale progression tunable through γ
– easy and direct interpretation (one wavelet subband per scale)

2.4   Feature space

In order to build the feature space for further classification of the ROIs, several measures are computed from the original image as well as from the wavelet coefficients of each subband. On the original images, values of pixels belonging to the ROIs are categorized into 22 bins of grey–level histograms of Hounsfield Units (H.U.) in $[-1050; 600[$. The distributions of the wavelet coefficients in each subband are characterized through the parameters of mixtures of two Gaussians, which have shown to characterize distributions of wavelet coefficients well in [18]. With fixed means $\mu_{1,2} = \mu$, the standard–deviations $\sigma_{1,2}$ are estimated using the expectation–maximization (EM) algorithm. Under the assumption that the global mean of the coefficient values is close to zero (according to the admissibility conditions of wavelets), using two standard–deviations allows a reasonable fit of the distributions (see Figure 1). The feature vector thus consists of 24 features for 8 levels of the quincunx transform. Features are extracted for orders $\gamma = 2, 3, 4$. To compare performances, 4 levels of the classical dyadic transform (using frames as well) were performed using B–Spline wavelets of degree $\alpha = 1, 2, 3$. Indeed, the equivalent order of derivatives γ corresponds to $\alpha + 1$. Using parameters of a mixture of two Gaussians for each subband, the feature vector contains 36 measures of the dyadic wavelet frames coefficients (more details can be found in [12]).

3   Results

Feature vectors from 843 ROIs containing healthy and five pathologic lung tissues are extracted. 674 instances (80%) are randomly drawn from the full dataset and
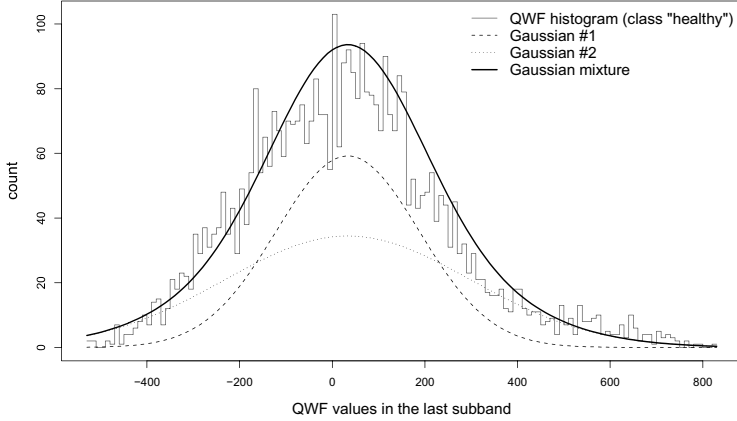
Fig. 1. Mixture of two Gaussians ($\mu_{1,2} = \mu$, $\sigma_{1,2}$) to modelize the distribution of wavelet coefficients within one subband.

used to train and optimize the parameters of Support Vector Machines (SVMs) with Gaussian kernel, which have shown to be effective to categorize texture in wavelet feature spaces in [19] and in particular lung tissue in [20]. The remaining 169 instances are used for testing. The global experimentation is repeated 30 times and means of the global classification accuracies along with means of class–specific accuracies are computed. A detailed description of the selection of the parameters of SVMs and a comparison of 5 common implementations of classifiers families can be found in [20]. Pairwise comparisons of classification accuracies using dyadic versus quincunx wavelet frames for several orders $(\alpha+1, \gamma)$ are shown in Table 2.

4   Discussion

Pairwise comparisons shown in Table 2 indicate that quincunx wavelets outperform dyadic ones in 91.7% of the comparisons (22 among 24). This global increase in performance is primarily due to the better isotropy properties of these non-separable wavelets, which is due to their close connection to the Laplacian. Indeed, the favored directions of the separable transform lead to noisy features breeding non–homogeneous clusters of instances belonging to the same class in the feature space, which decreases global classification performance. Although having influence on global accuracy as well, the finer scale progression allowed by the quincunx subsampling scheme increases the precision of the classification;

Table 2. Mean accuracies in % with experiments repeated 30 times. Isotropic polyharmonic B–spline wavelets with order $\gamma = 3$ allowed a mean of 94.3% of correct predictions among the six lung tissue classes with high precision (geometric mean = 89%).

class		$\alpha = 1, \gamma = 2$	$\alpha = 2, \gamma = 3$	$\alpha = 3, \gamma = 4$
healthy	dyadic	91.1	93.6	92.5
	quincunx	95.9	98.1	92.4
emphysema	dyadic	97.2	98.7	97.7
	quincunx	100	100	99.7
ground glass	dyadic	84.2	88.3	86.3
	quincunx	85.7	89	87.7
fibrosis	dyadic	95.8	95.2	96.5
	quincunx	96.5	96.3	94.5
micronodules	dyadic	89.8	93.3	88.8
	quincunx	94.1	95.2	91.7
macronodules	dyadic	40.3	48	46.9
	quincunx	54.2	55.5	48.5
geometric mean	dyadic	83.1	86.2	84.8
	quincunx	87.8	89	85.7
global mean	dyadic	90.6 ± 2.6	92.5 ± 1.4	91.4 ± 2.2
	quincunx	93.3 ± 1.6	$\mathbf{94.3 \pm 1.6}$	92 ± 1.9

i.e., by avoiding confusion between patterns with well–defined object sizes, such as *micro–* and *macro– nodules*.

Global accuracy values are trustworthy for further usage in clinical routine as the six classes of lung tissue pattern tested allow for diagnosing a wide variety of ILDs [2]. Compared to other studies on lung tissue analysis in HRCT data, our system is closer to clinical routine as the distributions of the classes are realistic contrary to [14] and we include healthy tissue (which is not the case in [4]). Indeed, healthy tissue is the most difficult to separate from others as the variety is by far the largest.

5   Conclusion

The ability of dyadic versus quincunx wavelet transforms to analyze lung tissue in HRCT data were evaluated on a high–quality dataset. Isotropic polyharmonic B–spline wavelets with optimized order allowed a mean of 94.3% correct predictions among six lung tissue classes associated with ILDs with high precision. Pairwise comparisons with a dyadic transform showed that the polyharmonic wavelets outperforms the classical separable frames 22 times among 24, which suggest that the latter is more appropriate for lung tissue analysis in HRCT data.

Further work has to be carried out in order to integrate clinical parameters for classifying the lung tissue regions, in the same way the radiologists interpret HRCT images. First experiments showed high potential for improving classification performances in [21].

6   Acknowledgments

References

1. Flaherty, K.R., King, T.E., Ganesh Raghu, J., Lynch III, J.P., Colby, T.V., Travis, W.D., Gross, B.H., Kazerooni, E.A., Toews, G.B., Long, Q., Murray, S., Lama, V.N., Gay, S.E., Martinez, F.J.: Idiopathic interstitial pneumonia: What is the effect of a multidisciplinary approach to diagnosis? American Journal of Respiratory and Critical Care Medicine **170** (July 2004) 904–910
2. Stark, P.: High resolution computed tomography of the lungs. UpToDate **September** (2007)
3. Kakinuma, R., Ohmatsu, H., Kaneko, M., Eguchi, K., Naruke, T., Nagai, K., Nishiwaki, Y., Suzuki, A., Moriyama, N.: Detection failures in spiral CT screening for lung cancer: Analysis of CT findings. Radiology **212** (July 1999) 61–66
4. Shyu, C.R., Brodley, C.E., Kak, A.C., Kosaka, A., Aisen, A.M., Broderick, L.S.: ASSERT: A physician–in–the–loop content–based retrieval system for HRCT image databases. Computer Vision and Image Understanding (special issue on content–based access for image and video libraries) **75**(1/2) (July/August 1999) 111–132
5. Müller, H., Michoux, N., Bandon, D., Geissbuhler, A.: A review of content–based image retrieval systems in medicine – clinical benefits and future directions. International Journal of Medical Informatics **73** (February 2004) 1–23
6. Sinha, U., Bui, A., Taira, R., Dionisio, J., Morioka, C., Johnson, D., Kangarloo, H.: A review of medical imaging informatics. Annals of the New York Acedemy of Sciences **980** (December 2002) 168–197
7. Haralick, R.M.: Statistical and structural approaches to texture. Proceedings of the IEEE **67**(5) (May 1979) 786–804
8. Reed, T.R., du Buf, J.M.H.: A review of recent texture segmentation and feature extraction techniques. Computer Vision, Graphics and Image Processing **57**(3) (May 1993) 359–372
9. Hsu, T.I., Calway, A.D., Wilson, R.: Texture analysis using the multiresolution fourier transform. In: Proc 8th Scandinavian Conference on Image Analysis, IAPR (May 1993) 823–830
10. Unser, M., Aldroubi, A.: A review of wavelets in biomedical applications. Proceedings of the IEEE **84**(4) (April 1996) 626–638
11. Unser, M.: Texture classification and segmentation using wavelet frames. IEEE Transactions on Image Processing **4**(11) (1995) 1549–1560
12. Depeursinge, A., Sage, D., Hidki, A., Platon, A., Poletti, P.A., Unser, M., Muller, H.: Lung tissue classification using wavelet frames. Engineering in Medicine and Biology Society, 2007. EMBS 2007. 29th Annual International Conference of the IEEE (August 2007) 6259–6262
13. Mojsilovic, A., Popovic, M., Markovic, S., Krstic, M.: Characterization of visually similar diffuse disease from b-scan liver images using nonseparable wavelet transform. IEEE Trans. Med. Imaging **17**(4) (August 1998) 541–549

14. Shamsheyeva, A., Sowmya, A.: The anisotropic gaussian kernel for SVM classification of HRCT images of the lung. In: Proceedings of the 2004 Intelligent Sensors, Sensor Networks and Information Processing Conference. (December 2004) 439–444

15. Müller, H., Marquis, S., Cohen, G., Geissbuhler, A.: Lung CT analysis and retrieval as a diagnostic aid. In: Medical Informatics Europe (MIE 2005), Geneva, Switzerland (August 2005) 453–458

16. Van De Ville, D., Blu, T., Unser, M.: Isotropic polyharmonic B-Splines: Scaling functions and wavelets. IEEE Transactions on Image Processing **14**(11) (November 2005) 1798–1813

17. Feilner, M., Van De Ville, D., Unser, M.: An orthogonal family of quincunx wavelets with continuously adjustable order. IEEE Transactions on Image Processing **14**(4) (April 2005) 499–510

18. Portilla, J., Strela, V., Wainwright, M.J., Simoncelli, E.P.: Image denoising using scale mixtures of gaussians in the wavelet domain. IEEE Transactions on Image Processing **12**(11) (November 2003) 1338–1351

19. Li, S., Kwok, J.T., Zhu, H., Wang, Y.: Texture classification using the support vector machines. Pattern Recognition **36**(12) (December 2003) 2883–2893

20. Depeursinge, A., Iavindrasana, J., Hidki, A., Cohen, G., Geissbuhler, A., Platon, A., Poletti, P.A., Müller, H.: A classification framework for lung tissue categorization. In Andriole, K.P., Siddiqui, K.M., eds.: Medical Imaging 2008: PACS and Imaging Informatics. Volume 6919., SPIE (April 2008) 69190C

21. Depeursinge, A., Iavindrasana, J., Cohen, G., Platon, A., Poletti, P.A., Müller, H.: Lung tissue classification in HRCT data integrating the clinical context. In: 2008 21st IEEE International Symposium on Computer-Based Medical Systems, Jyväskylä, Finland (June 2008) 542–547

Comprehensive Emphysema Subtype Diagnosis Using Structured Expert Knowledge

Jaron Schaeffer[1], Mamatha Rudrapatna[1], Arcot Sowmya[1], Peter Wilson[2]

[1] University of New South Wales, Sydney, NSW 2052, Australia
[2] I-Med Network, Sydney, NSW 2000, Australia

Abstract. Computerised emphysema quantification has received a lot of research attention due to the mass availability of CT. Yet, to our knowledge, no existing method is able to recognise all common subtypes of the disease, a diagnosis routinely given by radiologists. In this paper, we present a HRCT-based Computer Assisted Diagnosis system for emphysema subtype diagnosis. The system first detects low-attenuation regions using adaptive density mask, a novel refinement to the classic density mask method. Detected regions are then classified individually and results combined in a bottom-up manner to achieve per-patient diagnosis and quantification. Expert knowledge necessary for classification decisions was acquired incrementally using a multi-level Ripple Down Rules system. Evaluation shows that the multi-level approach well reflects the pathological characteristics of the subtypes, and RDR knowledge management provided robust diagnosis using very little training data.

1 Introduction

Emphysema is a chronic obstructive pulmonary disease (COPD) typically caused by exposure to tobacco smoke. Alveolar walls break down due to inflammatory responses to the particles inhaled and the affected regions show low attenuation on high-resolution computed tomography (HRCT) scans. Patients suffer from limited respiratory capabilities as the disease progresses.

Automated emphysema detection has been researched for almost as long as CT has existed; Sluimer *et al.* recently provided an overview and performance comparison in their survey on computerised CT analysis of the lung [1]. Despite its age, density mask (DM), a method introduced by Muller *et al.* in 1988 [2], is still considered the de-facto standard in computerised quantification [3]. DM calculates the percentage of the lung showing below-normal attenuation and has been shown to correlate well with pulmonary function tests, the standard for emphysema diagnosis in pre-CT times [4, 5]. However, DM and other emphysema quantification methods are unable to detect disease subtypes, a diagnosis made routinely by radiologists when diagnosing scans manually.

To build a computer assisted diagnosis (CAD) system capable of comprehensive subtype recognition, extensive expert knowledge must be transferred into the system systematically and in a time-efficient way. Ripple Down Rules

(RDR), a knowledge engineering technology for expert systems introduced by Compton and Jansen [6], is suitable for the task. RDR are ordered lists of rules with exceptions. If a misclassification is detected for an input case while reviewing classification results, an exception is appended to the firing rule to correct the conclusion; a new rule is added to the knowledge base if no rules fire at all; and no change is made for correctly classified cases (*review/approve-or-create-rule cycle*). This makes RDR an *incremental* or *per-case* knowledge acquisition method, where knowledge base testing and rule creation merges into a single working step to make use of every sample reviewed. New rules are validated against the existing knowledge base to ensure knowledge consistency. For any new rule or exception created, the human expert is responsible for formulating a rule condition based on some features of the case investigated. The choice of these conditions is a trade-off between rule specificity and the desired generalisation to similar but unseen future cases.

Feature extraction and feature design are crucial for the success of any CAD system, and often, new features are added as the expert finds that the features currently available are insufficient. If the diagnosis model was represented using a machine learning approach, this would require re-training; with RDR being an incremental learner, the expert would simply incorporate new features into the newly built rules. Another important property of RDR is that the system is built while being already in use. RDR can be seamlessly integrated into the day-to-day workflow of the experts as demonstrated by the commercially successful version of RDR for pathology domain [7, 8]. It is known that RDR produce knowledge bases (KB) similar in size to those developed by machine learning [9] and the time taken to add a rule remains roughly constant regardless of the size of the KB [8]. An enhanced version of RDR known as Multiple Classification RDR (MCRDR) [10] can provide multiple conclusions to a given case.

In this paper, we present a CAD system for classification and quantification of *centrilobular* (*centriacinar*), *panlobular* (*panacinar*) and *paraseptal* (*bullous*) emphysema (Fig. 1). The system design is inspired by the way we observed the radiologist in our group handle the task manually: raw emphysema regions are detected (Sect. 2) and classified individually based on a set of region features. Results are combined to provide input to a higher-level classification step (Sect. 3). Both steps use RDR rule bases for expert knowledge management and reasoning, forming a novel multi-level RDR classification system. We discuss evaluation methods and some results in Sect. 4 and summarise remaining challenges and ideas for future work in Sect. 5.

2   Emphysema Detection and Feature Extraction

Raw emphysema region detection is based on density mask, as this method is widely accepted by radiologists and studies have been carried out to select good threshold values [4, 5]. However, a pure DM approach fails in our case for later stage emphysema cases (*confluent emphysema*), where large regions of very serrated appearance are detected that provide no meaningful shape information
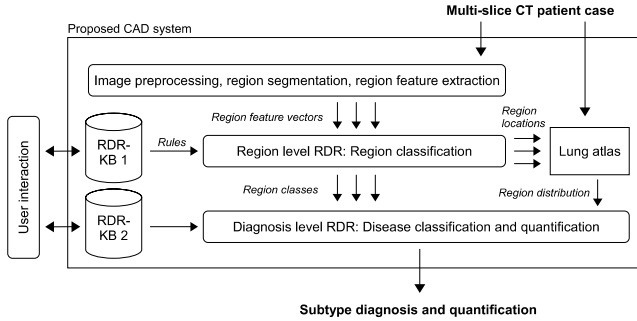
Fig. 1. CAD System overview

(see Fig. 2). To compensate, we have developed a DM variant called Adaptive Density Mask (ADM), somewhat similar to but computationally more efficient than the standard Watershed transform [11]. ADM was designed to give results similar to standard DM except for confluent cases, where it creates more and smaller unconnected regions more suitable for subsequent region analysis.

2.1 Adaptive Density Mask

The idea behind ADM is to select a binarisation threshold for each emphysema region automatically to maximise the number of distinct regions while simultaneously maximising their area, both within the limits imposed by minimum/maximum density masks. Traditional minima detection methods capable of achieving this [11, 12] are expensive, over-segment the image and require advanced de-noising beforehand. ADM provides a simple, fast and robust tradeoff between sensitivity and simplicity for the drawback of having to select a number of additional parameters.

It is worth pointing out that, by definition, DM (and any method based on it such as ADM) does not produce false positives in the traditional sense; every detected image pixel/voxel is considered as indication of the presence of emphysema and thus clinically relevant. Fig. 3 illustrates ADM schematically for a real-valued signal; in particular, the method works as follows:

Given a set of thresholds $t_1, \ldots, t_n$, a DM is calculated on the input image for each threshold (for e.g. -965 HU to -945 HU using 5 HU steps in Fig. 3). For each resulting distinct region (solid horizontal black lines), the minimum attenuation is determined and the corresponding x/y position is associated as an unique index to the region (squares and vertical lines). We can think of the different density mask outputs as an ordered stack of layers as indicated by the solid horizontal lines (binary one) and the dotted horizontal lines (binary zero) in Fig. 3. Each region is recursively visited by the following procedure, starting on the topmost layer:

(a)　　　　　　(b)　　　　　　(c)　　　　　　(d)
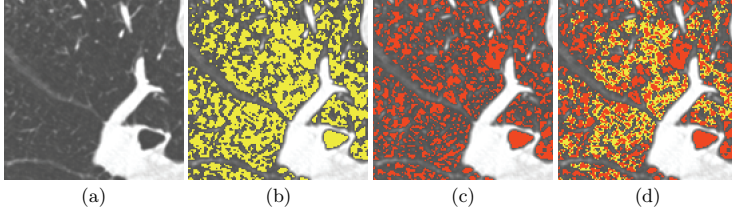
Fig. 2. Standard density mask (DM) compared to the adaptive density mask (ADM) method for raw emphysema segmentation. (a) Original image region, a late case of centrilobular emphysema. (b) DM output using a threshold of -950 HU (yellow). (c) ADM output using -975 HU, -970 HU, ..., -945 HU thresholds (red). (d) Overlay of the DM (yellow) and ADM (red) segmentations from (b) and (c). For large segments of late panlobular and centrilobular emphysema, ADM creates more and smaller unconnected regions as compared to DM that results in large, serrated regions.
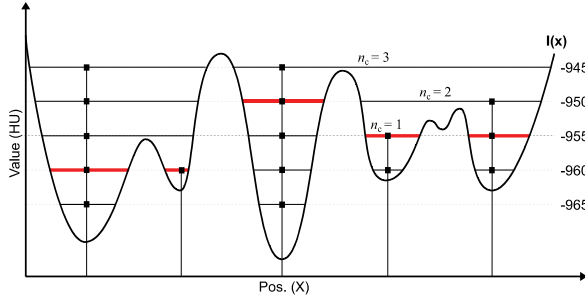
Fig. 3. Adaptive Density Mask: schema for a real-valued signal. The density mask outputs for the different thresholds are depicted in solid (binary one) and dotted (binary zero) lines, resulting emphysema regions are highlighted in bold red. For some regions, the number of regions below the current layer n_c is given.

- Find the number n_c of regions below the current layer covered by the current region. For the rightmost region in Fig. 3 on the -945 HU density mask, $n_c = 3$.
- If $n_c \leq 1$, the current region on the current layer will cover a maximum area while simultaneously maximising the number of regions: recursing to a lower layer will not increase the number of regions but will decrease the region area. Mark the current region in the result mask as binary one and stop recursion.
- If $n_c > 1$, visit the covered region in the layer below.

We observe the following two properties of ADM:

1. *Local minima detection:* The image is re-quantised using $t_1, \ldots, t_n$, and every local minimum detectable in the re-quantised image creates a distinct region

in the binary result image. This ensures that even for late confluent cases, shape information is retained.

2. *Maximum region area:* While the number of regions in the binary result image is determined by the number of detected local minima, each such region has maximum area. This makes ADM behave similarly to DM where the emphysema regions are generally well formed and isolated (non-confluent).

2.2 Region Feature Extraction

We extract multiple features for each low-attenuation region: Shape (compactness and elongation using image moments), first-order statistics (mean HU, HU standard deviation and uniformity using min/max normalised standard deviation), connectivity to the lung boundary (*pleura*), edge-radius-symmetry transform [13] for wall detection and a morphological wall detection approach in which the region is dilated and first order statistics are calculated for the new region pixels, resulting in a total of 12 features per region.

3 RDR Based Multi-level Classification

Our multi-level RDR system design is motivated by the finding that presence and distribution of the different emphysema region classes in a patient are commonly used by radiologists for subtype diagnosis [3]. Emphysema region identification and classification is a task on its own and must not be confused with the higher-level per-patient subtype classification.

Each disease subtype is dominated by a special type of low attenuation region [3] (see Table 1 for a brief characterisation). Accordingly, the extracted region feature vectors are used to classify regions into one of these region classes using a *region level* RDR layer. In a second step, a *diagnosis level* RDR layer builds on the region level results and produces the desired subtype diagnosis as outlined in the schema in Figure 1.

Table 1. Region class characterisation and relation to emphysema subtypes (characterisation extracted from [3])

Region class	Appearance	Predominant in subtype
Bullous	Round, compact, any size. Very low density, typically touching the pleura, visible walls	Paraseptal
Centrilobular	Small/medium size; typically rounded shape	Centrilobular
Diffuse	Large, often smoothly outlined by interlobular fissures and pleura; uniform density with smaller vessels within	Panlobular

Expertise for the region classification decision was transferred into the system using the common RDR cycle: new rules or exceptions to existing rules are added

by the expert either to correct a misclassified sample or to remove false positives. The radiologist in our group, a specialist on lung CT with more than 30 years of experience, created a total of 22 rules for the region classification rule base. Effectively, 22 regions from a number of typical cases were thus used as the gold standard (see Figure 4 for a sample rule created).

- TouchingPleura = 1, Area > 6.95 mm^2, Compactness > 40%, MeanDensity < -974 HU → **Bullous**
  except
    • WallsMorphMean < -890 HU → **Centrilobular**
      except
        ∗ WallsMorphStddev < 100 HU, Uniformity > 85% → **Bullous**

- EmphysemaDistribution={Diffuse,Bilateral}, EmphysemaTotalLungPct > 30% → **Moderate panlobular emphysema**

Fig. 4. Two sample RDR rules created during system development: a region level rule with an exception that in turn has an exception (top) and a diagnosis level rule (bottom).

Percentage involvement of all the three region types detected in the previous step are calculated for various lung regions. These percentages as well as distribution and predominance attributes (such as e.g. diffuse/focal, unilateral/bilateral distributions or apical/middle/basal predominance, calculated for each region class separately using [14]) are presented as features at the diagnosis level RDR. We use multi-classification RDR [10] to allow for co-existence of different emphysema subtypes, and each subtype is individually quantified as either *absent*, *mild*, *moderate* or *severe*. Figure 4 shows a sample diagnosis level rule based on these features.

4   Classifier Performance Evaluation

Evaluations were carried out separately for the region level and the diagnosis level layers. In addition, in order to examine whether our region features are powerful enough for a clear discrimination between the three region classes, we compared region classification results obtained using standard machine learning methods to expert opinion, where high agreement would indicate sufficiently powerful region features.

A dataset consisting of 4,514 manually labelled regions (176 bullous, 4,193 centrilobular and 145 diffuse) from 9 scans of different patients was created using a designated labelling tool to train and evaluate the following classifiers: a decision tree (C4.5), a naive Bayesian, decision tables using the Inducer of Decision Table Majority (IDTM) induction algorithm [15] and a fully connected single hidden layer perceptron classifier, all through the WEKA data mining

suite for Java [16]. Sample sizes for each class were chosen to roughly reflect the frequency distribution of the different region classes. Classifiers were trained using 10-fold stratified cross validation using 10 repetitions. Table 2 shows high agreement for all classifiers tested and proves that the proposed region features are sufficiently powerful for region discrimination.

For comparison, Table 3 displays the confusion matrix for the region-level RDR classifier after the creation of the 22 rules mentioned above for the same dataset. Very good performance is observed for the centrilobular and the diffuse region classes, while almost 33% of the bullous class are missed. Apparently, an insufficient number of samples were reviewed to trigger rule creation and knowledge transfer necessary for reliable classification of this region type. Also, since the same expert hand-labelled the region dataset and created the RDR rules, we cannot eliminate the possibility of some bias and thus over-estimation of the RDR classifier performance.

Table 2. Class-specific and overall classification results for standard machine learning region level classifiers using 4514 training regions. *F-Measure* is defined as the harmonic mean between classifier precision and recall; the *Total F-Measure* column contains a weighted sum of the class-specific F-Measures (weights proportional to class sample size).

	Class-specific F-Measure (%)			
Classifier	Bullous (176 inst.)	Centrilobular (4,193 inst.)	Diffuse (145 inst)	Total F-Measure
C4.5	67.50	98.50	89.80	97.01
Naive Bayes	50.40	94.80	64.20	92.09
Decision Tables	63.70	98.40	89.10	96.75
Perceptron	65.20	98.20	85.00	96.49

Table 3. Confusion matrix for RDR region level classifier using 22 training regions. The value at the bottom right shows the combined F-Measure for the RDR classifier as a result of a weighted average of the class-specific F-Measures (class weight proportional to class sample size).

↓ Label/classified →	Bullous	Centrilobular	Diffuse	F-Measure (%)
Bullous	119	55	2	66.48
Centrilobular	57	4,133	3	98.49
Diffuse	6	12	127	91.70
Combined				97.02

For the diagnosis level layer, the radiologist in our team reviewed an additional 25 scans (different from the ones used to create the region dataset) manually. Since raw emphysema detection is out of scope for this work, we only used scans that appeared to contain emphysema through their relatively high DM coverage (31.3% ± 8.4% for the selected scans using a standard -950 HU threshold). Emphysema subtypes present were identified and a severity diagnosis was given for each case. 5 scans were diagnosed not to show any emphysema. Among the others, 1 showed mild paraseptal (bullous) emphysema, 7 centrilobular (mild: 1, moderate: 2, severe: 4) and 14 panlobular (mild: 10, moderate: 2, severe: 2) emphysema. The diagnosis level RDR knowledge base was initially empty, and during the course of the review, the radiologist added new rules whenever the system came to a wrong or no conclusion.

Altogether, 12 rules were created this way; we recorded the order in which scans were reviewed and whether one or more rules were created to adjust the system behaviour. Two different visualisations of the process can be seen in Fig. 5. Figure 5 (a) shows that 12 rules were needed on the whole to classify the 25 scans correctly. Figure 5 (b) clearly highlights a plateau effect (solid line) and a decay (dotted line) as the number of scans reviewed increases: 10 out of the 12 rules (83.3%) were added for the first 17 scans (68%), confirming the well-known RDR benefit of robust classification results after using only a small number of training samples. Rather than testing on unseen cases, every new case is used for training in the RDR tradition.

(a)   (b)
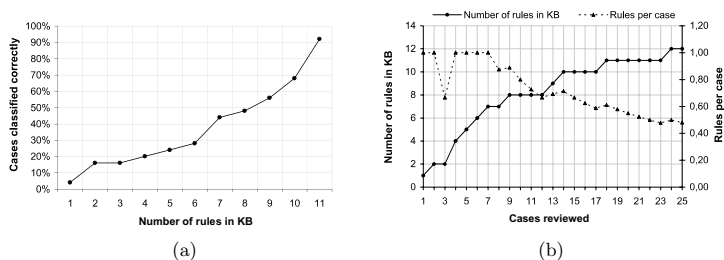
Fig. 5. Results of diagnosis level RDR evaluation. (a) Number of rules in the knowledge base and overall classifier performance on the 25 scan test set. (b) Scans reviewed and number of rules in diagnosis level RDR KB for the 25 scan test set used for evaluation (left axis)/average number of rules per scan (right axis).

5 Conclusion and Outlook

We have presented the first CAD system capable of a comprehensive emphysema diagnosis including disease subtypes. The RDR subtype diagnosis is parameterless and only influenced by the knowledge transferred into the system by the

radiologist(s), suggesting a scenario for a routine clinical use: a fraction of the scans passing through the CAD system in daily use can be reviewed manually and inappropriate classification behaviour can be adjusted. Evaluation of the RDR region classification showed results comparable to state-of-the-art machine learning techniques using only a fraction of training samples.

Longer-term clinical application is one of our goals and still the only way to demonstrate the system's practicability. However, for an incremental system such as the proposed one, results can be expected to improve as more knowledge is gradually added to the adapting system. The presented region detection and multi-level RDR classification may also be used as a generalised framework for subtype classification of other lung diseases (e.g. asbestos-related diseases using pleural plaque, diffuse thickening and pleural rind regions).

For the raw emphysema detection step, we have presented a refined version of density mask called Adaptive Density Mask (ADM) as an extension to the standard density mask method. Only ADM makes it possible to analyse the shape of low-attenuation regions even with late cases of emphysema.

Data for this work was recorded using scan intervals of 10-15 mm, preventing accurate 3D processing. However, extending the proposed system to 3D should be a straight forward task and can be expected to increase robustness and accuracy.

Throughout this paper, we assume that the CT scans under investigation contain emphysema and no other lung disease characterised by low-attenuation regions. This assumption might not hold in practice, for e.g. fibrosis often co-occurs with emphysema and might influence ADM detection results. In addition to that, non-emphysema low attenuation regions (such as e.g. the bronchial tree) should be segmented beforehand and excluded to improve the reliability of the computed volumes.

References

1. Sluimer, I., Schilham, A., Prokop, M., van Ginneken, B.: Computer analysis of computed tomography scans of the lung: a survey. IEEE Transactions on Medical Imaging **25**(4) (April 2006) 385–405
2. Muller, N., Staples, C., Miller, R., Abboud, R.: Density mask. An objective method to quantitate emphysema using computed tomography. Chest **94**(194) (1988) 782–787
3. Webb, W.R., Mueller, N.L., Naidich, D.P.: High-Resolution CT of the Lung. 3rd edition edn. Lippincott Williams & Wilkins (2001)
4. Stern, E., Frank, M.: CT of the lung in patients with pulmonary emphysema: diagnosis, quantification, and correlation with pathologic and physiologic findings. American Journal of Roentgenology **162**(4) (1994) 791–798
5. Kinsella, M., Muller, N., Abboud, R., Morrison, N., DyBuncio, A.: Quantitation of emphysema by computed tomography using a" density mask" program and correlation with pulmonary function tests. Chest **97**(2) (1990) 315
6. Compton, P., Jansen, R.: Knowledge in context: a strategy for expert system maintenance. In: AI '88: Proceedings of the second Australian joint conference on Artificial intelligence, New York, NY, USA, Springer-Verlag New York, Inc. (1990) 292–306

7. Edwards, G., Compton, P., Malor, R., Srinivasan, A., Lazarus, L.: Peirs: A pathologist-maintained expert system for the interpretation of chemical pathology reports. Pathology **25**(1) (1993) 27–34

8. Compton, P., Peters, L., Edwards, G., Lavers, T.: Experience with Ripple-Down Rules. Applications And Innovations in Intelligent Systems XIII: Proceedings of AI-2005, the Twenty-fifth SGAI International Conference on Innovative Techniques and Applications of Artificial Intelligence, Cambridge, UK, December 2005 (2006)

9. Compton, P., Preston, P., Kang, B., Yip, T.: Local patching produces compact knowledge bases. A Future for Knowledge Acquisition/EKAW '94 **94**

10. Kang, B., Compton, P., Preston, P.: Multiple Classification Ripple Down Rules: Evaluation and Possibilities. Proceedings 9th Banff Knowledge Acquisition for Knowledge Based Systems Workshop (1995) 17–1

11. Roerdink, J., Meijster, A.: The Watershed Transform: Definitions, Algorithms and Parallelization Strategies. Mathematical Morphology **41** (2000) 187–S28

12. Fairfield, J.: Toboggan contrast enhancement for contrast segmentation. Pattern Recognition, 1990. Proceedings., 10th International Conference on **1** (Jun 1990) 712–716 vol.1

13. Chabat, F., Hu, X.P., Hansell, D.M., Yang, G.Z.: ERS transform for the automated detection of bronchialabnormalities on CT of the lungs. IEEE Transactions on Medical Imaging **20**(9) (September 2001) 942–952

14. Zrimec, T., Busayarat, S., Wilson, P.: A 3D Model of the Human Lung. Medical Image Computing and Computer-Assisted Intervention-MICCAI (2004) 1074–1075

15. Kohavi, R.: The power of decision tables. Proceedings of the Eighth European Conference on Machine Learning (1995) 174–189

16. Witten, I.H., Frank, E.: Data Mining: Practical machine learning tools and techniques. 2nd edition edn. Morgan Kaufmann (June 2005)

Modelling the respiratory motion of the internal organs by using Canonical Correlation Analysis and dynamic MRI

Gang Gao[1], Jamie McClelland[1], Segolene Tarte[2], Jane Blackall[1], David Hawkes[1]

[1] Centre for Medical Image Computing, University College London, London, WC1E 6BT
[2] Faculty of Classics, Oxford University, OX1 3LU

Abstract. Radiotherapy (RT) treatments to lung tumours subject to significant respiratory motion have been proved to be difficult. By studying the respiratory motion of lung tumour from imaging modalities such as ultra-fast MRI and 4DCT, the ultimate research task is to model the tumour's respiratory motion and to use the model to predict the tumour motion. In this paper, we are proposing a method to build such a model by using a statistical technique called canonical correlation analysis. We built the model from dynamic MR volumes acquired from five volunteers. The leave-n-out (n=12) technique was used to evaluate the accuracy of the motion prediction. The motion prediction results were compared to the motion fields generated by using a B-Spline based non-rigid registration algorithm. The mean absolute differences between the two motion fields are 3.40±3.20mm, 3.62±3.08mm, 3.68±3.50mm, 4.62±3.97mm and 4.29±3.14mm. Our method is novel and efficient. Consider the model was built from low-resolution (5x5x5mm) MR volumes, the results were satisfactory. More thorough evaluations will be carried out on clinical data.

Keywords: Radiotherapy, MRI, respiratory motion model, CCA

1   Introduction

Radical radiotherapy (RT) is one of the primary treatments to non-small cell lung cancer. It uses high-energy X-Ray to kill cancer cells by causing irreparable damage to their DNA. Computed Tomography (CT), which provides high resolution anatomy and contains X-Ray attenuation need for dose calculation is usually used to plan RT procedures. However lung tumours may exhibit significant respiratory motions, limiting the accuracy of dose calculation and delivery in RT procedures. Treating lung tumours that are subject to respiratory motion has been a very active research topic in the last five years. Approaches include treating at breath hold, target delineation and dose calculations, gated treatment and tracked treatment. Except treating at breath hold, which relied on the reproducibility of the breath hold, the other three approaches require a good understanding of lung

tumour motion during respiratory cycles. However, directly tracking the tumour motion in real-time during RT treatments is difficult. The uses of implanted markers are invasive and can only measure the motions in a few locations. By studying the respiratory motion of lung tumours and its surrounding anatomy from imaging modalities like X-Ray imaging, 4DCT and MRI, it is possible to create a mathematic model which could ultimately be used to predict the tumour position [1]~[6]. Approaches to create such models have been reported in the past years. Low et al. proposed a linear model based on two continuous respiratory parameters, the volume and flow measured by spirometry [1]. The location of an internal point of interest (POI), normally in the tumour will depend on its location in the planning CT volume and the current values of the volume and air flow. A simple linear model can be built from the correlation the respiratory parameters and the positions of the POI. Khamene et al also build a model based on two respiratory parameters [2]. However the two parameters were yielded from a larger number of respiratory signals by applying Principal Component Analysis (PCA). The first two principal components (PC) were used as the respiratory parameters. In Khamene's paper, PCA was only used to reduce data dimensionality. Zhang et al extended Khamene's idea by using PCA to reduce the data dimensionality as well as to characterise the internal organ motion from the current height of the diaphragm and its height 1.5 second earlier [3]. Besides the extended application of PCA, Zhang used a free-form deformation field computed from a non-rigid registration algorithm to represent the lung tumour motions. Indeed, the non-rigid motion field does not only represent the tumour motion but also represent the motions of other internal organs, including the lungs and the diaphragm. Similar approach was used in [4][5][6]. McClelland et al found a B-Spline transformation based non-rigid registration with control points spacing of 20x20x20mm was capable of accurately representing the lung tumour motion [6]. He used a B-Spline cyclic function to model the respiratory parameters (B-Spline deformation) by using motion signals measured from a skin marker. Relatively good results were achieved.

In this paper, we propose a novel algorithm to model the respiratory motions of the internal organs and link these motions to externally monitored surrogates such as motion of the abdominal and thoracic skin surface. The algorithm involves two key steps: 1). Create the free-form deformation fields by non-rigidly registering a breath-hold reference MR volume to a set of free-breathing dynamic MR volumes; 2). Build the motion prediction model by using a statistical technique called Canonical Correlation Analysis (CCA).

In statistic, CCA is a well-known method developed by Hotelling in 1936 [7] to study the relationship between two multidimensional variables. Since its invention, CCA has been widely used in many areas including psychology, neuroscience and etc. In recent years it was used to detect neural activity in function MRI [8][9]. In this paper, we will demonstrate that CCA can also be a successful modelling technique to predict the respiratory motion. Experiments were carried out on five volunteer data sets. The results and discussion will be presented at the last sections of this paper.

2   Method

2.1 Canonical correlation analysis (CCA)

Given a p-dimensional random variable X and a q-dimensional random variable Y ($q < p$), both of which have zero mean, CCA seeks linear combinations aX and bY such that they are maximally correlated.
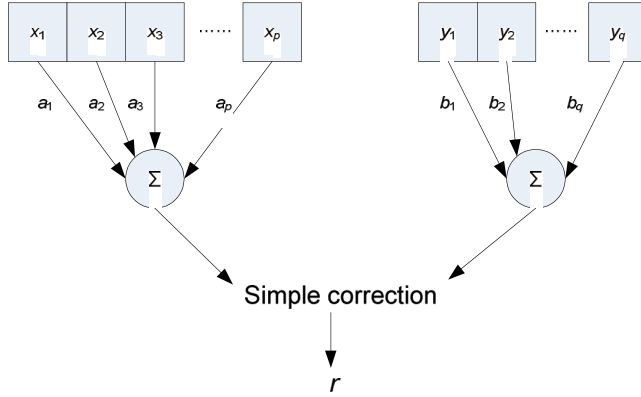
Fig.1. CCA finds the linear combination of coefficients a_1, $a_2...a_p$ and b_1, $b_2...b_q$ to gives the maximum correlation between X and Y.

Fig.1 illustrates how CCA works. Canonical correlation r between X and Y can be found by solving the eigenvalue equations (equation 1).

$$C_{XX}^{-1}C_{XY}C_{YY}^{-1}C_{YX}a = r^2a$$
$$C_{YY}^{-1}C_{YX}C_{XX}^{-1}C_{XY}b = r^2b \tag{1}$$

where C_{XY} is the covariance matrix between vectors X and Y. The eigenvalue r^2 is the squared canonical correlation. a and b are the eigenvectors to the matrices $C_{XX}^{-1}C_{XY}C_{YY}^{-1}C_{YX}$ and $C_{YY}^{-1}C_{YX}C_{XX}^{-1}C_{XY}$. They are also referred to as canonical weights.
Given U and V where

$$U = aX \quad V = bY \tag{2}$$

U and V are called canonical variates. Equation (2) is called canonical function. Up to q canonical functions can be found between p-dimensional variable X and q-dimensional variable Y ($q < p$). In principle, they must be non-correlate (orthogonal) to each other.
   Canonical correlation is the maximum possible correlation between two multi-dimensional variables. The relationship between the two variables can be further investigated by analysing their canonical loadings (CL) and canonical cross loadings (CCL). By definition, CL is the covariance matrix between the original variable X or Y and its canonical variate U or V. Similar to CL, CCL is the covariance matrix between X

or Y and its counterpart's canonical variate V or U. Equation (3) and (4) calculate CL and CCL. CL and CCL explain the amount of variable X explained by variable Y or vice versa.

$$C_{UX} = \text{Cov}(U, X) \quad C_{VY} = \text{Cov}(V, Y) \tag{3}$$

$$C_{VX} = Cov(V, X) \quad C_{UY} = Cov(U, Y) \tag{4}$$

Equation (4) is equivalent to matrix equation (5)

$$V = C_{VX}X \tag{5}$$

From equation (5) we have

$$X = C_{VX}^{-1}V \tag{6}$$

By combining equation (6) and equation (2), we have

$$X = C_{VX}^{-1}\text{b}Y \tag{7}$$

Equation (7) is considered as a prediction model which characterises vector X by using vector Y. Considering vector X as the internal organ motion signals and vector Y as the surrogate signals measured from the skin surface, equation (7) estimates the internal organ motion from the surrogate signals. By using equation (7), we presume that the correlation between the internal motion and the external motion does not change over time. A higher CL or CCL means more elements in X or Y contribute to the construction of the model, suggesting the possibility of a more accurate prediction. This theory will be proved by our experimental results.

2.2   Relationship between CCA and PCA

CCA and PCA are similar to each other in two ways. Firstly, both CCA and PCA are linear subspace methods. Secondly, both of them solve the same equation (equation 8).

$$B^{-1}Aw = r^2w \tag{8}$$

In fact equation (8) is equivalent to equation (1) given

$$\mathbf{A} = \begin{bmatrix} 0 & C_{XY} \\ C_{YX} & 0 \end{bmatrix}, \mathbf{B} = \begin{bmatrix} C_{XX} & 0 \\ 0 & C_{YY} \end{bmatrix} \text{ and } w = \begin{pmatrix} a \\ b \end{pmatrix}$$

However, PCA are fundamentally different with CCA because it solves equation (8) by using different matrices **A** and **B**. In PCA, matrix **A** equals to C_{XX}, matrix **B** is a constant I.

2.3   Internal organ motion

Dynamic MR volumes had high temporal resolution (0.5s for a full 3D volume) and a large field of view (480x480x265mm) showing the skin surface and the internal organs, but had a low spatial resolution (5x5x5mm). A breath-hold high resolution MR volume (1.875x1.875x2mm) was acquired when the subject was instructed to hold the breath. Dynamic MR volumes were non-rigidly registered to the reference volume using a B-Spline registration algorithm [11]. The B-Spline control point grid contained 2250 control points (15x15x10 with 40x40x40mm spacing), each of which had a 3D displacement (dx, dy, dz). The displacement vector D for the registration result is given by:

$$D = [dx_1, dy_1, dz_1, \ldots, dx_n, dy_n, dz_n]$$

Here, $n = 2250$.

Applying CCA directly to the displacement vector D is computationally prohibited. PCA, a dimensionality reduction technique was performed to reduce the number of the internal organ motion variables to n principal components (PC). Zhang et al chose the first two PCs which covered more than 83% of variance to build their model [3]. However, for the datasets we used, we need three PCs to represent around 80% of the variations.

Given the control point motion D, the PCs P. Equation (9) estimates the original variables.

$$D` = \overline{D} + \sum_{k=1}^{K} \lambda_k P_k \approx D \tag{9}$$

where λ is the principal component coefficient.

2.3   Skin motion

The skin surface is clearly visible in the dynamic MR volumes, and its position can be measured automatically using a threshold algorithm similar to the method used in [4] (fig.2). 6 control points were manually picked from the thoracic and the abdominal areas. For each control point, current position m_t and its precursor position m_{t-2} were measured to form the skin motion signals.

$$M = [m_{1,t1}, m_{1,t2}, \ldots\ldots m_{6,t1}, m_{6,t2}]$$

where

$$m_{i,t2} = \begin{pmatrix} m_{i,t-2} (t \geq 2) \\ m_{i,t+j-2} (t \leq 2) \end{pmatrix}$$

By adding a precursor position $m_{i,t-2}$ we had the advantage that it incorporated temporal correlations into the model, distinguishing the inspiration and expiration portions of the breathing cycle.
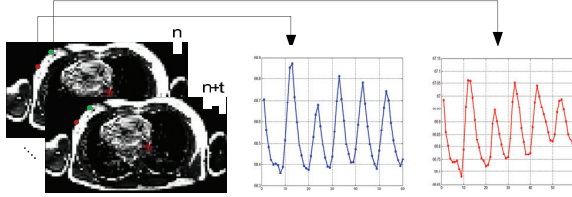
Fig.2. Skin motions were measured from the dynamic MR volumes by using a threshold based segmentation algorithm.

The input variables to CCA must be zero mean. Equation (10) calculates the standarised motion signals.

$$M^* = \frac{(M - \overline{M})}{std(M)} \tag{10}$$

Given the PCs of the internal organ motions P, and the skin motions M, CCA calculates CCL C and canonical weight b between P and M. From equation (7) we have

$$P = C^{-1}bM^* \tag{11}$$

By using equation (9) and (11), the estimated control point displacements D' can be calculated from the real-time measurement of the skin motion.

3   Experiment results

3.1   Spectrum of PCA Eigenvalues

For all the volunteers, PCA was applied to the internal control point displacements to reduce the number of variables from 6750 to 3 PCs. For the five volunteers, the first three eigenvalues account for 86.3%, 82.0%, 76.5%, 95.9% and 85.4% of the total variance in the data.

3.2   CCA

The CCA function was implemented in Matlab (The MathWorks, USA) statistics toolbox. The two parameters of CCA were the PCs of the internal motions and the skin motions. For all the five volunteer datasets, the skin motion variables had strong canonical correlations with the internal motion variables. The mean correlations between the three

pairs of canonical variables were 0.983 ± 0.007, 0.825 ± 0.08 and 0.595 ± 0.07. CL and CCL measured the strength of overall relationships between the skin motions and the internal organ motions. The mean cumulative sums of CL for the five datasets were 0.83, 0.89, 0.87, 0.65 and 0.80 respectively. The mean cumulative sums of CCL or the five datasets were 0.70, 0.83, 0.83, 0.54 and 0.77. The higher the cumulative sum of CL or CCL is, the more likely the model can deliver a satisfactory result. In the fourth dataset, we observed considerably lower values of CL and CCL. This indicated the motion predicted by the model built from the fourth dataset might not be as accurate as the models built from the other datasets.

3.3   Model evaluation

Five separate CCA models were built from the volunteer MR data. The registration results were assessed visually by an expert. Four out of five dataset were considered to be registered successfully. For one dataset, the registration results exhibit more than two voxel misalignment in the borders of the chest and the lungs. The leave-n-out strategy was adopted to evaluate the performance of the CCA models. From each of the datasets, 12 volumes which cover at least one respiratory cycle were dropped. The models were built from the other 48 volumes and used to predict the deformation fields of the missing volumes. Fig.3 shows the deformation field at an arbitrarily selected point inside the dynamic volume and the predicted deformation field at this point from one of the five CCA models. Generally the organ respiratory motions in the anterior-posterior direction (Y) and the left-right direction (Y) are not as considerable as the motion in the foot-head (Z) direction. Therefore the motion signals in x and y directions were easier to be contaminated by noise. Fig.3a and fig.3b show our attempt to model and predict the noisy signals. Fig.3c shows the predicted deformation field is closely matched the measured deformation field in Z direction. We calculated the mean absolute difference (MAD) between the predicted deformation fields and the deformation fields generated by a B-Spline-based non-rigid registration algorithm. For the five datasets, the MADs are 3.40±3.20mm, 3.62±3.08mm, 3.68±3.50mm, 4.62±3.97mm and 4.29±3.14mm. Fig.4 shows the MAD maps generated from the B-Spline control points around the lungs in the mid-coronal slices (Y=0). It is clear that most of the errors are around 5mm. In some extreme cases, the differences can be up to 12mm because of the noise, the image registration errors and the surrogate signal measurement errors. The motion prediction errors around the lungs in the 4[th] volunteers are bigger than the errors in other subjects because of two major reasons. 1). The volunteer has a fast respiratory rate and the temporal resolution of the dynamic MRI (0.5s/volume) is insufficient to produce an accurate motion prediction model; 2). Considerable errors were produced by the image registrations.

The cumulative sums of CL and CCL of the fourth dataset are considerably lower than those of the other four datasets. The experimental data prove the theory that the signal

prediction from the CCA model with high cumulative sum of CL or CCL is not as accurate as the prediction from the model with high cumulative sum of CL or CCL.

(a)   (b)   (c)
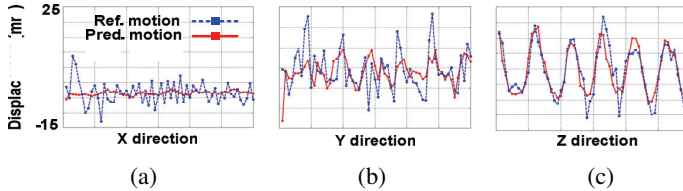
Fig.3. 12 volumes were dropped successively from the dynamic MR volumes. A leave-n-out strategy was used to predict the motion signals measured from the missing 12 volumes. From an arbitrarily selected point inside the MR volume, the predicted motion signals are closely matched the measured motions in Z direction. In X, Y and Z directions, the mean errors are 2.86 ± 2.40mm, 3.62 ± 3.16mm and 2.89 ± 2.37mm.

Generally, the results are satisfactory considering our models were built from low-resolution MRI with 5x5x5mm voxel spacing. Therefore, the errors correspond to about one voxel.

Subject 1   Subject 2   Subject 3   Subject 4   Subject 5
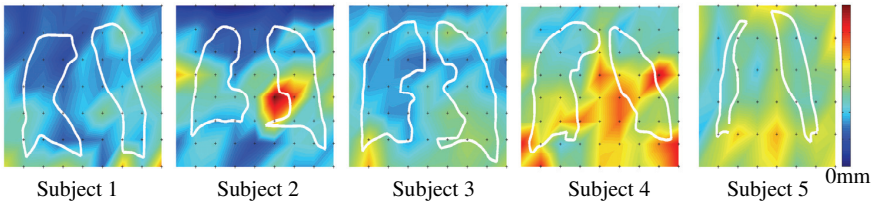
Fig.4 The MAD maps were calculated from the mid-coronal slices (Y=0) showing that most of the prediction errors around the lungs are approximately 5mm. The motion prediction errors around the lungs in the 4[th] subject were higher because of the large image registration errors. The white contours outline the boundary of the lungs and the dark spots suggest the locations of the B-Spline control points.

4   Discussion and conclusions

We proposed an algorithm to model the respiratory motion of the internal organs by using CCA and MRI. In this pilot study, CCA models were built from dynamic MR volumes acquired from five volunteers. By using the leave-n-out strategy, the deformation fields of the dynamic MR volumes were successfully predicted by using our model (average error is less than 1 voxel). Although we only used the skin motions to build and drive the model, it did not mean the skin surface was the only surrogate source possible with our

model. Without any modification, the model could be adapted to more parameters due to the nature of the CCA method.

Similar to Khamene's method, our model was built from 3D dynamic MR volumes while many other reported methods involved the uses of cine mode CT [1][3][6]. CT volumes acquired in cine mode often contain discontinuities in the data between adjacent couch position due to sorting errors and inter-cycle variation. Extra procedures are needed to correct the artifacts caused by these discontinuities [3]. By using an ultra-fast MR sequence, it is possible to acquire the whole 3D anatomy of the lung from one scan. Furthermore, by using MRI, it is possible to conduct a relatively long scan. 30 seconds MR data were acquired from each of the datasets in this study. Compared to other studies (24 seconds in [2], 20 seconds in [6], 11 seconds in [1] and one respiratory cycle in [3]), our data cover more respiratory cycles sampling more inter-cycle variation. Potentially, more data can be acquired with MR if required. But for CT studies, it is difficult to increase the data acquisition time due to the radiation dose limit.

Many of the methods mentioned in the introduction use polynomial or cyclic function to model the respiratory cycles. By using these methods, a presumption was made that the end positions of inspiration and expiration remained the same over time. However, this presumption is obviously wrong as the end respiratory positions can vary more than 10mm (fig.3) in reality. Our method did not make such presumption. It modeled the whole respiratory signals including the inter-cycle variation. In this study, our models were built from 48 dynamic MR volumes and used to predict the internal organ motions shown in the other 12 MR volumes. Reasonable results were achieved. Without the radiation dose limit, we can acquire more MR volumes and potentially can build a even more accurate model. Zhang et al. reported a method to build the respiratory motion model by using PCA [3]. Similar to CCA model, PCA model can also model the inter-cycle variation of the respiratory motion. But the use of CT has efficiently limited the acquisition time. In fact, Zhang's CT data only covered one respiratory cycle.

Although our approach has many unique advantages and the uses of MRI to plan RT procedures are beneficial (see Introduction), we have not yet evaluated our model on clinical data. Therefore, the clinical accuracy of our model remains unknown. Having an efficient and successful respiratory motion model is a big step toward. But it is still a significant challenge to put this model into clinical use. Other known issues such as tumour baseline variation and morphology changes [11] should be accounted for in any clinical system and we are exploring how our model may be quickly and accurately updated using interfraction imaging (kV or MV fluoroscopy or cone-beam CT).

References

1.   Low, D. A, Nystrom, M, Kalinin E, et al., "A method for the reconstruction of four-dimensional synchronized CT scans acquired during free breathing," Med. Phys. 30, 1254-1263 (2003).

2. Khamene, A, Warzelhan, J K, Vogt ,S, et al., "Characterization of internal organ motion using skin marker positions," MICCAI 2004, 526-533 (2004).
3. Zhang, Q H., Pevsner, A, Hertanto, A, et al., "A patient-specific respiratory model of anatomical motion for radiation treatment planning", Med. Phys. 34, 4772 (2007).
4. Blackall, J M, Ahmad, S, Miquel, M E, et al., "MRI Based Measurement of Respiratory Motion Variability and Assessment of Imaging Strategies for Radiotherapy Planning," *Phys. Med. Biol.* 51, 4147-4169 (2006).
5. Zeng, R, Fessler, J A, and Balter, J M, "Estimating 3-D respiratory motion from orbiting views by tomographic image registration," *IEEE Trans. Med. Imaging* 26, 153-163 (2007).
6. McClelland, J R, Chandler, A G, Blackall, J M, et al., "A Continuous 4D Motion Model from Multiple Respiratory Cycles for Use in Lung Radiotherapy", Med. Phy., 33, 3348-3358 (2006)
7. Hotelling, H., "Relations between two sets of variates", Biometrika, 28, 321–377 (1936)
8. Worsley, K., J., Poline, J-B, Friston, K., J., Evans, A., C., "Characterizing the response of PET and fMRI data using multivariate linear models", Neuroimage, 6, 305–319 (1997).
9. Friman, O., Carlsson, J., Lundberg, P., et al., "Detection of neural activity in functional MRI using canonical correlation analsyis", Magnetic Resonance Medicine, 45, 323-330 (2001).
10. Rueckert, D, Sonoda, L I, Hayes, C, et al., "Nonrigid registration using free-form deformations: application to breast MR images," *IEEE Trans. Med. Imaging* 18, 712–721 (1999).
11. Sonke, J J, Lebesque, J, and van Herk, M, "Variability of four-dimensional computed tomography patient models," *Int. J. Radiat. Oncol., Biol., Phys.* 70, 590-598 (2008).

Integral Geometry Descriptors for Characterizing Emphysema and Lung Fibrosis in HRCT Images

Michal Charemza[1], Elke Thönnes[1,2], Abhir Bhalerao[3], and David Parr[4]

[1] Centre for Scientific Computing, University of Warwick, UK,
m.t.charemza@warwick.ac.uk,
[2] Department of Statistics, University of Warwick, UK, e.thonnes@warwick.ac.uk,
[3] Department of Computer Science, University of Warwick, UK,
abhir.bhalerao@dcs.warwick.ac.uk,
[4] Department of Radiology, University Hospital Coventry and Warwickshire NHS
Trust, Coventry, UK, david.parr@uhcw.nhs.uk

Abstract. Integral geometry descriptors are used to characterize local textural properties of lung parenchyma from HRCT images. These quantities, known as Minkowski functionals, describe the morphology and topology (connectivity) of 2D and 3D binary structures. They have been shown to be effective in describing properties of complex and disordered media, such as open foams. In this paper we describe the Minkowski functionals and propose their use for detecting and grading emphysema and fibrosis in HRCT images. We show which relevant combinations of the four quantities correlate well with three degrees of severity of emphysema and states of fibrosis. We present some illustrative results and make proposals for the use of these descriptors in a larger validation study.

1   Introduction

Emphysema and pulmonary fibrosis are common respiratory disorders which destroy the lung and reduce its ability to oxygenate blood. The primary cause is smoking, but other more subtle risk factors, such as genetic predisposition, are known to be involved. In terms of health care costs, it ranks among the top five western world diseases. The development of robust quantitative methods for its early diagnosis, monitoring of its treatment and understanding of its disease process is therefore of great importance.

High Resolution Computed Tomography (HRCT) is used regularly to assess lung function and structure, and has become an essential tool in the detection and assessment of emphysema and related conditions such as lung fibrosis. The disease is characterized by the destruction of the lung's alveolar sacs and the collapse of their walls which results in localized fibrosis. This leads to a corresponding loss of the respiratory capacity of the lung.

In HRCT images, the loss of lung tissue manifests itself as a reduction of the mean lung density, which can be measured. One simple way is to locate

lung regions below a given low threshold value (e.g. -910 Hounsfield Units) for patients holding their breath at full lung capacity [1]. Histogram based analysis has also been proposed to better determine a cut-off threshold. Nevertheless, the use of a single, global density index has been shown to be unreliable during early or mild stages of the disease, or when the emphysematic regions are dispersed in normal parenchyma [2]. Also, when the pathologic process is mixed, such as in the presence of inflammation and fibrosis, the local density can actually increase. Emphysema and fibrosis can create complex patterns of the lung, whose radiological appearance is commonly known as ground-glass opacities (GGO), honeycombing (HC) and irregular linear or reticular infiltrates [3].

Since the appearance of the disease progression is *textural*, it is unsurprising that image texture analysis techniques have been applied to the problem of robust quantification [2, 4, 1, 5, 3]. The early work of Uppaluril et al. [2] used first and second-order texture features together with fractal dimension (see [6]) to characterize emphysema from CT images. Using pattern recognition, they performed feature selection and trained a classifier to achieve an accuracy of approximately 90% and showed it to be significantly better than mean lung density (MLD) and histogram analysis. Chabat et al. [4] proposed the use of grey-level co-occurrence matrix (GLCM) texture measures to produce a 13 dimensional feature vector containing common features of the GLCM, such as energy and entropy. They also included a number of local shape features by considering the number of primitives at a given grey level (quantized to 16), of a predefined radius. These were then used to estimate the "emphasis", size (short and long) and uniformity, which would capture the connectivity and size of the primitives. This approach is empirical but could be interpreted in the integral geometry framework being presented here. The use of GLCM and the related shape and connectivity measures suggested by Chabat are used again in more recent work of Xu et al. [5]. This time, 3D voxel neighbourhoods were considered and the classifier was set up to grade the pathology into severe, mild and normal lung appearance. Across 34 subjects, they achieved significantly better discrimination over an equivalent 2D approach. In reference [1] and recently Zavaletta et al. [3], spatial maps of lungs have been produced which graphically show the results of the texture classification and can be visually related to the degenerative states of the disease. The type of classifiers used vary from simpler kNN ones, as used in [3], to Bayesian [5] and to neural network approaches [7].

Other work which is of relevance is the study of distributions of air-spaces in the parenchyma and their progression [8], something which could be readily estimated by integral geometry descriptors. Notably, in-vitro histopathological analysis of lung architecture is reported by Ochs et al. [9]. They calculated Euler-Poincaré characteristics from histological samples and were able calculate the total number of alveoli in the lung based on an estimate of the shape coefficient of the alveolus. In this paper, we propose the use of integral geometry descriptors: the Minkowski functionals, with the aim of producing calibrated maps of the emphysema and related fibrosis, localizing and grading the lung parenchyma into several stages from normal to severe. After a brief introduction to the math-

ematical foundations behind Minkowski functionals, we present experiments on synthetically generated data (based on a binary point-process model), and axial HRCT slices of a subject. We discuss the results finding and make proposals for further work.

2   Integral Geometry Descriptors

Healthy and emphysematic lung tissue present textures of distinct morphological properties. To describe these we need geometrical and topological characteristics that describe not only volume and shape but also connectivity. Integral geometry provides a suitable family of such descriptors: *Minkowski functionals or intrinsic volumes*. In 3D there are four such functionals and these are proportional to more commonly known quantities such as volume, surface area, mean breadth and the Euler-Poincaré characteristic.

Intrinsic volumes have been successfully used in material science to characterize and discriminate morphology of various media [10, 11]. In material science the structure of interest is usually homogeneous. In this paper we are interested in characterizing a medium that is composed of various locally homogeneous but texturally differing regions. To do so we compute Minkowski functionals for equal-sized regions of voxels rather than the whole image.

Minkowski functionals are unbiased, stereological estimators and provide local and global morphological information. In contrast to more standard characterisation methods such as pair-correlation functions or cord-length distributions, they incorporate information from higher order correlations. Minkowski functionals have nice mathematical properties such as C-additivity, see (4), which means that they can be computed efficiently based on simple configuration counts.

Mathematically, the Minkowski functionals are defined for a convex, compact set $K \subset \mathbb{R}^3$ via Steiner's formula. Let $K \oplus B_r$ be the dilation of the set K by a closed ball of radius r centered on the origin. Then the volume V of $K \oplus B_r$ can be written as a polynomial function of r as follows:

$$V(K \oplus B_r) \quad = \quad \sum_{k=0}^{3} \binom{3}{k} W_k(K) r^k. \tag{1}$$

Here W_k is the kth Minkowski functional. For example, if C is a cube of side-length a then

$$V(C \oplus B_r) \quad = \quad a^3 + 6a^2 r + 3a\pi r^2 + \frac{4\pi}{3} r^3 \tag{2}$$

and so $W_0(C) = a^3, W_1(C) = 2a^2, W_2(C) = a\pi$ and $W_3(C) = 4\pi/3$. This also illustrates the relation of between the Minkowski functionals and the more common descriptors volume V, surface area S, mean breadth B and Euler-

Poincaré characteristic χ:

$$W_0(K) = V(K), \qquad W_1(K) = \frac{1}{3}S(K), \qquad (3)$$

$$W_2(K) = \frac{2}{3}\pi B(K), \qquad W_3(K) = \frac{4}{3}\pi\chi(K).$$

As the reader is more likely to be familiar with the latter quantities, we will present our results in terms of V, S, B and χ. The Minkowski functionals are proportional to these quantities and so we simply refer to these descriptors as Minkowski functionals.

3 Method

To compute a Minkowski functional for a binary image their definition is extended from compact convex sets to sets that are finite unions of such. This is done exploiting C-additivity of Minkowski functionals, that is, for compact convex K_1 and K_2 we have

$$W_k(K_1 \cup K_2) \quad = \quad W_k(K_1) + W_k(K_2) - W_k(K_1 \cap K_2). \qquad (4)$$

Further their definition is extended to the interior A_n° of a compact n-dimensional geometrical object A_n in 3D space as follows [12]:

$$W_k(A_n^\circ) \quad = \quad (-1)^{3+n+k}W_k(A). \qquad (5)$$

These extensions are required as each voxel is assumed to be a cube of edge length 1, and the image is defined to be a number of white voxels in regular lattice of a black background. Each white voxel is subdivided into its component geometrical objects, the interiors of which are disjoint. Using (4) and (5), the computation of the Minkowski functionals now reduces to the problem of counting these objects: i.e. finding the numbers of open cubes n_3, open faces n_2, open edges n_1, and vertices n_0 ensuring that for adjacent white voxels each shared geometrical object is only counted once.

By repeated application of (4) and (5), and using the known values of the Minkowski functionals for a cube, face, edge and vertex, the Minkowski functionals can then be shown to be [12]:

$$V = n_3,$$
$$S = -6n_3 - 2n_2,$$
$$B = \frac{3}{2}n_3 - n_2 + \frac{3}{2}n_1,$$
$$\chi = -n_3 + n_2 - n_1 + n_0. \qquad (6)$$

For our data, to find n_3, n_2, n_1 and n_0, we use the method in [13], referred to as the *algorithm of equations*. A black voxel does not give any contribution to the sums of geometrical objects. Given a white voxel we find its contribution to

the total of geometrical objects, not counting any shared geometrical objects of
thirteen of the twenty-six adjacent voxels. Indeed, Blasquez and Poiraudean [13]
refer to these voxels as the *preceding* voxels, however their contributions are
independent of the order in which they are found.

4   Experiments

To analyse a binary image, the image was split into a number of identical cuboids
of a given size in order to describe the heterogeneity of the data. Each cuboid
was then embedded in a black background, to ensure the boundary geometrical
objects are all counted in a consistent manner. Then the Minkowski functionals
were found for each cuboid and maps of the Minkowski functionals were pro-
duced. We investigated the correlation of the 4 quantities by principal component
analysis (PCA) and produced grey-scale maps of the principal variation and the
first three principal components mapped to the green, red and blue channels
respectively.

4.1   Synthetic Data

We first use the described method to test whether Minkowski functionals can
differentiate between idealized, and fairly crude, models of diseased and healthy
lung tissue.

   The models are sampled in a window of size $320 \times 500 \times 30$ containing 3
partially overlapping regions of equal overlaps. The first two models are Boolean
models [14] which use a random point process to place different structuring
primitives or grains, see Figure 1. The third model is a Voronoi tesselation [14].
In a Boolean model the locations of grains are independent and identically dis-
tributed uniformly on the region. For the healty tissue, the grains are $2 \times 2 \times 2$
cubes and 3500 grains are placed in total. For the fibrotic tissue, 4000 points
are chosen, and at each point is placed a cuboidal segment. The segments have
lengths that are uniformly distributed on the interval $[15, 20]$, and gradients that
vary between ± 0.8. Their thickness is fixed and is 5 in the z-axis and 2 in the
xy-plane. Finally, for the emphysematic (honeycombed) tissue 800 points are
placed at random and the corresponding Voronoi tesselation is computed. The
thickness of the walls of each Voronoi cell is chosen to be one voxel. Figure 1
shows slice 15 of the synthesized volume.

   We produced a scatter plot of the 4 functionals projected onto the first three
principal components of the covariance of the output, see Figure 2. Data from
the five vertical bands of the synthetic image are plotted and clearly show the
distribution in the Minkowski sub-space and good separation of the different
structures. We used the same dimensions to map the descriptors onto the green,
blue and red colour channels to produce the descriptor map in Figure 1 (bottom-
right).
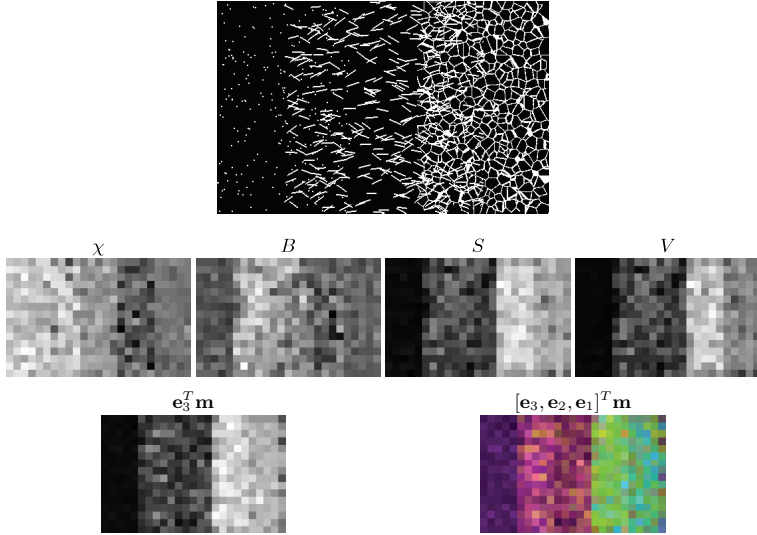
Fig. 1. Minkowski functionals found on synthesized data generated using Boolean and Voronoi models. Top: Slice 15 of synthesized volume size $320 \times 500 \times 30$ (parameters of models are given in section 4.1). Second row: Minkowski functionals found on voxel regions of size $20 \times 20 \times 10$ from slices 11 to 20. Bottom row: Projection of functionals onto principal mode of variation $\mathbf{e}_3$ from PCA of data $\mathbf{m} = (V, S, B, \chi)^T$; RGB mapping of three principal modes of variation.

4.2 HRCT Lung Images

An axial HRCT scan of a subject with lung fibrosis was acquired using a GE Lightspeed Plus CT scanner at a tube current of 150mA. The slice images were of size 512×512 with in in plane resolution of 0.684×0.684 mm and a slice thickness of 0.625 (Figure 3). We calculated the Minkowski functionals on thresholded lung data, using a threshold of -600 Hounsfield units, and taking voxel regions of size $6 \times 6 \times 4$.

In order to visualize the results, we performed principal component analysis on the output from axial slices 29 to 33, and produced colour descriptor maps by mapping the 4 measures on to a 3D sub-space spanned by the first three principal axes of the PCA. Illustrative results for three slices are shown in Figure 4. The colours are, approximately: normal tissue appears as green; fibrosis as a blue to purple hue and honey-combing as brown and red.
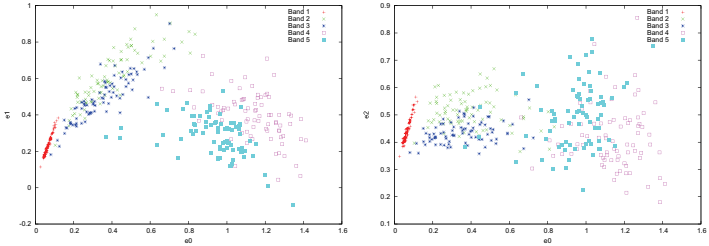
Fig. 2. Scatter plot of measures from vertical bands of synthetic image projected onto pairs of the first thee principal components of variation, $[\mathbf{e}_0, \mathbf{e}_1]$ and $[\mathbf{e}_0, \mathbf{e}_2]$. Bands are numbered from left-to right (with respect to images in Figure 1).
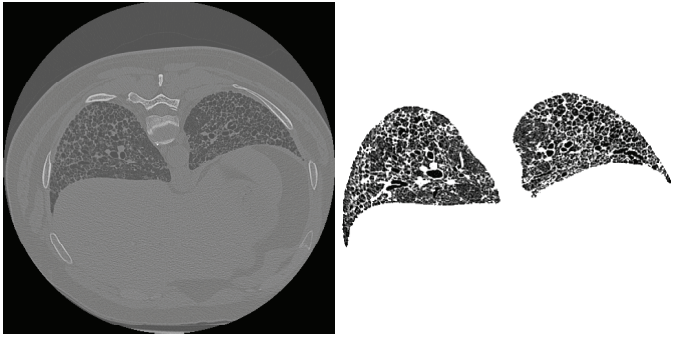
Fig. 3. Original and close-up of 'stripped' data of axial slice 40, size 512×512, with in-plane voxel resolution of 0.684×0.684mm. This slice clearly shows the honeycombing artefacts and fibrosis.

5   Discussion and Conclusions

The results of Figure 4 show that Minkowski functionals can discriminate between different stages of degeneration in emphysematic/fibrotic lung. Minkowski functionals not only have the advantage of being descriptors that have a clear geometric/topological interpretation, see Table 1, but are also easy and fast to compute. They are fully 3D descriptors that go beyond simple local lung density estimation.

   After submitting this paper we discovered concurrent but independent work on the use of Minkowski functionals for grading of emphysematic/fibrotic lung by Boehm et al. [15, 16]. The authors select specific volumes of interest that have been labelled by experts as either normal, emphysematic or fibrotic lung tissue. The Minkowski functionals of these volumes are then computed as a function of

Table 1. Interpretation of Minkowski functionals for homogeneous structures. If W is the volume of the sampling window then $V_v = V/W, S_v = S/W, B_v = B/W$.

Specific fibre length (edge density) L_V	$B_v/\pi(1 - V_v)$
Number of particles	Isolated particles - χ
Mean thickness of edges	$S_v/\pi L_v$
Mean section area of edges	V_v/L_v
Porosity (volume fraction) p	$1 - V_v$

the threshold value used in the binarisation of the image. Using an integrative filtering procedure a single numerical quantity is produced that then can be used for classification. As in our study, the authors show that Minkowski functionals are effective tools for the grading of emphysematic/fibrotic lung and provide information that differs from densitometric measures. In contrast to Boehm et al. we do not classify volumes of interest, but rather aim at the segmentation of the whole lung into the various stages of degeneration. This is achieved by using localized versions of the Minkowski functionals that can then be displayed as spatial maps, see Figure 4. In further work we are planning to examine distribution of these localized functionals for the various stages of pathology. Modelling the spatial distribution of the localized topological measures as a mixture distribution then provides the basis of segmentation of the lung.

In Section 4.1 we present a crude model for some of the structures that may be encountered in diseased lung. More advanced probabilistic model development is needed as important groundwork for statistical inference and the analysis of longitudinal data. The computation of Minkowski functionals is a first step towards model fitting as explicit expressions of these functionals are known for various models, for example the Boolean model.

While the promise of the proposed descriptors is clearly illustrated with our results, supervised validation on a larger survey is needed and planned as further work. In this paper, a simple PCA illustrates how the Minkowski functionals differentiate between textures. Further work will exploit these characteristics in more sophisticated method for feature selection [1] as well as classification and learning [7].

Acknowledgments

References

1. Hoffman, E.A., Reinhardt, J.M., Sonka, M., Simon, B.A., Guo, J., Saba, O., Chon, D., Samrah, S., Shikata, J., Tschirren, J., Palagyi, K., Beck, K., McLennan, G.: Characterization of the Interstitial Lung Diseases via Density-Based and Texture-Based Analysis of Computed Tomography Images of Lung Structure and Function. Academic Radiology **10** (2003) 1104–1118
2. Uppaluri, R., Theophano, M., Sonka, M., Hoffman, E.A., McLennan, G.: Quantification of Pulmonary Emphysema from Lung Computed Tomography Images. American Journal of Respiratory and Critical Care Medicine **156** (1997) 248–254
3. Zavaletta, V.A., Bartholmai, B.J., Robb, R.A.: High resolution multidetector CT-aided tissue analysis and quantification of lung fibrosis. Academic Radiology **7** (2007) 769–771
4. Chabat, F., Yang, G., Hansell, D.: Obstructive lung diseases: Texture classification for differentiation at CT. Radiology **228** (2003) 871–877
5. Xu, Y., Sonka, M., McLennan, G., Hoffman, E.A.: MDCT-based 3-D texture classification of emphysema and early smoking related lung pathologies. IEEE Transaction on Medical Imaging **25** (2006) 464–475
6. Shimizu, K., Johkoh, T., Ikezoe, J., Ichikado, K., Arisawa, J., Nakamura, H., Tamura, S., Nagareda, T.: Fractal analysis for classification of ground-glass opacity on high-resolution CT: an in vitro study. Journal of Computer Assisted Tomography **21**(6) (1997) 955–961
7. Friman, O., Borga, M., Lundberg, M., Tylén, U., Knutsson, H.: Recognizing Emphysema - A Neural Network Approach. In: Proceedings of 16th Int. Conf. on Pattern Recognition - ICPR'02. (2002) 10512
8. Nakano, Y., Saiki, H., Muro, S., Hirai, T., Oku, Y., Nishimura, K., Mishima, M.: Comparison of low attenuation areas on computed tomographic scans between inner and outer segments of the lung in patients with chronic obstructive pulmonary disease: incidence and contribution to lung function. Thorax **54** (1999) 384–389
9. Ochs, M., Nyengaard, J.R., Jung, A., Knudsen, L., Voigt, M., Wahlers, T., Richter, J., Gundersen, H.J.G.: The Number of Alveoli in the Human Lung. American Journal of Respiratory Care Medicine **169** (2004) 102–124
10. Ohser, J., Mücklich, F.: Statistical Analysis of Microstructures in Materials Science. Statistics in Practice. John Wiley & Sons, New York (2000)
11. Arns, C., Knackstedt, M., Mecke, K. Lecture Notes in Physics 600. In: Charactersing the Morphology of Disordered Materials. Springer (2002)
12. Michielsen, K., De Raedt, H.: Integral-geometry morphological image analysis. Physics Reports **347** (July 2001) 461–538
13. Blasquez, I., Poiraudeau, J.F.: Efficient processing of Minkowski functionals on a 3D binary image using binary decision diagrams. Journal of MSCG **11**(1) (2003)
14. Stoyan, D., Kendall, W., Mecke, J.: Stochastic Geometry and its Application. Probability and Statistics. John Wiley & Sons, New York (1995)
15. Boem, H., Fink, C., Becker, C., Reiser, M.: Automated Characterization of Normal and Pathologic Lung Tissue by Topological Texture Analysis of Multi-Detector CT. In Giger, Karssemeijer, eds.: Medical Imaging 2007: Computer-Aided Diagnosis, Proceedings of SPIE Vol 6514. (2007) DOI: 10.117/12.702697
16. Boehm, H., Fink, C., Attenberger, U., Becker, C., Behr, J., Reiser, M.: Automated classification of normal and pathologic pulmonary tissue by topological texture features extracted from multi-detector CT in 3D. European Radiology **18** (2008) DOI: 10.1007/s00330-008-1082-y.
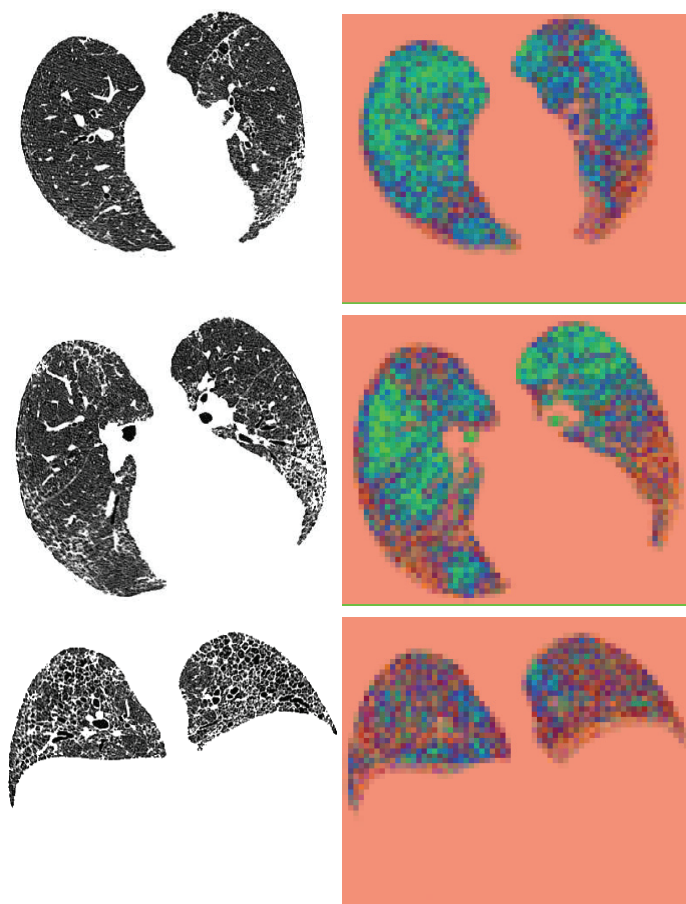
Fig. 4. Results on axial HRCT slices of subject with severe pathology. Left column: original HRCT data. Right column: Results of PCA projection of first 3 principal modes of variation from covariance analysis of slice 30 (second row). Slices 20, 30 and 40 shown. Normal tissue colours as green; fibrosis colours as a blue-purple hue; honey-combing as brown and red.

Automatic Detection and Segmentation of Large Lung Cancers from Chest CT Images

Jinghao Zhou[1], Sukmoon Chang[1,2], Qingshan Liu[1], Dimitris N. Metaxas[1], Binsheng Zhao[3], Michelle S. Ginsberg[3], and Lawrence H. Schwartz[3]

[1] CBIM, Rutgers, The State University of New Jersey, NJ, USA,
jhzhou@eden.rutgers.edu, {qsliu, dnm}@cs.rutgers.edu
[2] Computer Science, Capital College, Penn State University, Middletown, USA,
sukmoon@psu.edu
[3] Department of Radiology, Memorial Sloan-Kettering Cancer Center, NY, USA
{zhaob, schwartl, ginsberm}@mskcc.org

Abstract. Response assessment is critical for cancer patient management and new drug approval. Traditional methods to assess the response are based on measuring tumor size changes in one or two dimensions on computed tomography (CT) before and after therapy, and can be biased. In order to investigate if changes in tumor volume can better assess therapy response, there is an urgent need to develop accurate and reproducible computer-aided tools. Automatic detection and segmentation of lung cancers is a difficult task as lung cancers are often large in size, irregular in shape, and can grow against surrounding structures of similar density and intensity. In this paper, we propose a novel method for automatic segmentation of lung areas that can be distorted by large lung cancers using robust active shape models. We also propose a novel method for automatic detection and segmentation of large lung cancers using a supervised learning framework followed by the analysis of 3D texture likelihood maps. Finally, we present promising results of our methods applied to different clinical applications. The proposed computer-aided methods may provide a new powerful tool for accurate and reproducible quantification of tumor volumes in lung cancer clinical trials.

1 Introduction

The evaluation of therapy response is critical for determining whether a particular treatment is effective on a specific cancer type in a patient. Traditionally, the ways to assess the response are based on measuring size changes of cancer in a transverse image using computed tomography (CT) before and after a treatment [1, 2]. However, the traditional uni-dimensional (maximal diameter of tumor) and bi-dimensional (product of maximal diameter and its perpendicular maximal diameter) measurements can be biased especially when a tumor is not spherical in shape and does not change its shape in a spherical fashion. The preliminary result in a lung cancer study [3] showed that the changes in

tumor volume could be determined as early as 3 weeks after a novel chemotherapy, whereas the changes of tumor volume measured in the traditional methods were significantly less sensitive in the same time period. In addition, manual delineation of tumor contours is time-consuming and lacks the reproducibility. Therefore, there is an urgent need for automatic detection and accurate segmentation methods for the volumetric assessment of therapy response.

Unlike small lung nodules, lung cancers to be treated are often large in size, present spiculate edges, and grow against surrounding structures such as the chest wall, the mediastinum, and blood vessels, which make automatic detection and segmentation difficult [3]. Thus, the algorithms developed for automatic detection and segmentation of small solid lung nodules are most likely to fail when applied to large lung cancers [4–13]. In those studies, larger lung lesions that were attached to the chest wall and mediastinum could be easily and mistakenly excluded from the segmented lungs in which the subsequent lesion detection would be performed [4–7]. Also, the existing segmentation algorithms often assumed that small lung nodules would possess spherical shape, which is not adequate for describing large lung cancers. Furthermore, inability to separate a larger lesion from its surrounding structures of similar intensities was another shortcoming of the existing segmentation algorithms.

In this paper, we propose novel methods for automatic segmentation of lung areas as well as automatic detection and segmentation of large lung cancers from CT images for the purpose of therapy response assessment. We first propose a robust active shape model for the accurate segmentation of lung areas that are distorted and occluded by large lung cancers. Next, we develop a classifier for the detection of cancers in the segmented lung areas by boosting a k-Nearest Neighbor (k-NN) classifier, whose distance measure is the Euclidean distance between the nonparametric density estimates of two regions. The statistical validation of the proposed classifier is also provided. Finally, the classified cancers are automatically segmented by analyzing 3D texture likelihood maps of the surrounding areas. We present the promising experimental results of our method applied to various clinical data. The proposed methods would provide a new powerful tool for automatic detection as well as accurate and reproducible segmentation of lung cancers for therapy response assessment in lung cancers.

2   Method

2.1   Robust active shape models (RASM) for lung area segmentation

Large lung cancers often grow against surrounding structures, such as the chest wall and mediastinum. Lung areas that are distorted and occluded by such lesions are hard to segment due to the similarity of the intensities between the cancers and the surrounding structures in CT images. In this section, we develop a robust method to accurately segment lung areas occluded by large cancers by improving the active shape model framework.

An active shape model (ASM) represents the shapes of interest as a Point Distribution Model (PDM) [14]. Then, it constructs a shape space as a set of
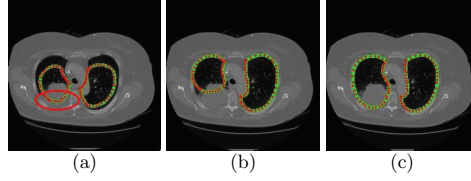
Fig. 1. The segmented lung area using RASM. Red points are active shape model and green lines are the connected contour. (a) Large cancers attached to the chest wall and mediastinum (large red circle) and the initialization of the ASM, (b) ASM finding the false boundary of lung, and (c) RASM finding the correct boundary of lung and the large white area in the left lung is a large lung cancer.

orthogonal basis $\boldsymbol{P}$ by applying the Principal Component Analysis (PCA) and finds an optimal shape for a new example of the shapes with PCA reconstruction. Given the shape space $\boldsymbol{P}$, the projection $\boldsymbol{C}$ of a new example shape $\boldsymbol{X}$ is given as $\boldsymbol{C} = \boldsymbol{P}^T d\boldsymbol{X}$, where $d\boldsymbol{X} = \boldsymbol{X} - \bar{\boldsymbol{S}}$ and $\bar{\boldsymbol{S}}$ is the mean shape from the aligned shapes of the training set. Based on the projection $\boldsymbol{C}$, we can easily find a corresponding shape in the shape space as $\tilde{\boldsymbol{X}} = \boldsymbol{P}\boldsymbol{C} + \bar{\boldsymbol{S}}$. For simplicity, we denoted $d\tilde{\boldsymbol{X}} = \boldsymbol{P}\boldsymbol{C}$. Since $\bar{\boldsymbol{S}}$ is constant, the accuracy of $\tilde{\boldsymbol{X}}$ depends on $\boldsymbol{C}$ which is related to $d\boldsymbol{X}$. In many applications, $d\boldsymbol{X}$ is often optimized with some low-level image features such as the gradient along normal directions to the boundary of an initial shape toward the strongest edge in the image [14].

The ASM method as described above, however, is not suitable for the accurate segmentation of lung areas with large cancers attached on their walls, since the cancers occlude the real boundary of the lung and appear as the strongest edge, as illustrated in Fig. 1(a) and (b). To overcome this difficulty, we develop a robust ASM (RASM) based on the robust M-estimator [15]. The goal is to recover the projection $\boldsymbol{C}$ with the majority of the correct $d\boldsymbol{X}$ and to restrain the outlier points of $d\boldsymbol{X}$. Mathematically, it computes $\boldsymbol{C}$ by minimizing the following robust energy function:

$$\boldsymbol{E}_{rpca}(\boldsymbol{C}) = \min_{\boldsymbol{C}} G\left(\|d\boldsymbol{X} - \boldsymbol{P}\boldsymbol{C}\|, \sigma\right) \tag{1}$$

where, $G(x,\sigma) = x^2/(x^2 + \sigma^2)$ is the Geman-McClure error function and σ is a scale parameter that controls the convexity of the robust function. The solution for $\boldsymbol{C}$ can be obtained by an iterative gradient descent search on $\boldsymbol{E}_{rpca}$:

$$\boldsymbol{C}^{(n+1)} = \boldsymbol{C}^{(n)} + \lambda \Delta \boldsymbol{C} \tag{2}$$

where, λ is a small constant that determines the step size and

$$\Delta \boldsymbol{C} = \frac{\partial \boldsymbol{E}_{rpca}}{\partial \boldsymbol{C}} = -2\boldsymbol{P}(d\boldsymbol{X} - \boldsymbol{P}\boldsymbol{C})\frac{\sigma^2}{(\|d\boldsymbol{X} - \boldsymbol{P}\boldsymbol{C}\|^2 + \sigma^2)^2}$$

(a)	(b)	(c)
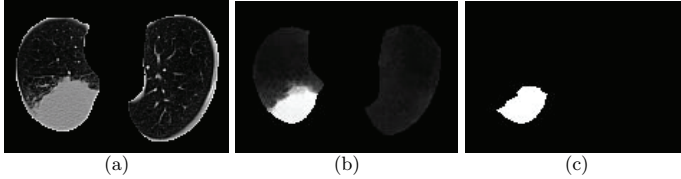
Fig. 2. Effects of F_{cyl} (only one slice from the whole volume is shown). (a) Original volume, (b) Volume filtered with F_{cyl}, and (c) Volume after thresholding.

The iterative process is performed until $\|\boldsymbol{E}_{rpca}^{(t+1)} - \boldsymbol{E}_{rpca}^{(t)}\| < \epsilon$, where ϵ is a pre-selected tolerance. Using the robust projection $\boldsymbol{C}^*$, we obtain a robust shape in the shape space as:

$$\tilde{\boldsymbol{X}} = \boldsymbol{P}\boldsymbol{C}^* + \bar{\boldsymbol{S}}$$

The result of this process is illustrated in Fig. 1(c), where the lung area occluded by a large lesion is accurately segmented.

2.2 Detection of large lung cancers

In this section, we present a novel method for automatic detection of large lung cancers from the segmented lung areas. The method is based on 3D texture analysis using a machine learning framework, i.e., boosting the k-NN classifier. However, the accuracy of the detection may be hindered by various structures within a lung. Thus, we first apply a 3D cylinder filter to suppress the intensity values of vessels and other elongated structures as well as noise inside a lung, while maintaining the intensity values of large lung cancers intact [16, 17]. The cylinder filter F_{cyl} is defined as:

$$F_{\text{cyl}}(x) = \min_{\theta} \left(\min_{y \in \Omega_\theta^x} I(y) \right)$$

where, Ω_θ^x is the domain of the cylinder filter centered at x with orientation θ. F_{cyl} is a hybrid minimum neighborhood filter that produces strong responses to large blob-like objects (e.g., large cancers). In this paper, we have selected the parameters of F_{cyl} empirically and used a cylinder with radii of 1, 2 and 3 voxels and length of 7 voxels at 7 different orientations. In Fig. 2(a) and (b), we can see that vessels and noise are effectively suppressed while the large lung cancers remains intact. After the filtering, we isolate the candidate regions for large lung lesions by simple thresholding (Fig. 2(c)). The threshold value is automatically determined by analyzing the histogram of the filter response image [16]. Each candidate region is then classified with a learning framework described below.

To apply a supervised learning framework, we collected volumetric samples for positive (lesion) and negative (non-lesion) examples manually. Let Ψ_M be the

region of a volumetric sample bounded by a sphere. We estimate the probability density function (pdf) of the intensity values of the interior of Ψ_M. We use a nonparametric kernel based method to approximate the pdf. Let $i \in [0, 255]$ denote the random variable for intensity values. The intensity pdf of Ψ_M is defined as:

$$P(i|\Psi_M) = \frac{1}{V(\Psi_M)} \iiint_{\Psi_M} \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left(-\frac{(i - I(y))^2}{2\sigma^2}\right) dy \qquad (3)$$

where, $V(\Psi_M)$ denotes the volume of Ψ_M, y are the voxels in the domain Ψ_M, and σ is the standard deviation of a Gaussian kernel.

For the candidate areas of large cancers isolated above, the learning for their classification has a discrete target function $f : \mathcal{R}^n \mapsto \{\oplus, \ominus\}$, with the label $\oplus$ for lesions and $\ominus$ for non-lesions. For k-NN, an instance $\boldsymbol{x}$ is represented as a point in $\mathcal{R}^n$ by a feature vector $\langle a_1(\boldsymbol{x}), \cdots, a_n(\boldsymbol{x})\rangle$, where $a_i(\boldsymbol{x}) = P(i|\Psi_M)$. The Euclidean distance is used as the distance measure between two instance vectors. Given a query instance $\boldsymbol{x}_q$ to be classified, k-NN returns $\hat{f}(\boldsymbol{x}_q)$, as its estimate of $f(\boldsymbol{x}_q)$, which is the most common value of f among the k training instances nearest to $\boldsymbol{x}_q$, that is, $\hat{f}(\boldsymbol{x}_q) = \arg\max_{v \in \{\oplus, \ominus\}} \sum_{i=1}^{k} \delta(v, f(\boldsymbol{x}_i))$, where $x_1, \cdots, x_k$ denote the k instances from training samples that are nearest to $\boldsymbol{x}_q$, and $\delta(a, b) = 1$ if $a = b$ and 0 otherwise. To obtain an accurate classification, k-NN requires a large training set, which results in slow classification due to the large number of distance calculations. We overcome this difficulty by boosting k-NN [18]. As in [18], our purpose for boosting k-NN is to improve the speed of k-NN by reducing the number of prototype instances and thus reducing the required number of distance calculation without affecting the error rate.

2.3 Segmentation of large lung cancers

We now segment the classified large lung cancers. Because of the hazy appearance and irregular shape of large lung cancers and the large overlap of intensity values between large lung cancers and surrounding vessels, simple thresholding and contour based segmentation method do not provide accurate segmentation. The proposed method involves the analysis of a 3D texture likelihood map using a nonparametric density estimation [19], followed by eigenanalysis of the Hessian matrix to accurately remove vessels overlapped with large lung lesions.

We extract the region of interest (ROI) surrounding a classified large lung cancer based on the detection of the large lung cancers. For each voxel in the ROI, we evaluate the likelihood of the voxel belonging to a large lung cancer by measuring the 3D texture consistency between the large lung cancer and a small spherical region (i.e., 3D texon) centered at the voxel.

Let Φ_M be the region of a volumetric sample of a classified large lung cancer bounded by a sphere. Using (3), we estimate the pdf of the intensity values of the interior of Φ_M, i.e., $p_M = P(i|\Phi_M)$. Similarly, let Φ_T be the region of the 3D texon centered at the given voxel in the ROI. Again using (3), we also estimate
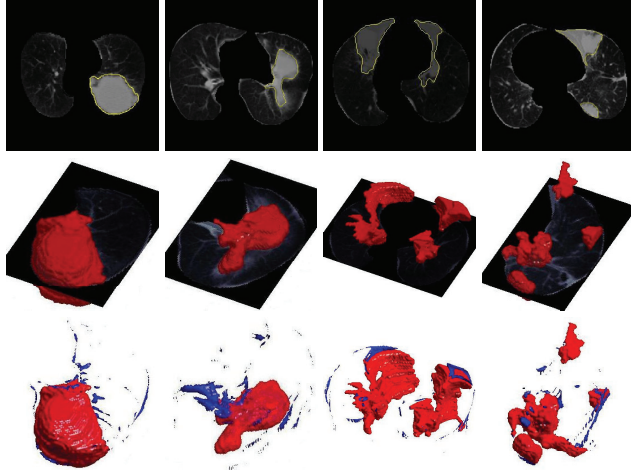
Fig. 3. Results. Segmented large lung cancers projected onto a slice (top), 3D reconstruction of large lung cancers (middle and bottom).

the pdf of the interior of Φ_T, i.e., $p_T = P(i|\Phi_T)$. To measure the similarity between the two pdfs, we use an information theoretic distance measure called Kullback-Leibler Divergence (KLD) [20]. The Bhattacharya distance, which is a symmetrized variation of KLD, between Φ_M and Φ_T is:

$$B(p_M\|p_T) = -\log\rho(p_M\|p_T) = -\log\int[p_M(i)]^{\frac{1}{2}}[p_T(i)]^{\frac{1}{2}}\,di$$

We evaluate the 3D texture likelihood of the 3D texon at every voxel in ROI. We define this likelihood using ρ, since it increases as the Bhattacharya distance between two distributions decreases. The radius of a 3D neighborhood sphere used in our paper is less than 3 voxels and the model interior texture is mostly homogeneous with some level of noise. Thus, it is not necessary to consider the spatial correlation between voxels. Finally, we remove the remaining vessels around large lung cancers in the 3D likelihood map by using the eigenanalysis of the Hessian matrix [21–23]

3   Results

We have 10 chest CT images containing 16 large lung cancers. To test the proposed method, we collected 500 volumetric samples, containing 300 training samples and 200 testing samples, from 4 training clinical chest CT images. The

samples were of size $15 \times 15 \times 3$ voxels from the CT volumes. Each sample was converted to an instance vector in $\mathcal{R}^{256}$, representing its nonparametric density estimate. For the boosted k-NN, we used the standard Euclidean distance as the distance measure between two instances as described in Section 2.2. We performed bootstrapping to estimate the generalization error of our large lesion detection method [24]. We trained and tested the proposed method on bootstrap samples. After 20 steps of boosting, the mean error rate converged to 3.50%.

We applied the trained classifier to all 10 CT volumes containing 16 large lung cancers. The CT volumes were acquired by multi-slice HRCT scanners with 5mm slice collimation. The number of slices in each CT scan ranged from 44 to 69 (and digitally resliced to obtain cubic voxels, resulting in 130 to 205 slices), each of which are of size 512×512 pixels, with in-plane resolution of 0.82mm. The classifier detected all 16 lesions successfully with no false negatives (Fig. 3). However, it also detected 2 false positive lesions, which the trained radiologists classified as atelectases.

The detected large lung cancers were then segmented using the method described in Section 2.3. Fig. 3 illustrates four representative cases of the segmented large lung cancers. In the figure, the 3D reconstruction of the segmented 3D large lung cancers (middle and bottom row) as well as their 2D projections on one of the slices (top row) are shown. From the figure, we can also see that the surrounding vessels are accurately removed from the large lung lesions segmented. Table 1 compares the greatest diameters, their greatest perpendicular diameters and tumor volumes of the 16 lung cancers from the results of the manual segmentation by experts and the automatic segmentation by the proposed method. The table shows that the mean relative error of the greatest diameter and its greatest perpendicular diameter are 2.8% and 2.2% and shows that the mean relative error of the tumor volume is 8.4%. We also compared the overlapping ratios of the tumor regions segmented manually and automatically, which ranged from 80.9% to 97.3%. The low overlapping ratios were resulted from the cases in which the cancers were heavily occluded by blood vessels, where the expert radiologists also found difficulty. The mean overlapping ratio was 90.9%. These results demonstrate the potential of our method to correctly segment occluded lung areas as well as the accuracy of the classification and segmentation of the large lung cancers. These results demonstrate the potential of our method to correctly segment occluded lung areas as well as the accuracy of the classification and segmentation of the large lung cancers.

4   Discussion

Lung cancers to be treated are often large in size and grow against surrounding structures such as chest wall, mediastinum, and blood vessels. Large lung cancers attached to such structures make it difficult to accurately segment lung areas from chest CT images, since they occlude the real boundary of the lungs and have similar intensity values to the surrounding structures. In this paper, we proposed a novel method for automatic and accurate segmentation of lung areas

Table 1. Comparisons. $d1$ and $d2$ are the greatest diameter and its greatest perpen-
dicular distance of each tumor. Vol is the volume of each tumor. Overlap ratio is the
volume overlap ratio of the manual segmentation results and automatic segmentation
results. Manual and Auto are the measurements on the manual segmentation results
and the automatic segmentation results, respectively.

Dataset	d1 (mm)		d2 (mm)		Vol (mm^3)		Overlap Ratio (%)
	Manual	Auto	Manual	Auto	Manual	Auto	
1	95	94	37	36	9676	9043	92.5
2	74	73	35	33	13357	12237	90.2
3	84	83	25	24	12704	11753	91.2
4	34	35	32	33	1643	1819	93.6
5	21	20	16	16	278	254	92.1
6	13	13	11	11	305	291	92.5
7	15	16	13	13	462	503	90.0
8	51	52	18	18	2152	2228	91.6
9	21	20	12	12	258	235	82.2
10	21	20	19	18	1402	1226	97.3
11	7	7	7	6	61	57	80.9
12	68	67	59	58	5988	5583	94.3
13	36	37	27	28	6268	6783	91.4
14	22	23	13	13	291	320	88.3
15	27	26	17	16	657	591	95.6
16	39	38	33	32	4583	4212	91.6
mean error (%)	2.8		2.2		8.4		90.9

that were distorted and occluded by large lung cancers using robust active shape
models.

We also proposed a novel method for the automatic detection and segmen-
tation of large lung cancers from chest CT images. The proposed method first
extracted candidate lung cancer areas by applying the 3D cylinder filter. Then,
each candidate region was classified by boosting the k-NN, whose distance mea-
sure was the Euclidean distance between the two intensity pdfs. We performed
bootstrapping to estimate the generalization error of the method and showed the
mean error rate of the method converged to 3.50%. Each cancer detected was
automatically segmented by analyzing the texture likelihood map of the region.

The very promising results of our methods applied to various clinical chest
CT images were also presented. Although the evaluation of therapy response is
critical for determining whether a particular treatment is effective on a specific
cancer type in a patient, the traditional methods such as uni-dimensional and bi-
dimensional measurements of tumor size are not sensitive enough to accurately
evaluate the changes in tumor volumes. In addition, the manual delineation of
cancer contours is time-consuming and lacks the reproducibility. The proposed
methods provides a new powerful tool for automatic detection as well as accu-

rate and reproducible segmentation of large lung cancers for therapy response assessment in lung cancers.

References

1. Miller, A., Hogestraeten, B., Staquet, M., et al.: Reporting results of cancer treatment. Cancer **47** (1981) 207–214
2. Therasse, P., Arbuck, S., Eisenhauer, E., et al.: New guidelines to evaluate response to treatment in solid tumors. J. Natl. Cancer Inst. **92** (2000) 205–216
3. Zhao, B., et al.: Lung Cancer: Computerized Quantification of Tumor Response—Initial Results. Radiology **241** (2006) 892–898
4. Armato III, S., Giger, M., MacMahon, H.: Automated detection of lung nodules in CT scans: Preliminary results. Medical Physics **28**(8) (2001) 1552–1561
5. Brown, M., et al.: Patient-specific models for lung nodule detection and surveillance in CT images. IEEE Trans. Med. Imag. **20**(12) (2001) 1242–1250
6. Lee, Y., et al.: Automated detection of pulmonary nodules in helical CT images based on an improved template-matching technique. IEEE Trans. Med. Imag. **20**(7) (2001) 595–604
7. Takizawa, H., et al.: Recognition of lung nodules from X-ray CT images using 3D Markov Random Field models. In: Int. Conf. Pattern Recog. Volume 1. (2002) 10099
8. Zhao, B., et al.: Automatic detection of small lung nodules on ct utilizing a local density maximum algorithm. J. Appl. Clin. Med. Phys. **4**(3) (2003) 248–260
9. Zhao, B., et al.: Two-dimensional multi-criterion segmentation of pulmonary nodules on helical ct images. IEEE Trans. Med. Imag. **22**(10) (2003) 1259–1274
10. Boscolo, R., Brown, M., McNitt-Gray, M.: Medical image segmentation with knowledge-guided robust active contours. RadioGraphics **33** (2002) 437–448
11. Ko, J.P., et al.: Small pulmonary nodules: volume measurement at chest ct—phantom study. Radiology **228** (2003) 864–870
12. Kostis, W., et al.: Three-dimensional segmentation and growth rate estimation of small pulmonary nodules in helical ct images. IEEE Trans. Med. Imag. **22**(10) (2003) 1259–1274
13. Chang, S., Zhou, J., Dimitris Metaxas, L.A.: False-positive elimination for computer-aided detection of pulmonary micronodules. In: SPIE Internation Symposium on Medical Imaging. (2006) 1751–1758
14. Cootes, T., Taylor, C., Cooper, D., Graham, J.: Active shape models—their training and application. Comp. Vis. Imag. Under. **61**(1) (1995) 38–59
15. De La Torre, F., Black, M.: A framework for robust subspace learning. Int. J. Comput. Vis. **54**(1-3) (2003) 117–142
16. Chang, S., Emoto, H., Metaxas, D., Axel, L.: Pulmonary micronodule detection from 3d chest ct. In: MICCAI. (2004) 821–828
17. Zhou, J., Chang, S., Metaxas, D., Zhao, B., Schwartz, L.H., Ginsberg, M.S.: Automatic detection and segmentation of ground glass opacity nodules. International conference on Medical Image Computing and Computer Assisted Intervention (MICCAI) (2006) 784–791
18. Freund, Y., Schapire, R.: Experiments with a new boosting algorithm. In: ICML. (1996) 145–156
19. Huang, X., Qian, Z., Huang, R., Metaxas, D.: Deformable-model based textured object segmentation. In: EMMCVPR. (2005) 119–135

20. Ali, A., Silvey, S.: A general class of coefficients of divergence of one distribution from another. J. Roy. Stat. Soc., **28** (1966) 131–142

21. Lorenz, C., et al.: Multi-scale line segmentation with automatic estimation of width, contrast and tangential direction in 2d and 3d medical images. In: CVRMed-MRCAS. (1997) 233–242

22. Zhou, J., Chang, S., Metaxas, D., Axel, L.: Vessel boundary extraction using ridge scan-conversion and the deformable model. In: ISBI. (2006) 189–192

23. Zhou, J., Chang, S., Metaxas, D., Axel, L.: Vascular structure segmentation and bifurcation detection. IEEE International Symposium on Biomedical Imaging (2007) 872–875

24. Efron, B.: Estimating the error rate of a prediction rule: Improvement on cross-validation. J. Amer. Stat. Assoc. **78** (1983) 316–331

Background estimation in CT lung images with applications to perfusion visualisation and lobe separation

Pieter Bruyninckx[1,2], Dirk Loeckx[1], Frederik Maes[1], Dirk Vandermeulen[1], and Paul Suetens[1]

[1] Medical Image Computing, Department of Electrical Engineering (ESAT), K.U.Leuven, Belgium *
[2] `Pieter.Bruyninckx@uz.kuleuven.ac.be`

Abstract. This article proposes a method for estimating the local properties of the lung background or parenchyma, i.e. local mean and standard deviation of the intensities, from CT images. The properties estimating method uses an iterative estimation scheme with outlier detection and exclusion applied in a sliding window. The iterative optimisation method allows for the inclusion of prior knowledge about the background characteristics, thereby increasing the robustness.

Two applications for the algorithm are presented. The first application is the visualisation of the background pattern in the lungs. The second application is the separation of the lungs lobes by using a soft vessel segmentation. The estimated background properties allow the detection of the smaller vessels without the need for a fixed threshold. The background estimation method has been validated on a set of 50 phantom images with and without simulated vessels. The applications have been tested on 3 data sets from 3 different patients. The results have been evaluated by a radiologist. He considers the perfusion method helpful for reading the images and finds the lobe segmentation results acceptable.

1 Introduction

Radiologists have to review an increasing amount of CT images of the thorax, due to increased scanner speed and resolution, and an increasing demand. The evaluation of these images is most often performed by looking at 2D slices, using the standard lung window (from -1300 HU to 300 HU). In order to enhance the quality of the diagnosis and reduce the time spent looking at images, CAR/CAD methods are useful. This article presents two methods for improving the reading of lung images: visualisation of the lung perfusion and of the lung lobes. Both of these methods are based on the estimation of the properties of the lung background. The lung voxels can be subdivided in low intensity background voxels (parenchyma) and higher intensity foreground voxels (mainly vessels).

The visualisation of the lung perfusion is mainly useful for the assessment of mosaic perfusion. Mosaic perfusion is a symptom common in the lungs, (mildly) present in about 80% of all healthy people. It manifests itself in a CT image by inhomogeneous background intensities. There are however multiple factors that can cause mosaic perfusion. [1] Diseases, e.g. pulmonary emphysema or small airways disease, may decrease the blood perfusion in certain areas of the lung, which causes a local decrease in intensity. Gravitation causes a slight increase of blood in the lower parts of the lung, which results in a different intensity pattern according to the patient's orientation (prone vs. supine). Other factors, e.g. the age of the patient and the phase of the respiratory cycle, globally affect the intensities. When the standard lung window is used to view the image it is not easy to perceive the inhomogeneous background pattern.

Herzog et al. have proposed a tool to visualize this background pattern in order to assess pulmonary embolism. [2]. They detect the background by using a global adaptive threshold that preserves the 80% of the lung voxels with the lowest intensities. A smoothed version of these background voxels is shown as a coloured overlay upon the original image.

We propose an improvement computing the properties of the background by using an iterative parameter estimation algorithm with outlier detection and exclusion, applied locally within a sliding window. The parameters under consideration are the mean and standard deviation of the intensities. Due to the iterative procedure this algorithm allows for the incorporation of prior knowledge concerning the lung background intensities.

As a first application of this algorithm the perfusion in the lungs can be visualised by showing the estimated mean as a coloured overlay.

We also propose to use the background knowledge to segment the lung lobes. A segmentation of the lung lobes is useful because it facilitates orienting oneself in the image and because lobe boundaries are often also a border to various pathologies.

Segmenting the lung lobes has traditionally been performed by segmenting the fissures themselves, like in [3]. This is also the method a physician applies when he looks for different lobes in the image. However the fissures are not always visible, or might even be absent [4]. The lobes can also be segmented using the vessel tree, as has been done by [5]. The authors first segmented the vessel tree, and looked for the surface that optimally separated the vessel trees. The exact location of this surface was refined using the fissure information if it was present in the image. Zhou et al. used a similar approach, but they also included the segmentation of the bronchi [6].

We present a lung lobe segmentation using the estimated background properties. First a soft vessel segmentation is performed, by detecting the vessels as outliers to the background. By applying the fuzzy distance transform (FDT) [7] starting from manually indicated points near the hilum the lung lobes are separated using their vessel trees.

The lobe segmentation method presented improves over the previously mentioned methods because it uses the estimated background properties to be robust

in case of mosaic perfusion. Hence it shows the advantage of using a suited background estimator.

The next section presents the methods. A validation of the background estimation method on phantom images and of the applications on real images is presented in Section 3. A discussion of the presented methods and hints for future work in Section 4 conclude this article.

2 Method

The algorithm consists of the following parts. The first step segments the lungs. Next the local background parameters, i.e. the local mean and the local standard deviation of the intensities, of the lung parenchyma are estimated. These parameters are used for two applications, i.e. visualising the lung perfusion and the separation of the lung lobes. All algorithms operate in three dimensions.

2.1 Lung Segmentation

In order to be able to focus on the relevant parts of the image a basic lung segmentation algorithm is applied. At first it selects all air in the image by applying a fixed threshold at -200 HU. The lungs are extracted using a connected component analysis of the whole image. The dominant connected component in the slice at one third from the top in the area $[\frac{1}{4}n_x, \frac{1}{2}n_x] \times [\frac{1}{4}n_y, \frac{3}{4}n_y]$ will denote the air in the lungs, with (n_x, n_y) the slice dimension. The smaller vessels are already included in this volume due to partial volume (PV) effects. To include also the larger vessels, a closing operator using a cuboid of size $10 \times 10 \times 10\text{mm}^3$ is applied. The speed of the lung segmentation is increased by applying it on a subsampled image. The image is subsampled by keeping the first of every two voxels in every dimension, reducing the total image size by a factor $1/8$. If the only interest is visualising the background, the algorithm can be speed up by omitting the closing operation.

2.2 Estimating the Background Properties

The background of the lungs consists of the low intensity (around -900 HU) voxels that make up the parenchyma. The foreground consists of everything else with a higher intensity. The vessels make up most of these higher intensity structures. The background intensities are modeled by a normal distribution, characterised by its mean μ and standard deviation σ.

The background properties μ and σ are computed in a sliding window. For each voxel a surrounding region is selected. The diameter of this region should be chosen larger than the largest vessel inside the lung mask. For this application we have chosen the diameter to be 10 mm. Such a region typically contains only two tissues, i.e. the background and the vessels, so the histogram is expected to consist of two Gaussian peaks. For a healthy lung the background peak is centered at about -900 HU, while the vessel peak is centered at about 200 HU,
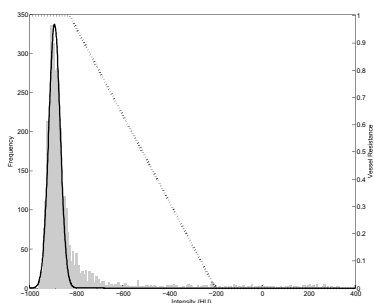
Fig. 1. Histogram of an area of the lung (bars). The full line is the Gaussian fitted to the background peak. The dotted line shows the vessel resistance function defined in (1).

assuming contrast images. However due to partial volume effects, the vessel peak can degenerate into a slope of the background peak. A typical histogram of such a region is shown in Figure 1.

To be able to separate this slope on the right from the background on the left we need to know the properties of the Gaussian at the left, which are in fact the afore mentioned μ and σ. This could be done using a maximum likelihood expectation maximisation (MLEM) that incorporates PV effects like [8].

A more efficient way to get this result is by calculating the mean and variance with outlier detection and exclusion, using an iterative algorithm. First μ and σ are initialised at their default values, i.e. -860 HU and 70 HU. These values are based on the measurement of lung intensities of healthy and pathological lungs. μ and σ are iteratively updated by computing the mean and standard deviation of the voxel intensities in the area that does not deviate more than $n\sigma$ from μ. The estimated standard deviation is biased because it is computed on voxels with a limited intensity range. To prevent a bias on the estimated standard deviation it should be multiplied by 1.13684 for n equal to two, as used in this article. The formula to compute this value is derived in the appendix. The iteration is repeated seven times. Empirical testing has shown that this suffices for convergence on clinical images of the thorax.

This algorithm is speed up by calculating μ and σ on a subgrid, and interpolating the result afterwards. The distance between the voxels included in the subgrid is half the size of the sliding window for each dimension. Since the algorithm acts as a low pass filter this approximation does not cause significant changes to the result. In a pathological image the average absolute differences were 3 HU and 4.3 HU per voxel for μ and σ respectively.

The perfusion is visualised by showing the computed mean μ as an overlay using a rainbow colour map onto the original image. This is shown in Figure 3.

2.3   Lobe Segmentation

The estimated background properties are used to calculate in each voxel a probability-like vessel resistance value f. This value is 1 for background voxels, and 0 for voxels that almost certainly belong to a vessel, while linearly interpolating in between. The following definition has been chosen to incorporate as much image information as possible while minimising the effects of the fissures present in the image.

$$
f = \begin{cases} 1 & x \leqslant \mu + k\sigma \\ 1 - \frac{x - \mu - k\sigma}{-200 - \mu - k\sigma} & x \in (\mu + k\sigma, -200) \\ 0 & x \geqslant -200 \end{cases} \tag{1}
$$

In this equation x is the voxel intensity, while μ and σ are the values computed in section 2.2. The parameter k determines the distance between the centre of the background peak and the start of the slope. This parameter should be chosen to eliminate the background peak from the slope, so a value of $k = 3$ is a safe choice. A lower value might increase the accuracy while a value too low will severely distort the result. For all figures in this paper k is set to 2.3. The end of the slope is fixed at -200 HU, so all voxels with a higher intensity are expected to belong to a vessel, and will have zero vessel resistance. A plot of f is shown in dots on Figure 1.

The lobes are segmented using a FDT of (1) by comparing the distance from multiple starting locations in the image to each lung voxel. The starting locations are placed manually at the start of large vessel trees, near to the hilus.

3   Results

The method for computing the background properties has been validated using a software phantom. The perfusion method as well as the lobe separation have been tested on CT data sets of three different patients, and the results were evaluated by an experienced radiologist.

3.1   Tests on synthetic data

In order to test the accuracy of the background estimation method a set of random phantom images imitating some of the lung's characteristics has been created. The generated phantom images consist of a background with slowly varying intensity (-900 HU to -700 HU, left to right) containing a set of non overlapping vessels of varying size (radius 0.25 voxels to 4.75 voxels, top to bottom). The vessel model is similar to the one presented in [9]. The vessels are modeled as cylinders. The model incorporates partial volume (PV) effects, a Gaussian point spread function (PSF), and different noise characteristics for vessels and background. The PSF is an isotropic Gaussian with standard deviation 0.75 voxels. The foreground (vessel) voxels have a mean intensity of 200 HU and a standard deviation σ_1 of 35 HU, while the background has a standard deviation σ_2 of

(a)   (b)   (c)
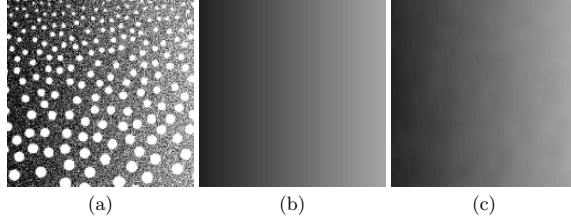
Fig. 2. Example of the calculation of the background properties on a phantom image. Phantom image (a), background ground truth (b), and estimated background (c). All images have a window level from -970 HU to -600 HU.

50 HU. The standard deviation of the PV voxels is $\sigma_\alpha = \sqrt{\alpha\sigma_1^2 + (1-\alpha)\sigma_2^2}$, where α is the vessel/voxel volume ratio for a certain voxel. A slice of one of the created phantom images can be seen in Figure 2(a).

The background estimation algorithm has been applied to a set of 50 phantom images. The difference between the estimated background intensity and the ground truth background was 7.2927 ± 2.5478 HU The difference between the estimated standard deviation of the background and the ground truth standard deviation was 3.9615 ± 1.5929 HU. The small positive bias of the mean is probably due PV effects. Changing the vessel intensity to 0 HU and 400 HU did not result in significantly different values. For comparison the same values for 50 background-only images were 0.0024 ± 0.7518 HU for the mean, and 0.1799 ± 0.7301 HU for the standard deviation. The standard deviations for all these values over all images was smaller than 0.3 HU.

3.2 Background Perfusion

We have applied the perfusion tool to three datasets of three different patients and presented the results to an experienced radiologist. He said the perfusion tool allows to visualise subtle diffusion variations, like those caused by gravitational effects. Without this tool a tedious amount of twiddling with the window level would be needed for these subtleties to be visible, so therefore he thought the tool would be able to increase his productivity. Since the algorithm is fully automated, it can be executed beforehand, thus enabling the radiologist to show the perfusion by using only one slider.

An example of the result of this algorithm is shown in Figure 3. The image on the left shows the original slice using the standard lung window. The image at the right shows the perfusion using a rainbow color map. The image in the center shows the perfusion map as an overlay onto the original image. In the original image a brighter background is visible when one has a good look at the bottom of the lung on the right hand side. In both other pictures this area is visible at first sight without effort. Moreover the right image shows a soft green

(a) original   (b) original with overlay   (c) overlay only
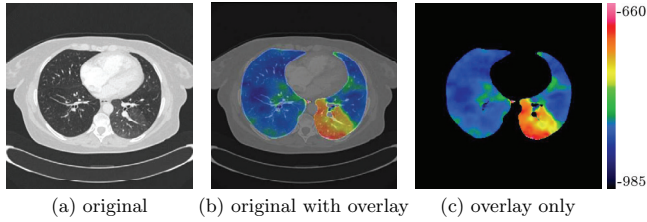
Fig. 3. Demonstration of the background properties. The coloured overlay shows the intensity of the parenchyma, thus allowing to visualise subtle diffusion variations.

(a) original   (b) original with overlay   (c) overlay only
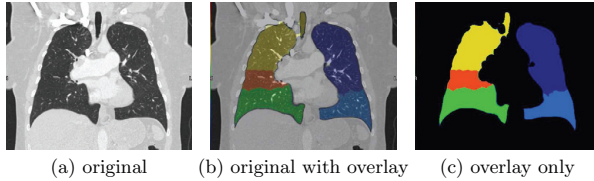
Fig. 4. Demonstration of the lobe segmentation. Different colours indicate different lung lobes.

glowing at the bottom of the lung on the left hand side, indicating a slightly higher perfusion due to gravitational effects.

The algorithm has also been executed with less optimal initial estimates for μ. The initial estimate does not influence the values to wich the iterations converge, since a difference by 100 HU yielded a mean absolute difference of less than 1 HU after fourteen iterations.

3.3 Lobe Segmentation

We have applied the lobe segmentation to three datasets of three different patients and presented the results to an experienced radiologist. According to him the lobe segmentation tool is adequate for clinical use. The most important advantage of this algorithm is that it is complementary to lobe segmentation by fissure detection, which is how the radiologist normally recognizes different lobes.

The algorithm does not always return the optimal result after the first manual initialisation. In this case an iterative procedure, in which additional starting points are added, is performed.

An example of the results returned by this algorithm can be seen in Figure 4. Figure (a) shows the original CT slice, and (c) shows a colour map depicting the different lung lobes. A combination of both is shown in (b). The lobe segmentation has been performed for $k = \{2.0, 2.3, 3.0\}$. The maximum volume

percent difference between the several results was smaller than 0.75%. These small differences between the different results show that the exact value of k is not critical.

3.4   Timing

The algorithms have been implemented in Matlab R2006b with C++ calls for the most time intensive parts. The Matlab image processing toolbox is used for performing the morphological operations of the lung segmentation step. On a Dell Precision workstation with dual Xeon 2.2 GHz/512k processors the lung segmentation took 37 minutes, estimating the background properties took 51 seconds, and separating the lobes took 270 seconds, for an image of $512{\times}512{\times}601$ voxels.

   If one is interested in the perfusion only, and not in the separation of the lung lobes, the efficiency of the lung segmentation step can be improved by omitting the time consuming morphological operations, which account for 36 minutes of processing time.

4   Discussion and Future Work

We have proposed a method for computing the properties of the lung background using a low pass filter with outlier detection and exclusion. The result of this filter can be visualised immediately, showing the lung perfusion. The accuracy of the results allows the physician to see otherwise nearly invisible variances in the background distribution, like those caused by gravitational effects. Showing the result as an overlay results in a simple tool, which was judged by a clinician as very helpful, e.g. because it can speed up the reading. By using the afore-mentioned background properties we have also proposed a lobe segmentation algorithm, which is robust in cases of mosaic perfusion.

   The algorithm has not yet been tried on non-contrast images, but the intensity gap between foreground and background should still be sufficient to yield acceptable results.

   To improve the visualisation of the perfusion an automatic detection of the adequate window level would be useful. A fixed window level is not adequate, since the amount of air in the lungs affects the background intensities. The initialisation of the lobe segmentation method could be automated by using a map like in [3]. Once the lobe segmentation has been performed, the perfusion detection can be recomputed on a per lobe basis, increasing the visibility of the contrast differences at the fissures.

   Our focus for the future will rather be on using the background estimation algorithm in applications, e.g. vessel segmentation, than on improving the algorithm as such. The methodology for segmenting the lung lobes can also be used to separate the arteries and the veins, yielding promising preliminary results.

A   Bias on Estimated Standard Deviation

Let $f(x)$ be a Gaussian probability density function with mean μ and standard deviation σ.

$$f(x) = \frac{1}{\sqrt{2\pi}\sigma} \exp\left(-\frac{(x-\mu)^2}{2\sigma^2}\right) \tag{2}$$

Suppose the mean and the standard deviation of this probability density function are estimated on a limited interval, i.e. in the interval $[\mu - n\sigma, \mu + n\sigma]$, with $n > 0$. It is obvious that the estimate mean μ' will be the same as the real mean μ. The estimated standard deviation σ' however will be biased.

The estimated standard deviation σ' is computed as

$$\sigma' = \sqrt{\frac{\int_{\mu-n\sigma}^{\mu+n\sigma}(x-\mu)^2 f(x)dx}{\int_{\mu-n\sigma}^{\mu+n\sigma} f(x)dx}} \quad . \tag{3}$$

From (3), $\frac{\sigma}{\sigma'}$ can be calculated.

$$\frac{\sigma}{\sigma'} = \left(-\frac{\left(\sqrt{2}n - \mathrm{erf}\left(\frac{\sqrt{2}}{2}n\right)\sqrt{\pi}e^{\frac{n^2}{2}}\right)e^{-\frac{n^2}{2}}}{\sqrt{\pi}\,\mathrm{erf}\left(\frac{\sqrt{2}}{2}n\right)}\right)^{-\frac{1}{2}} \quad . \tag{4}$$

References

1. Webb, W.R., Müller, M.L., Naidich, D.P.: High-Resolution CT of the Lung. 3 edn. Lippincott Williams & Wilkins (2001)
2. Herzog, P., Wildberger, J.E., Niethammer, M., Schaller, S., Schoepf, U.J.: CT perfusion imaging of the lung in pulmonary embolism. Acad Radiol **10**(10) (October 2003) 1132–1146
3. Zhang, L., Hoffman, E.A., Reinhardt, J.M.: Atlas-driven lung lobe segmentation in volumetric X-ray CT images. IEEE Trans Med Imaging **25**(1) (January 2006) 1–16
4. Dandy, W.E.: Incomplete pulmonary interlobar fissure sign. Radiology **128**(1) (July 1978) 21–25
5. Kuhnigk, J.M., Hahn, H., Hindennach, M., Dicken, V., Krass, S., Peitgen, H.O.: Lung lobe segmentation by anatomy-guided 3D watershed transform. Volume 5032., SPIE (2003) 1482–1490
6. Zhou, X., Hayashi, T., Hara, T., Fujita, H., Yokoyama, R., Kiryu, T., Hoshi, H.: Automatic segmentation and recognition of anatomical lung structures from high-resolution chest CT images. Computerized Medical Imaging and Graphics **30**(5) (July 2006) 299–313
7. Saha, P.K., Wehrli, F.W., Gomberg, B.R.: Fuzzy distance transform: Theory, algorithms, and applications. Computer Vision and Image Understanding **86**(3) (June 2002) 171–190
8. Van Leemput, K., Maes, F., Vandermeulen, D., Suetens, P.: A unifying framework for partial volume segmentation of brain MR images. IEEE Trans Med Imaging **22**(1) (January 2003) 105–119

9. Wörz, S., Rohr, K.: Segmentation and quantification of human vessels using a 3-D cylindrical intensity model. IEEE Transactions On Image Processing **16**(8) (August 2007) 1994–2004

A Simultaneous Segmentation and Regularization Framework for Vessel Extraction in CT Images

Gang Song[1], Alonso Ramirez-Manzanares[1,2] and James C. Gee[1]

[1] Penn Image Computing and Science Laboratory (PICSL), Department of
Radiology, University of Pennsylvania, Philadelphia, PA 19104, USA
http://www.picsl.upenn.edu
[2] Univ. de Guanajuato, Facultad de Matematicas, Callejon Jalisco S/N, Valenciana,
Guanajuato, Gto. Mexico. C.P. 36240

Abstract. A new framework is proposed for segmenting the pulmonary vessel tree while simultaneously estimating vessel orientations from lung Computed Tomography (CT) images. The problem is formulated as a joint optimization for both the segmentation and the orientation of the vessel tree. We propose to use a histogram vector to describe vessel orientation, which avoids explicit classification of branching points from a vessel tree. The objective function encodes the orientation information by defining a neighboring relationship between voxels, and is solved iteratively by alternately optimizing the segmentation and regularizing the orientation. The validation on manually labeled datasets suggests the potential value of our algorithm.

1   Introduction

Pulmonary vessel extraction is an important step in performing a quantitative analysis for lung CT images. Blood vessels inside lungs have a salient tree structure over the entire lung. The radii of vessels get decrease as they extend from the center to the periphery of the lung. The vessels typically have a higher intensity than the lung parenchyma. However, it is not sufficient to classify the vessels from the rest of the lung volume only by using the voxel intensities. The imaging process adds noise to the data volume and under current limitations of scanning resolution, the vessel structure in the images may become disconnected. These factors make vessel extraction a challenging problem.

In order to achieve a robust vessel segmentation, researchers have proposed various approaches. One is to preprocess the volume by enhancing the vesselness. For example, Shikata et al ([1]) first enhanced the vessels by a line-filter based on the Hessian matrix, after which the vessel tree was segmented by thresholding the vesselness and filling local gaps. Another method turns the problem into a tracking problem on the vessel tree after getting the initial segmentation using some heuristic thresholding. In [2], the initial estimation of the vessels was obtained from maximum intensity projection, and a rule-based scheme was adopted

to separate the vessel tree by extracting the center lines. Wu et al ([3]) used the method in a more principled way. They proposed a regulated morphology approach to produce a set of fuzzy spheres; a tracking algorithm subsequently generated a set of connected trees based on constraints such as for instance collinearity and size. There are also level-set based approaches such as those in [4][5], which optimize the vessel boundary as the zero level set of a deformable curve model.

In terms of optimization, the variable that all these methods try to optimize is only the segmentation label. Other information, especially the orientation of the vessel tree, is fixed after precomputation from either the Hessian matrix ([1]) or the morphology ([2][3]). In contrast to these previous approaches, we exploit the orientation information in a new way. An accurate estimation of orientation can provide better tracking in segmentation. Meanwhile, a good segmentation eliminates noise and improves the orientation information. In this paper we propose a new method of vessel extraction to compute both the segmentation and the orientation simultaneously.

In our framework, the orientation is regarded as a variable to be optimized rather than a precomputed feature for segmentation. We formulate our approach as an optimization problem and propose an iterative solution. First, a new description of orientation is applied for vessel extraction. We employ a histogram to accommodate multiple directions at tree branching points. This is different from using the direction vector, which is not suitable for describing bifurcation; for instance, in previously reported methods ([1][2]) special tracking rules have to be defined for the branching points. To initialize the orientation response in multiple directions, we design a bank of elongated second-order filters, which can detect more than one dominant direction.

Second, we propose a new formulation to alternately optimize both the orientation information and the segmentation label. We apply the graph cut method in segmentation, which has become a popular method for medical image segmentation ([6][7][8]). This method guarantees the global optimal solution and is computationally efficient. It can model both the data likelihood and the neighboring relationship between two voxels. We introduce a novel term defined in terms of the orientation of neighboring voxels. This leads us to a new way to integrate the tracking and segmentation procedures.

This paper is organized as follows: Section 2 presents our simultaneous segmentation and regularization framework. Section 3 shows both quantitative and qualitative results on High-Resolution Computed Tomography ($HRCT$) ([9]) data. Finally, we give our conclusion in Section 4.

2   Method

The goal of our method is two-fold: to segment the pulmonary vascular vessels, and to estimate the vessel orientation at each voxel. We propose a new framework for optimizing both the segmentation label and the orientation estimation. Given the image volume I, for each voxel i with coordinates s_i, we want to assign a

(a)   (b)

(c)   (d)
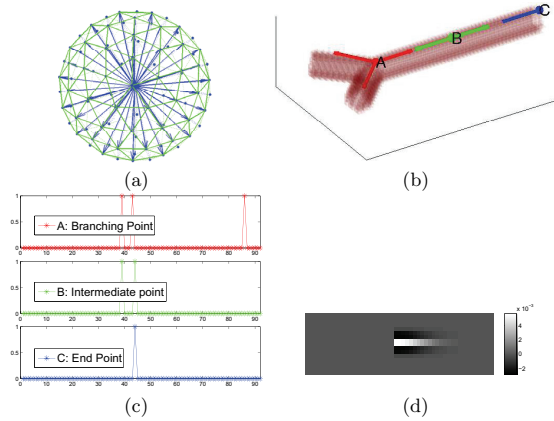
Fig. 1. (a) The 92 predefined directions ψ_k illustrated on a 3D sphere. (b) Three types of points on a Y-junction. (c) The orientation vectors o of for A) branching points, B) intermediate points and C) end points. (d) 2D view of the orientation filter (Equ. (2)).

label l_i indicating whether it belongs to the vessel tree ($l_i = 1$) or the background ($l_i = 0$). Moreover we want to estimate the orientation of the vessel voxel at s_i.

The orientation is described by a histogram vector $\boldsymbol{o}_i = \langle o_i^1, \ldots, o_i^K \rangle$. Each bin in the histogram corresponds to one of the predefined K unit normal vectors, $\{\boldsymbol{\psi}_1, \ldots, \boldsymbol{\psi}_K\}, \|\boldsymbol{\psi}_k\| = 1$. Figure 1(a) shows the 92 directions used in the paper. The value of bin o_i^k is a continuous value in the range of $[0, 1]$, representing the likelihood to the k-th direction of the vessel. Such a histogram vector is capable of describing multiple dominant vessel directions at the same location; thus it can generalize the description of the three types of points in the vascular tree: the branching points with three non-zero bins, the intermediate points with two, and the end points with one. In general $\boldsymbol{o}_i$ is a sparse vector $\in [0, 1]^K$ (see Figure 1(b)(c) for an illustration on a Y-junction).

Such a vector description is different from the one used by most previous methods ([1][2]), in which the orientation is only defined on the intermediate points, and the branching points need to be classified differently from other points in order to initiate vessel tracking. In contrast, our histogram representation is capable of describing orientation information in all parts of the vascular tree with the same format. Therefore, there is no need for classifying and tracking branching points like in [3]. This description is more robust to local noise and broken vessels.

The data term, D_i^S, denotes the negative log-likelihood of assigning label l_i to the pixel i by the intensity at site s_i. $W_{i,j}^S$ is the neighboring connectivity

term referring to the compatibility between i and j. $\mathcal{N}$ is the second order spatial neighborhood of 26-connectivity. D_i^R and $W_{i,j}^R$ are the terms for orientation. W^S and W^R ensure that the segmentation labels and the orientation of the vessels change smoothly along the tree. Our algorithm computes the minimum of the objective function:

$$C(L, \boldsymbol{O}) = \sum_{i \in I} D_i^S + \lambda_S^W \sum_{i \in I} \sum_{l_j \neq l_i | j \in \mathcal{N}_i} W_{i,j}^S + \sum_{i \in I} D_i^R + \lambda_R^W \sum_{i | l_i = 1} \sum_{l_j = 1 | j \in \mathcal{N}_i} W_{i,j}^R. \quad (1)$$

Each term is a function of the segmentation label field $L = \{l_i\}$, and the orientation field $\boldsymbol{O} = \{\boldsymbol{o_i}\}$. The first half of the formulation, $\sum_i D_i^S + \sum_i \sum_j W_{i,j}^S$, is the popular segmentation energy function, which can be efficiently optimized by the min-cut/max-flow algorithm ([6]). In contrast, the variable $\boldsymbol{o_i}$ in the second half, $\sum_i D_i^R + \sum_i \sum_j W_{i,j}^R$, is a continuous variable, which leads us to a different regularization approach as explained in Section 2.3.

2.1 Initial Estimation of Orientation

To locate the multiple dominant directions, we apply a bank of directional filters tuned along the set of directions $\{\boldsymbol{\psi}_k\}$. This is different from using the Hessian matrix ([1]), which is suitable for only one dominant direction. To detect thin structures, we generalize the elongated second-order derivative filter in [10][11] to 3D as:

$$F(\alpha, \beta, \gamma) = \delta_{\tilde{x} \geq 0} \frac{1}{C} \frac{\partial^2}{\partial \tilde{y}^2} \exp\left(\frac{\tilde{x}^2}{\lambda_x \sigma^2} + \frac{\tilde{y}^2}{\sigma^2} + \frac{\tilde{z}^2}{\lambda_z \sigma^2}\right)\Bigg|_{(\tilde{x}, \tilde{y}, \tilde{z})^T = \mathbf{R}(\alpha, \beta, \gamma)(x, y, z)^T}. \quad (2)$$

The direction of the filter is controlled by the Euler angle (α, β, γ). The predefined direction $\boldsymbol{\psi}_k$ determines α and β. The third Euler angle γ comes from n evenly distributed angles from 0 to π. We empirically choose $n = 4$ as a tradeoff between accuracy and running time. $\delta_{\tilde{x} \geq 0}$ makes the filter respond only to the forward direction. σ controls the scale of the filter, λ_x controls the elongated scale along x-axis, and λ_z controls the thickness of the thin structure (see Figure 1(d)). We set $\sigma = 1$, $\lambda_x = 16$ and $\lambda_z = 1$ in our experiments. $\mathbf{R}$ is the rotation matrix defined by (α, β, γ). By rotating the filter in 3D we get a bank of filters $\{F(\alpha, \beta, \gamma)\}$ for different directions. So an image volume needs to be convoluted with $K \times n$ filters. The initial estimation of the orientation bin $\hat{o}^k$ is the maximum response out of the n filters of $\boldsymbol{\psi}_k$, normalized to the range $[0, 1]$ divided by the maximum response of the whole volume.

2.2 Segmenting the Frontier Band

The vascular vessels are typically brighter than the rest of the lung region. However, in the regions towards the boundary of the lung, the vessels become as thin as only one or two voxels. Small motion blur and imaging noise make it difficult to identify these vessels. In contrast to voxel intensity, the orientation information is relatively more robust because the orientation filter typically covers a
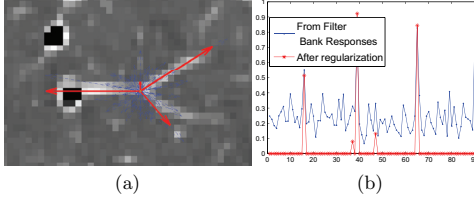
Fig. 2. (a) A patch from HRCT lung volume. *Blue arrows*: the initial orientation $\hat{o}_i$ estimated from convolution with the filter bank. *Red arrows*: the orientation after regularization, o_i. (b) The histogram vector of $\hat{o}_i$ and o_i.
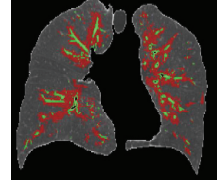
Fig. 3. The active band (red region) and the segmentation (green region) for the first iteration.

large neighborhood. We introduce the orientation into the connectivity term as follows:

$$
\tilde{W}_{i,j}^S = \sum_{k=1}^{K} w_{ij}^k o_i^k, \quad w_{ij}^k = \begin{cases} \langle \boldsymbol{\psi}_k, \boldsymbol{s}_j - \boldsymbol{s}_i \rangle & \text{for} \quad \langle \boldsymbol{\psi}_k, \boldsymbol{s}_j - \boldsymbol{s}_i \rangle \geq t \\ 0 & \text{otherwise} \end{cases}, \quad (3)
$$

where $\langle \cdot, \cdot \rangle$ is the vector dot product and t is a threshold ($t = 0.9$ for our experiments). In this definition, the spatial regularization weight w_{ij}^k is high if the voxel j is along any of the detected directions at voxel i. $W_{i,j}^S$ in Equ. (1) is defined to be symmetric: $W_{i,j}^S = \max(\tilde{W}_{i,j}^S, \tilde{W}_{j,i}^S)$. We set $\lambda_S^W = 3/7$ in Equ. (1).

The data likelihood term is defined as $D_i^S = -\log \Pr(I_i|l_i)$. This likelihood is computed by the local Parzen-window nonparametric density estimation. A small neighborhood around each voxel is sampled by an isotropic Gaussian PDF. The implementation details for optimizing the segmentation labels L given Equ. (3) are illustrated in [8] by using the min-cut/max flow algorithm.

2.3 Regularization of Orientation Vectors

The initial estimation of vessel orientation depicted in Section 2.1 is very noisy. We desire a sparse solution for vector $\boldsymbol{o}_i$, in which each vector should have strong responses only in no more than three bins. Also the estimation should be spatially coherent. For this purpose we apply the multi-modal regularization framework introduced in [12]. The value of $\boldsymbol{o}_i$ is regularized from the initial estimation $\hat{\boldsymbol{o}}_i$. The data term in Equ. (1) is defined as:

$$
D_i^R = \|\boldsymbol{o}_i - \hat{\boldsymbol{o}}_i\|^2 + \kappa(c\,\overline{o}_i^2 - \sum_{k=1}^{K}(o_i^k)^2), \quad \text{in which } \overline{o}_i = \frac{1}{K}\sum_{k=1}^{K} o_i^k. \quad (4)
$$

The first term restricts the regularization results $\boldsymbol{o}_i$ to be close to the input $\hat{\boldsymbol{o}}_i$. The second term enhances predominant orientations and attenuates the spurious ones, while c controls the number of non–zero peaks (see details in [12]).

We set $c = 6$ and $\kappa = 0.5$ in this paper to favor 1 to 3 non-zero different peaks. For the spatial regularization term, $W_{i,j}^R$ in Equ. (1) is defined as: $W_{i,j}^R = \sum_{k=1}^{K} w_{ij}^k (o_i^k - o_j^k)^2$. As normalized bin values in $\boldsymbol{o}_i$ are non-negative real-valued, we minimize this stage by solving a constrained system of linear equations. The details for optimizing the orientation $\{\boldsymbol{o}\}$ with the Gauss-Seidel algorithm are given [12]. We set λ_R^W to be 6. Figure 2(a)(b) shows an example of the regularization results.

2.4 Iterative Optimization of the Objective Function

We propose an iterative approach to alternately optimize the first and the second half of the objective function Equ. (1). In each iteration, we maintain a region called *Active Band* (*AB*). The segmentation and regularization is only applied sequentially inside this region, and the labels outside AB are fixed.

In the initialization stage, AB is obtained by thresholding the response of multiple orientation filters (see Figure 3). Next, orientation is regularized within AB. After the regularization, segmentation labels within AB are computed by the min-cut/max flow algorithm. At the end of each iteration, AB is updated by dilating current vessel tree labels (we use a dilation radius of 2).

The tubular-like structure of AB prevents the segmentation from leaking into the lung parenchyma. This improves computation efficiency without sacrificing accuracy. Such an idea is also adopted in [7]; but in contrast, we do not use Laplacian pyramids to obtain the active band. The convergence point of iteration gives a sub-optimal solution to the objective function of both L and $\boldsymbol{O}$.

3 Results

The two datasets were acquired from one patient with no known pulmonary disease. Using a volumetric expiratory HRCT protocol ([9]), the patient was scanned once at end-inspiration and again at maximal end-expiration with a 4-detector CT scanner (GE Lightspeed, 2.5 mm collimation, 120kVp, 240mA, 0.5s gantry rotation time, 15mm per rotation). Images were reconstructed to 1.25 mm-thick slices with a 512×512 matrix of 0.63-mm in-plane resolution and were further downsampled by half-resolution to $256 \times 256 \times 200$. Before vessel extraction, a rough mask of the whole lung was estimated via semi-automatic level-set segmentation ([13]) to remove extra parts of the bones, heart, and large airways.

Three iterations were run for each volume to obtain convergence. Figure 4(a) shows the final vessel 3D mask on one volume. Most of the extracted vessels are thin (i.e. not over segmented) and connected without manual initialization. We evaluated our segmentation results using a similar approach to that used in [1]. About 3000 points uniformly distributed within vessels were manually labeled for each volume. For verification, the lung region is divided into five distinct regions (see Figure 4(b)) depending on the distance from two seed points manually placed on the border of the left and right lung around the hilum, where most vessels enter the lung region. The peripheral region E contains mostly thin

(a)   (b)
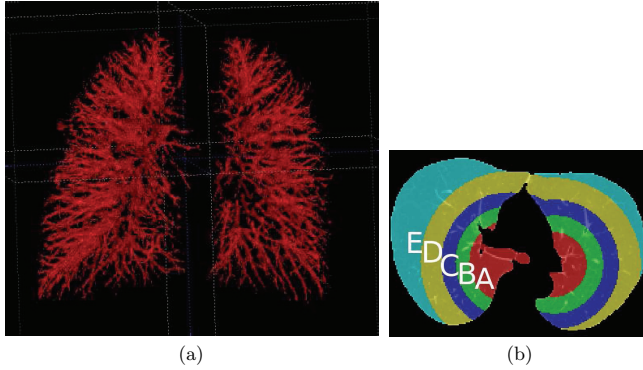
Fig. 4. (a) The extracted 3D vessel mask from the 256×256×200 volume, with no leakage apparent from visual inspection. (b) The five divided regions according to the distances from the two seed points.

region	A	B	C	D	E	all regions
case 1	127/129	300/307	491/503	794/821	557/585	96.8%
case 2	321/327	384/393	544/564	554/577	201/215	96.5%
total	98.3%	97.7%	97.0%	96.4%	94.8%	96.7%

Table 1. Validation on different regions. The left in -/- is the number of points correctly extracted; the right is the number of points available for validation in each region. The *true negative* is mainly from thin vessels around the lung periphery due to the restriction of the resolution. Note that *false positive* ratios are, however, unable to be reported, as in [1], due to the absence of the ground truth of the entire vessel segmentation.

vessels, while thick vessels exist in region A. For each region we counted the number of correctly classified samples. Table 1 lists the true positive rates averaged over different regions and different volumes. While our results are comparable to those reported in [1] for regions A, B and C, our accuracy decreases from the hilum to the periphery of the lungs (region D and E), where the vessels are almost one pixel thin and get blurred by both imaging noise and low resolution.

Due to the absence of the ground truth of orientation, we only provide qualitative results for the estimated orientation. Figure 5(a) shows an example of the initial orientation estimated from the multi-orientation filter bank in a small region. Figure 5(b) is the final regularized orientation, which is more smooth and sparse, i.e. the noise was eliminated.

(a)   (b)
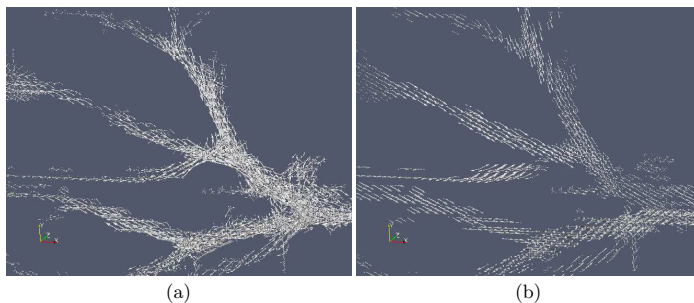
Fig. 5. A partial region from the entire mask. (a) The initial orientation from the filters. (b) The final orientation, which is much more regularized and has less noise than (a).

4   Conclusion

In this paper a new method is proposed for simultaneously segmenting the pulmonary vessel trees and estimating the vessel orientation from lung CT images. The orientation information is represented by histogram vectors, which unifies the representation for branching, intermediate and end points in a vessel tree. The orientation is also a variable to be optimized in the objective function, in contrast to being fixed in existing methods such as those reported [1][3]. The approach has two phases for each iteration: in the first phase, the estimated orientation is regularized by fixing the segmentation mask; in the second phase, the segmentation mask is updated by a min-cut/max-flow algorithm.

Future plans include applications of our method for nodule and abnormality detection. Also we want to examine quantitative validation of orientation and connectivity, and compare them with existing methods in a large scale test. Regarding computation efficiency, current unoptimized implementation took about 40 minutes on the $256{\times}256{\times}200$ volumes for one iteration on a Intel Xeon 3GHz CPU. This can be improved, since the method, especially the regularization step, is highly parallelizable. And about 60% of the time is spent sampling the Gaussian PDF for D^S (see [8]), which can be reduced by a better sampling function.

Acknowledgment

References

1. Shikata, H., Hoffman, E.A., Sonka, M.: Automated segmentation of pulmonary vascular tree from 3d ct images. Volume 5369., SPIE (2004) 107–116

2. Tozaki, T., Kawata, Y., Niki, N., Ohmatsu, H., Moriyama, N.: 3-d visualization of blood vessels and tumor using thin slice ct. In: IEEE Nuclear Science Symposium and Medical Imaging Conference. Volume 3. (1995) 1470–1474
3. Wu, C., Agam, G., Roy, A.S., Armato, III, S.G.: Regulated morphology approach to fuzzy shape analysis with application to blood vessel extraction in thoracic CT scans. In: Medical Imaging 2004: Image Processing. Volume 5370. (2004) 1262–1270
4. Lorigo, L.M., Faugeras, O.D., Grimson, W.E.L., Keriven, R., Kikinis, R., Nabavi, A., Westin, C.F.: Curves: Curve evolution for vessel segmentation. Medical Image Analysis **5** (2001) 195–206
5. Nain, D., Yezzi, A.J., Turk, G.: Vessel segmentation using a shape driven flow. In: Medical Image Computing and Computer-Assisted Intervention (MICCAI). Volume 1. (2004) 51–59
6. Boykov, Y., Funka-Lea, G.: Graph cuts and efficient n-d image segmentation. International Journal of Computer Vision **70**(2) (2006) 109–131
7. Sinop, A.K., Grady, L.: Accurate banded graph cut segmentation of thin structures laplacian pyramids. In: Medical Image Computing and Computer-Assisted Intervention (MICCAI). Volume 2. (2006) 896–903
8. Song, Z., Awate, S.P., Licht, D., Gee, J.C.: Clinical neonatal brain MRI segmentation using adaptive nonparametric data models and intensity-based Markov priors. In: Medical Image Computing and Computer-Assisted Intervention (MICCAI). Volume 1. (2007) 883–90
9. Nishino, M., Boiselle, P., Copeland, J., Raptopoulos, V., Hatabu, H.: Value of volumetric data acquisition in expiratory high-resolution computed tomography of the lung. J Comput Assist Tomogr. **28**(2) (2004) 209–214
10. Morrone, M.C., Owens, R.A.: Feature detection from local energy. Pattern Recognition Letters **6**(5) (1987) 303–313
11. Leung, T.K., Malik, J.: Contour continuity in region based image segmentation. In: ECCV '98: Proceedings of the 5th European Conference on Computer Vision-Volume I, London, UK, Springer-Verlag (1998) 544–559
12. Ramirez-Manzanares, A., Rivera, M., Kornprobst, P., Lauze, F.: A variational approach for multi-valued velocity field estimation in transparent sequences. In: Proceedings of International Conference on Scale Space and Variational Methods in Computer Vision. LNCS 4485, Ischia, Italy (2007) 227–238
13. Yushkevich, P.A., Piven, J., Hazlett, H.C., Smith, R.G., Ho, S., Gee, J.C., Gerig, G.: User-guided 3D active contour segmentation of anatomical structures: Significantly improved efficiency and reliability. Neuroimage **31**(3) (2006) 1116–1128

Automatic Non-rigid Lung Registration Method for the Visualization of Regional Air Trapping in Chest CT Scans

Jeongjin Lee[1, 2], Namkug Kim[1], Joon Beom Seo[1, *], Ho Lee[2], and Yeong Gil Shin[2]

[1]Department of Radiology and Research Institute of Radiology,
University of Ulsan College of Medicine, Asan Medical Center
registration@dreamwiz.com, namkugkim@gmail.com,
seojb@amc.seoul.kr,
[2]School of Electrical Engineering and Computer Science, Seoul National University,
San 56-1 Shinlim 9-dong Kwanak-gu, Seoul 151-742, Korea
{holee, yshin}@cglab.snu.ac.kr

Abstract. In this paper, we propose an automatic non-rigid lung registration method for the visualization of regional air trapping in full inspiration and expiration chest CT scans. To cope with intensity differences between CT scans, we segment the lung vessel and parenchyma in each scan. Then, we match them without referring any intensity information. We globally align two lung surfaces by affine transformation. Then, locally deformable transformation model is developed for the subsequent non-rigid registration. Subtracted quantification results are visualized by pre-defined color map. Experimental results show that proposed registration method is able to correctly align the full inspiration and expiration CT images in 10 patients. Our method can model the global and local deformation between full inspiration and expiration CT scans. Our non-rigid lung registration method may be useful for the assessment of regional air trapping by providing intuitive color-coded information of quantification results.

1 Introduction

Regional air trapping has been qualitatively assessed by comparing the intensity change of lung parenchyma during the respiration between full inspiration and expiration chest CT images. However, due to the change in body posture, the complex respiratory motion of lung, and heart beating, it is difficult for doctors to find correspondences manually between two CT data sets. Therefore, automatic lung registration methods, which align two images, and establish correspondences, are much helpful for radiologists to find areas of regional air trapping.

Several methods have been suggested for lung registration in temporal chest CT scans. Betke et al. [1] developed an automatic registration method of temporal chest CT images for the nodule registration. They detected the trachea, sternum and spine

---

* Corresponding author.

as anatomical landmarks using an attenuation-based template matching. The optimal rigid-body transformation that aligned the corresponding landmarks was found for the initial registration. Then, the initial surface alignment was refined step by step in an iterative closest-point process. However, the transformation model of this method assumed there was no local deformation between two scans so that this method cannot be applied to the registration between full inspiration and expiration scans. Fan et al. [2-3] proposed a volumetric lung registration method between CT images obtained at different stages of breathing. Both feature points and lung surfaces at consecutive frames were incorporated as a priori knowledge for 3D warping to derive an initial sparse comprehensive displacement field. This field was then interpolated over the entire volume in an iterative fashion governed by a model derived from continuum mechanics and 3D optical flow. However, their method was only validated for data sets, which had small volumetric differences. Dougherty et al. [4-5], and Torigian et al. [6] proposed an optical flow-based method to register images and visualize changes between temporal CT scans. Their method was applied to lung nodule assessment, evaluation of pulmonary enhancement, and functional changes due to air trapping. However, their methods could be applied to data sets, which had small volumetric differences.

In this paper, we propose an automated non-rigid registration method for regional air trapping visualization in full inspiration and expiration chest CT scans. To cope with intensity differences between two scans, we segment the lung vessel and parenchyma in each scan. We match them without referring any intensity information. We globally align two lung surfaces by affine transformation. Then, locally deformable model is used for the subsequent non-rigid registration. Our registration method is implemented in multi-resolution scale using a Gaussian pyramid. Our method can model the global and local deformation between full inspiration and expiration chest CT scans.

2   Methods

2.1   The Generation of a Binary Image

We segment lung parenchyma and vessel as follows. For the preprocessing, we remove the background area by applying 2D seeded region growing [7] on every slice with outermost seed points of each slice in the threshold range of -1024HU(Hounsfield Unit) to -200HU. The airway is extracted by applying 3D seeded region growing with the automatically detected seed point, which is the center of the uppermost airway in the threshold range of -1024HU to -950HU. Lung parenchyma is extracted by applying 3D seeded region growing [7] with the same seed point and subtracting the airway. The threshold range for the full inspiration CT is from -1024HU to -400HU. And the threshold range for the full expiration CT is from -1024HU to -200HU. From the hole filled lung parenchyma, we can extract lung vessel in the threshold range of -400HU and 3095HU. Then, the intensity of lung parenchyma is replaced as 0. And the intensity of lung vessel and other areas is replaced as 1. Subsequent registration procedure is performed on this binary image.

2.2 Initial Lung Surface Registration

In a surface registration algorithm, the calculation of the distance from a surface boundary to a certain point can be rapidly calculated using a preprocessed distance map [8]. To generate a 3D distance map, we approximate the global distance computation with chamfer distance transform [9], which is repeated propagation of local distances. Chamfer distance transform can be computed by performing a series of local operations while scanning image twice. We explain chamfer distance transform in 2D coordinate for an illustration. In forward scan, we compute $f_1(p)$ for all $p \in$ image in a single standard scan of image. For each p, f_1 has already been computed for all of the qs in B(p). If p has coordinates (x, y), B(p) contains (x, y+1), (x-1, y), (x-1,y+1) and (x+1, y+1).

$$f_1(p) = \begin{cases} 0 & \text{if } p \in \text{boundary} \\ \min\{f_1(q)+1 : q \in B(p)\} & \text{if } p \notin \text{boundary} \end{cases}. \qquad (1)$$

In backward scan, we compute $f_2(p)$ for all $p \in$ image in a single reverse standard (right-to-left, bottom-to-top) scan of image. A(p) contains the remaining neighbors of p, which are not contained in B(p).

$$f_2(p) = \min\{f_1(p), f_2(q)+1 : q \in A(p)\}. \qquad (2)$$

After the generation of chamfer distance map, the average distance between two surfaces can be calculated as follows.

$$\textbf{average distance} = \frac{1}{N} \sum_{x \in \text{Image}_1} \text{DistanceMap}_{\text{Image}_2}(Transform(x)), \qquad (3)$$

where the point in image 1 is transformed into the point, $Transform(x)$ in image 2. $\text{DistanceMap}_{\text{Image}_2}(x)$ is the distance value at the position, x of chamfer distance map. N is the number of points in the overlapped area between image 1 and 2. Transformation model is composed of translation, rotation, and scaling. The transformation is calculated as follows using three translation, rotation, and scaling parameters: T_x, T_y, T_z, R_x, R_y, R_z, S_x, S_y, S_z [10].

$$P_{\text{Image}_1} - C_{\text{Image}_1} =$$

$$R_x(\theta_x) \cdot R_y(\theta_y) \cdot R_z(\theta_z) \cdot S_x \cdot S_y \cdot S_z \cdot (P_{\text{Image}_2} - C_{\text{Image}_2})$$

$$+ T(T_x, T_y, T_z), \qquad (4)$$

where P_{Image_1} and C_{Image_1} are the position of voxel and center in image 1. P_{Image_2} and C_{Image_2} are the position of voxel and center in image 2. To search for transformation parameters, Powell's directional method [11] is applied as the optimization technique.

2.3 Non-rigid Registration based on Locally Deformable Model

Due to local deformation of lung, a point inside lung moves locally during the respiratory motion. Therefore, we can model the local movement of a point inside lung as the summation of linear transformation and translation. We calculate parameters of linear and translational transformation at each point, p around their neighborhood, Ω to minimize the following error functions.

$$\mathrm{E} = \sum_{\Omega \in P} [\mathrm{Image}_1(P) - \mathrm{Image}_2(L_p P + T_p)]^2 \ , \tag{5}$$

where L_p is 3×3 matrix including 9 linear transformation parameters at a point, p. And T_p is 3×1 matrix including 3 translational transformation parameters at a point, p. Ω can be approximated as 5×5×5 regular hexahedron. If we find the optimal transformation parameters according to Eq. (5), holes can occur when deformation parameters spatially change a lot. Therefore, we add the smoothness term [12] into locally deformable model to smooth the spatial change of deformation parameters. Eq. (5) with the smoothness term is optimized by Newton-Raphson iterative technique.

The range of Ω with respect to a point is limited in the neighborhood regular hexahedron centering the point. The degree of modeled deformation is limited by the size of a regular hexahedron. Therefore, multi-resolution approach is applied in this paper. Gaussian pyramid [13] for each image is generated. In Gaussian pyramid, the average intensity of 2×2×2 voxels in the current level is assigned as the intensity of one voxel in the next level. In this paper, five levels of hierarchical Gaussian pyramid are generated.

3   Results

Our method was performed on Intel Core2 Quad 2.4 GHz PC. Our method has been applied to ten pairs of clinical data sets, which were full inspiration and expiration chest CT images. Each image had a matrix size of 512 x 512 pixels with 0.75 mm slice thickness and spacing. For each scan, a stack of 450-550 contiguous slices were acquired.

The performance of our method was evaluated from the aspects of visual inspection, accuracy, and color mapping of subtraction images. First, 3D volume rendering images of each lung surface before and after the registration are shown in Fig. 1. Fig. 1(a), and (c) show a large amount of initial misalignments. This initial misalignment is completely corrected by proposed method. The lung boundaries of the full inspiration and expiration CT scans are aligned exactly, as in Fig. 1(b), (d).
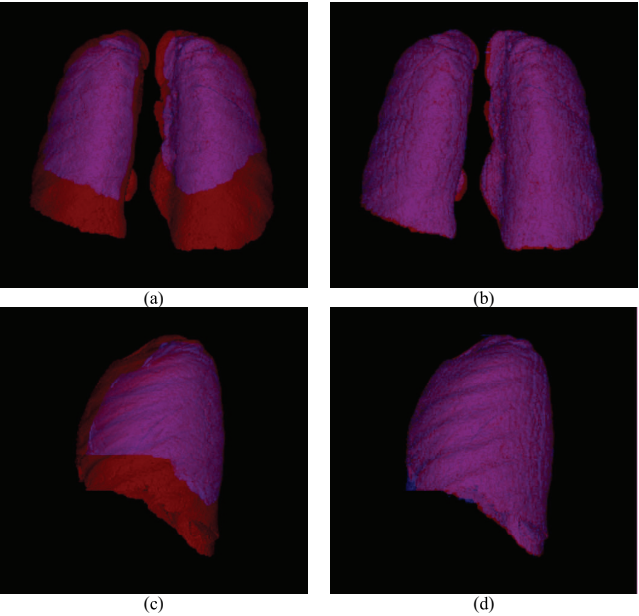
Fig. 1. 3D volume rendering images after the registration of full inspiration(red) and expiration(blue) CT images (a) before the registration(the front view) (b) after the registration(the front view) (c) before the registration(the side view) (d) after the registration(the side view)

We evaluate the registration accuracy of proposed method by measuring normalized mutual information(NMI) before and after the registration. Average NMI value measured inside the lung for ten patients was 1.00310 (before registration), 1.00394 (after initial registration), and 1.01246 (after non-rigid registration) as shown in Fig. 2. Average NMI value was the maximum after non-rigid registration.
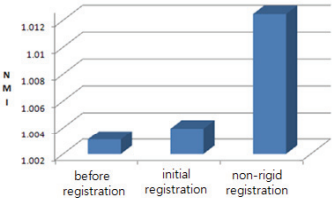
Fig. 2. Average NMI value

The joint histogram to measure NMI was shown in Fig. 3. The intensity value of CT images, which is in the range of -1024HU and 3095HU, was scaled in the range of 0 and 255. The joint histogram before the registration was dispersed as shown in Fig. 3(a). The joint histogram after the registration was congested as shown in Fig. 3(b), which meant the alignment of full inspiration and expiration CT images.

(a)   (b)
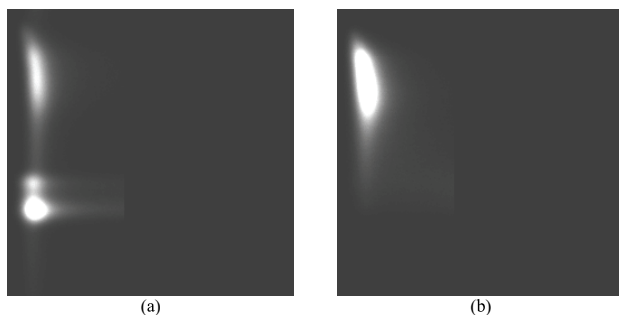
Fig. 3. Joint histogram (a) before the registration (b) after the registration

Finally, we evaluated the clinical efficiency of proposed method by color mapping of the subtraction image of lung parenchyma after the registration. Fig. 4(a) showed the subtraction image between the full inspiration and expiration CT images after the registration. Fig. 4(b) showed the color mapped subtraction image using a rainbow color table. Blue colored regions represented that the intensity difference between the full expiration and inspiration CT images was small. These regions meant air trapping regions. The average processing time per each patient was about fifty minutes on Intel PC with Core2 Quad 2.4 GHz CPU and 2GB RAM.

(a)   (b)
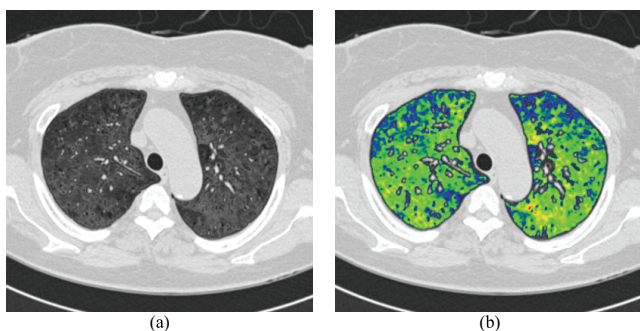
Fig. 4. Subtraction image after the registration (a) subtraction image (b) color mapped subtraction image

4 Conclusion

We have proposed an automatic non-rigid lung registration method for the visualization of regional air trapping in full inspiration and expiration chest CT scans. To cope with intensity differences between CT scans, we segmented the lung vessel and parenchyma in each scan. Then, we matched them without referring any intensity information. We globally aligned two lung surfaces by affine transformation. Then, locally deformable transformation model was developed for the subsequent non-rigid registration. Subtracted quantification results were visualized by pre-defined color map. Experimental results showed that proposed registration method was able to correctly align the full inspiration and expiration CT images in 10 patients.

References

1. Betke, M., Hong, H., Thomas, D., Ko, J.P., Landmark detection in the chest and registration of lung surfaces with an application to nodule registration, Medical Image Analysis, Vol. 7, No. 3 (2003) 265-281
2. Fan, L., Chen, C.W., 3D warping and registration from lung images, Proceedings of SPIE Medical Imaging 3660 (1999) 459 - 470
3. Fan, L., Chen, C.W., Reinhardt, J.M., Hoffman, E.A., Evaluation and application of 3D lung warping and registration model using HRCT images, Proceedings of SPIE Medical Imaging 4321 (2001) 234 - 243
4. Dougherty, L., Asmuth, J.C., Gefter, W.B., Alignment of CT lung volumes with an optical flow method, Academic Radiology, Vol. 10 (2003) 249-54
5. Dougherty, L., Torigian, D.A., Affusso, J.D., Asmuth, J.C., Gefter, W.B., Use of an optical flow method for the analysis of serial CT lung images, Academic Radiology, Vol 13, No. 1 (2006) 14-23
6. Torigian, D.A., Gefter, W.B., Affuso, J.D., Emami, K., Dougherty, L., Application of an optical flow method to inspiratory and expiratory lung MDCT to assess regional air trapping: a feasibility study, American Journal of Roentgenology, Vol. 188 (2007) 276-280
7. Gonzalez, R.C., Woods, R.E., Digital image processing, Prentice Hall, (2008).
8. Lee, J.J., Kim, N.K., Lee, H., Kang, S.H., Park, J.W., Chang, Y.I., Automatic skull segmentation and registration for tissue change measurement after mandibular setback surgery, Lecture Notes in Computer Science, Vol. 4319 (2006) 322-331.
9. Borgefors, G., Distance transformations in arbitrary dimensions, Computer Vision, Graphics, and Image Processing, Vol. 27 (1984) 321-345
10. Maes, F., Collignon, A., Vandermeulen, D., Marchal, G., Suetens, P., Multimodality image registration by maximization of mutual information, IEEE Transactions on Medical Imaging, Vol. 16, No. 2 (1997) 187-198
11. Maes, F., Vandermeulen, D., Suetens, P., Comparative evaluation of multiresolution optimization strategies for multimodality image registration by maximization of mutual information, Medical Image Analysis, Vol. 3, No. 4 (1999) 373-386
12. Horn, B.K.P., Robot vision, MIT Press, (1986).
13. Adelson, E.H., Anderson, C.H., Bergen, J.R., Burt, P.J., Ogden, J.M., Pyramid method in image processing, RCA Engineer, Vol. 29, No. 6 (1984) 33-41

Quantitative Assessment of Registration in Thoracic CT

K. Murphy, B. van Ginneken, J.P.W. Pluim, S. Klein, and M. Staring

University Medical Center, Utrecht, The Netherlands.

Abstract. A novel method for the quantitative evaluation of registration systems in thoracic CT is utilised to examine the effects of varying system parameters on registration error. Regional analysis is implemented to determine whether registration error is more prevalent in particular areas of the lungs. Experiments on twenty-four CT scan-pairs prove that in many cases significant reductions in processing time can be achieved without much loss of registration accuracy. More difficult cases require additional steps in order to achieve maximum precision. Larger errors appear more frequently in the lower regions of the lungs close to the diaphragm.

1 Introduction

The accurate registration of intra-patient thoracic CT scans has a variety of motivating clinical applications including improved ease of visual comparison, quantitative or automatic analysis of pathology progression, and in the case of inspiration/expiration pairs, analysis of lung function. In radiotherapy planning, registration information can be used to construct pulmonary motion models in order to propagate the location of the target region [4].

Although many promising registration algorithms exist, the quantitative evaluation of these techniques poses a further challenge due to the lack of an established reference standard. Without a means for quantitative assessment the improvement and optimisation of a registration algorithm is extremely difficult. Visual analysis of registered images is a time-consuming and subjective process and particularly in 3D images it is impossible to visually quantify subtle differences between results from various systems.

In this work a registration reference standard for thoracic CT pairs is formulated in an efficient semi-automatic manner, resulting in a well-distributed mesh of corresponding landmarks throughout the lung volumes to be registered. This reference standard is used to evaluate a parametric intensity-based registration algorithm under varying conditions. Regional error analysis is implemented to determine whether registration error is more prevalent in specific areas of the lungs.

2 Materials

All scans used in this work form part of an experimental lung cancer screening programme. Twenty-four patients (22 male, 2 female, ages 54-79yrs), each with a baseline and a follow-up scan (3-9 months apart) were chosen randomly from the database. All scans were obtained at full inspiration and without contrast injection on a 16 detector-row scanner (Mx8000 IDT or Brilliance 16P, Philips Medical Systems). They have a per-slice resolution of 512×512, with the number of slices per scan varying from 383 to 551. Slice thickness is 1mm with slice-spacing of 0.7mm. Pixel spacing in the X and Y directions varies from 0.55mm to 0.8mm.

All registration experiments were carried out on a standard desktop PC with an Intel Core 2 Duo processor, 2.4GHz.

3 Methods

3.1 Reference Standard Construction

In this section an overview of the reference standard construction method will be provided. The technique used is described in detail in [5].

The first step in the construction of the reference standard is the determination of landmark locations in the baseline scan. A fully automatic system has been designed which identifies 100 well-dispersed points throughout the lungs. These points are required to be sufficiently distinctive to enable them to be matched in the corresponding follow-up scan and points on the pleural surface itself are therefore excluded. A projection view of all the landmarks selected for a scan is shown in figure 1(a).

A semi-automatic system was developed to accurately match the voxels identified as landmarks in the baseline scan with voxels at the corresponding anatomic locations in the follow-up scan. Each scan pair was processed twice by independent observers (medical students). The observers were required to match at least 20 of the 100 landmarks manually using a custom-made graphical interface. The ordering of the points presented to the users was designed such that each subsequent point was well-distanced from its predecessors. During this phase the system utilised a thin-plate-spline (TPS) [1] and the thus-far annotated point pairs to model the relationship between the two images. The TPS model was evaluated at each new point by attempting to predict the correct correspondence and comparing this prediction with the subsequent user annotation. When 20 points were manually matched the system handled the remaining points automatically, provided that the TPS model had been validated by successful predictions of the user annotated matches. The annotation procedure took 20-30 minutes per scan-pair and did not require observers with significant experience of pulmonary anatomy.

(a)   (b)
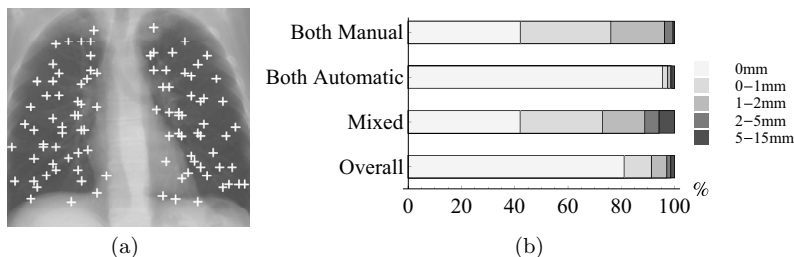
Fig. 1. (a)A set of automatically determined landmarks projected in the coronal direction. (b) Inter-observer differences categorised by match-types.

3.2 Registration Methods

Prior to registration the baseline and follow-up scans were down-sampled in order to reduce memory consumption. The down-sampling was by means of block-averaging such that the matrix size of $512{\times}512$ in the original images is reduced to $256{\times}256$, with the number of slices being reduced accordingly by a factor of 2. The calculated transform from the registration procedure was subsequently applied to the full resolution follow-up scan.

The registration procedure consisted of an initial affine registration step followed by an elastic registration to handle the non-rigid deformations of the lung tissue. Both registration steps involved a multi-resolution strategy using a Gaussian image pyramid. A mutual information cost function [7] was used in both cases along with a stochastic gradient descent optimizer [3]. The elastic registration deformations were modelled by a B-Spline grid [6]. The grid-size varied per resolution-level with the finest grid at the last level having a spacing of 8 voxels in each dimension.

In this work only the anatomy within the lungs is registered and all other structures are masked out. Previous experiments [5] have determined that this gives more accurate registration of the structures within the lungs. The mask used to distinguish the lungs from other anatomy was created by means of an automatic lung segmentation procedure based on the work of Hu et al. [2].

A number of experiments have been carried out in order to test the effects of tuning various parameters in the registration system. In particular the number of resolutions in the multi-resolution scheme and the number of iterations in the stochastic gradient descent optimizer are varied to determine their importance and optimal values. Registration error is analysed for each scan-pair at the various settings and regional error analysis is carried out over the entire dataset

4 Results

4.1 Reference-Standard Construction Results

Inter-observer differences The inter-observer differences were analysed to verify the ability of observers and of the system to find reproducible corresponding anatomic locations for the landmarks. In figure 1(b) the inter-observer differences in mm are illustrated, categorised by match-type. The match-type indicates whether the point was marked manually by one or both observers or whether it was chosen by the system. Regardless of match-type, 97% of all points had an inter-observer difference below 2mm. As is expected, points which were marked automatically by both observers are considerably more likely to have differences of 0mm than those which were marked manually, however manual observations are within 2mm of each other in 96% of cases.

Dispersal of Reference Points In order to verify that the reference points were evenly distributed around the lung volume, and later to analyse registration error in a regional manner, each lung was divided into 4 equally sized volumetric regions as follows: (See figure 2(a)). The centre of mass, c of both lung volumes together was identified. A sphere s, centred on c, was constructed such that 25% of the left lung volume was enclosed by s. This 25% represented the portion of left lung around the mediastinum. The remainder of the left lung was divided into 3 equally sized volumes by cutting at the appropriate slices. The right lung was then divided in an analogous fashion.

The number of points in each region over all images is illustrated in figure 2(b). It is clear that the points are well distributed over all areas. The right lung has slightly more points than the left which is to be expected due to its larger size, and the mediastinal area has a slightly higher concentration of landmarks since it is generally a much more distinctive region than any of the peripheral areas.

(a)   (b)
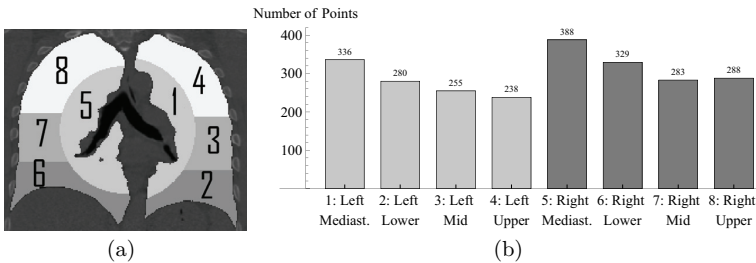
Fig. 2. (a)A slice showing a cross-section of the 3D lung partitions calculated. (b)The dispersal of points among the partitioned regions.

4.2   Registration Results

For each image-pair the computed transform T which maps from locations in
the deformed follow-up scan to locations in the original follow-up scan is applied
to each of the landmark points l from the baseline scan. It is clear that for
an accurate registration we expect $T(l) \approx lm$, where lm is the matching point
marked during reference standard formulation.

For all points lm_{obs1} marked (manually or automatically) by observer1 the
Euclidean distances $\delta(T(l), lm_{obs1})$ between $T(l)$ and lm_{obs1} were calculated us-
ing the appropriate T for the scan-pair. These distances δ were used as a measure
of registration accuracy.

Varying number of Iterations Reducing the number of iterations performed
in the stochastic gradient descent procedure is one way to considerably improve
the speed of the registration system. In order to determine the importance of this
parameter on the registration results the 24 scan-pairs were registered first with
512 iterations, then with 256 and finally with 128. All other parameters were
kept fixed, with 4 resolution levels during the affine registration and 5 during the
elastic. The registration errors $\delta(T(l), lm_{obs1})$ are shown in box-whisker plots for
each scan-pair at each setting in figure 3. In most cases the median error increases
slightly with fewer iterations although in a few instances, particularly those scans
with the lowest error measures the number of iterations has little effect.

The time to register a down-sampled image pair was reduced from approx-
imately 10 minutes with 512 iterations to 5 minutes with 256 iterations or 3
minutes with 128. Consideration must be given to balancing the registration
accuracy against the amount of time required to complete a registration since
many clinical applications demand results within a specified timeframe.

Figure 4 shows an example of a difficult case (the sixth case from figure 3).
Although the images are reasonably well aligned there are clearly some errors in
the vessel structure. Subtraction images shown in figures 5(a) (subtraction after
registration with 128 iterations) and 5(b) (subtraction after registration with
512 iterations) illustrate the difficulty of visually assessing registration results.

Varying number of Resolutions In this experiment the number of iterations
was kept fixed at 512 while the number of resolutions was varied. The registra-
tions were carried out firstly with 4 resolutions in the affine step and 5 in the
elastic step and secondly with 3 resolutions in the affine step and 4 in the elastic
step. The registration errors $\delta(T(l), lm_{obs1})$ for each scan-pair are shown in fig-
ure 6. In most cases the reduction in numbers of resolutions had little effect on
the registration error, however in a single case the registration result with fewer
resolution levels is so poor that the box showing the interquartile range of errors
cannot be seen at the scale shown in figure 6. The scans to be registered in this
case were so disparate that they required extra low-resolution steps in order to
overcome the large-scale differences early in the procedure. By adding the extra
resolution step back into the affine phase only, the median error is reduced from
approximately 18mm to just 0.5mm.
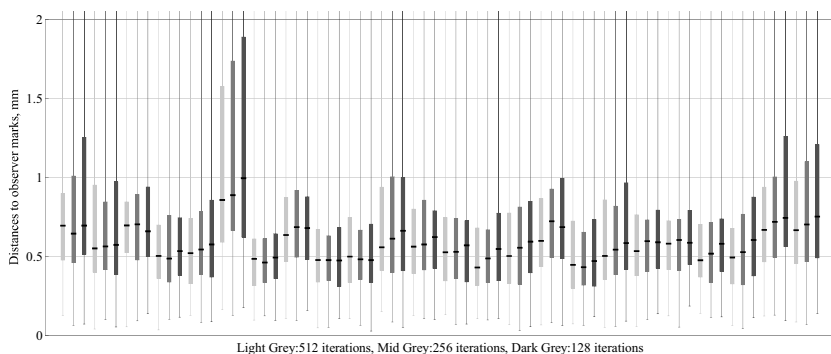
Fig. 3. Registration errors per scan-pair for varying numbers of iterations.
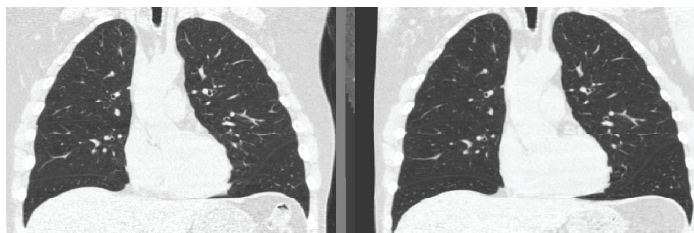
Fig. 4. A difficult case, corresponding slices from the fixed image and the deformed moving image.
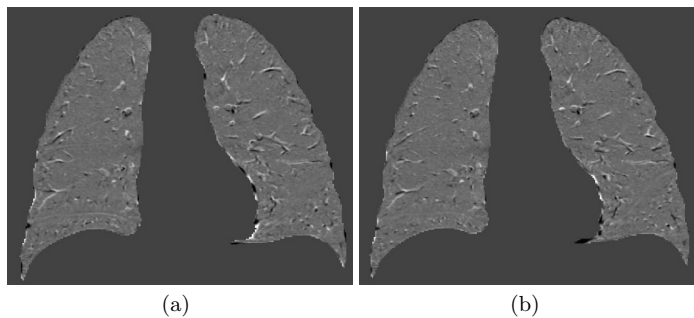
Fig. 5. The same difficult case as shown in figure 4. Subtraction image after registration with (a) 128 iterations and (b) 512 iterations.

The reduction of the number of resolution steps had a minimal effect on the time required to complete a registration, saving only in the order of 30 seconds of the 10 minutes.
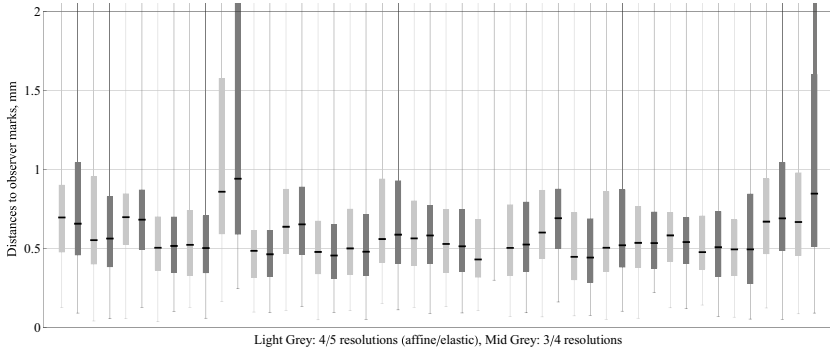
Fig. 6. Registration errors per scan-pair for varying numbers of resolutions

Registration Error per Region As described in section 4.1 and illustrated in figure 2(a) each lung was divided into 4 equal volumes to enable regional analysis. In figure 7 the error per region is shown for the system at 3 different settings. Based on the range of errors above the 0.75 quantile value it is clear that in all cases the largest errors for each lung are seen in the lower sections (region labels 2 and 6). This is to be expected since the motion of breathing affects the lower lungs much more significantly than the upper. Similarly, in all but one case, the upper section of the lung (region labels 4 and 8) has less error than any other section.

In figure 8 a closer view of the median values of the same box-plots is shown. The differences in median error values between regions are of the order of 0.1mm showing that for the majority of points there is little difference between regions. Median values for the regions close to the mediastinum and the diaphragm (region labels 1,2,5,6) tend to be slightly higher than those for the peripheral areas of the middle and upper lung.

5 Conclusion

A semi-automatic system for reference standard formulation has been used to generate a well-distributed mesh of corresponding landmark points in intra-patient thoracic CT scan pairs. The scan pairs have been non-rigidly registered using a parametric intensity based registration algorithm with various parameter
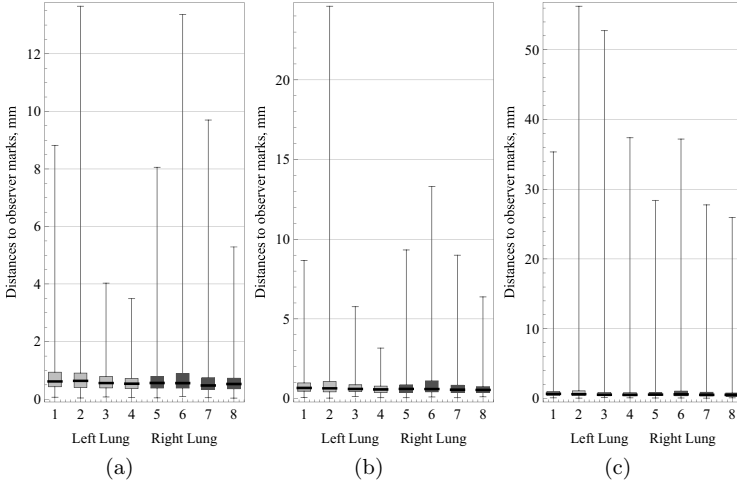
Fig. 7. Plots showing registration error per region for various system settings. The labels on the X-axes refer to the lung regions as shown in figure 2(a). (a) System with 512 iterations and 4/5 (affine/elastic) resolutions (b) System with 128 iterations and 4/5 resolutions (c) System with 512 iterations and 3/4 resolutions
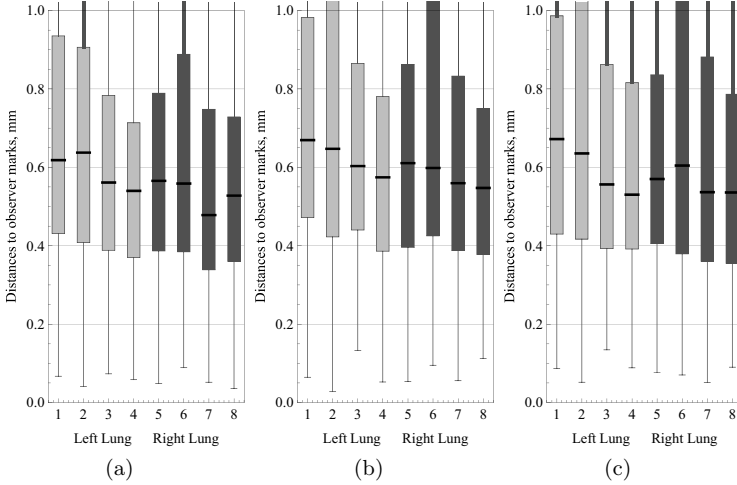
Fig. 8. A closer view of the median regions from the plots of figure 7.

settings. The constructed reference standard enabled the quantitative compari-

son of results from different system settings and the detection of subtle disparities in registration accuracies.

Regional analysis was also possible due to the regular distribution of the landmark points. It appears that larger errors are more likely to occur in the lower sections of the lung close to the diaphragm.

The results of parameter testing confirm that registration accuracy is generally negatively affected by the reduction of the number of iterations in gradient descent optimisation. However the accuracy difference in terms of millimetres is usually small or even negligible for easier cases (those where a good result is already achievable with fewer iterations, probably due to a good initial alignment of the baseline and follow-up images). Depending on the application and required accuracy the reduction in processing time may be more important than a negligibly small gain in accuracy.

The number of resolutions to be used in the multi-resolution scheme was shown to be an important factor in registration accuracy. In cases where the initial difference between scans is large the reduction of the number of resolutions proved to be detrimental to the system accuracy while providing little in terms of processing speed improvement.

These initial results suggest that optimisation of the registration algorithm may best be achieved by means of a feedback strategy whereby easier registration tasks may be completed with a minimal number of iterations while more difficult cases would be identified in the early stages and treated accordingly with additional steps.

References

1. F.L. Bookstein. Principal Warps: Thin-Plate Splines and the Decomposition of Deformations. *IEEE Trans. PAMI*, 11:567–585, 1989.
2. S. Hu, E. A. Hoffman, and J. M. Reinhardt. Automatic lung segmentation for accurate quantitation of volumetric X-ray CT images. *IEEE Trans. Med. Imaging*, 20(6):490–498, 2001.
3. S. Klein, M. Staring, and J.P.W. Pluim. Evaluation of optimization methods for nonrigid medical image registration using mutual information and B-splines. *IEEE Trans. Image Proc.*, 16(12):2879–2890, Dec 2007.
4. J. McClelland, J. Blackall, and S. Tarte. A continuous 4D motion model from multiple respiratory cycles for use in lung radiotherapy. *Medical Physics*, 33, 2006.
5. K. Murphy, B. van Ginneken, J.P.W. Pluim, S. Klein, and M. Staring. Semi-automatic reference standard construction for quantitative evaluation of lung CT registration. In *Medical Image Computing and Computer-Assisted Intervention*, 2008.
6. D. Rueckert, L.I. Sonoda, C. Hayes, D.L.G. Hill, M.O. Leach, and D.J. Hawkes. Nonrigid registration using free-form deformations: application to breast MR images. *IEEE Trans. Med. Imaging*, 18(8):712–721, 1999.
7. P. Thévenaz and M. Unser. Optimization of mutual information for multiresolution image registration. *IEEE Trans. Image Proc.*, 9(12):2083 – 2099, December 2000.

Registration of Hyperpolarized Helium-3 Diffusion MR Images for the Assessment of Changes with Albuterol Treatment in COPD Patients

PD Burstein[1], JP Mugler III[2], EE de Lange[2], J Mata[2], K Ruppert[2], TA Altes[2]

[1]School of Biomedical Engineering, Science and Health Systems, Drexel University, 3141 Chestnut St., Philadelphia, PA 19111

[2]Department of Radiology University of Virginia, Box 800170, Charlottesville, VA 22908

pablo.d.burstein@drexel.edu, taa2c@virginia.edu

Abstract

Our preliminary results show that albuterol administered to patients with chronic obstructive pulmonary disease (COPD) has no significant effect on the local ADC (apparent diffusion coefficient) values. To prevent confounds stemming from airway dilatation to affect albuterol's effect on COPD, we compute ADC and ADC difference maps over ventilated regions only. We compare intra-subject ADC values of pre and post albuterol administration conditions through the computation of RMSE, a voxel-wise based metric, in order to avoid compensation likely to occur in the computation of the (global) mean-ADC difference. Voxel-wise comparison is achieved through coregistration and normalization of the b0 images previously masked with corresponding combined ventilation masks (zero in unventilated voxels), both in the pre and the post albuterol reference spaces. The ADC and ADC difference maps are also computed for ventilated-only voxels. We also show that the registration process does not affect global ADC metrics.

Introduction

Hyperpolarized helium-3 is a gaseous contrast agent for MR imaging that, when inhaled, provides high temporal and spatial resolution images of the airspaces of the lung (1, 2). By measuring the spin density of the inhaled helium, images of lung ventilation can be obtained (3). Focal areas of poor ventilation have been demonstrated in a variety of lung diseases including asthma (4), chronic obstructive pulmonary disease (COPD) (2, 5), and cystic fibrosis (6, 7). In addition to ventilation imaging, the size and connectedness of the alveoli/distal airspaces can be assessed using diffusion techniques, similar to those used in brain MRI (8, 9). Elevated apparent diffusion coefficients (ADC) have been found in patients with COPD (10) and in animal models of emphysema (11, 12). These elevated ADC values are thought to correspond to the distal airspace enlargement that characterizes emphysema (12). To date, analysis of hyperpolarized helium-3 ADC maps has largely been confined to histogram based approaches, adequate for studies in which each subject is imaged a single time to assess for disease presence/severity, but that do not

take advantage of spatial or regional information contained within the ADC maps (8, 11, 13, 14). In order to assess changes on a regional basis, more sophisticated image analysis methods are required if the subjects are imaged more than once. One potential difficulty in developing a method to assess regional changes in the ADC maps is that the pattern of ventilation in the lung may vary with time or treatment (15). Since the ADC can only be measured in ventilated regions of the lung, comparison of ADC values at different time points is only meaningful at sites where the voxels are ventilated at the examined time-points.

The purpose of this study is to assess the effect of albuterol on COPD (emphysema) patients through the local analysis of the corresponding ADC maps. Our approach intends to circumvent the problem arising from the use of global metrics, such as difference of means, requiring no alignment of the pre and post-albuterol administration images, but suffering from possible compensation due to complementary changes, potentially leading to misinterpretation of the results (e.g.., false negative). In the following, we describe a newly developed paradigm for the computation of always-ventilated ADC difference maps. Our approach allows for the computation of locally based metrics such as RMSE. We present preliminary results when this algorithm is applied to a small group including healthy (normal) subjects and COPD patients.

Methods

Acquisition Protocol

Hyperpolarized helium diffusion MR imaging was performed in 6 subjects (2 healthy and 4 with COPD) with each subject being imaged before and after the administration of inhaled albuterol on two successive days (Day1, and Day2). Thus, 4 ADC maps were obtained for each subject. Albuterol is a bronchodilator and as such could change the regional pattern of ventilation but is not expected to affect alveolar morphology, i.e, ADC. All subjects underwent spirometry immediately prior to the helium MR scan on each of the two imaging days. The hyperpolarized helium gas was administered with approval from the Food and Drug Administration (FDA) as an Investigational New Drug (IND # 57,866), and this study was approved by our local institutional review board (IRB) with written informed consent was obtained from all subjects.

The helium-3 gas was polarized in a commercial helium polarizer (Model IGI9600 Helium Polarizer; Magnetic Imaging Technologies Inc., Durham, NC) by the collisional spin-exchange method and polarizations of 30 to 40% were typically achieved. Approximately 300 mL of polarized helium-3 gas was mixed with approximately 700 mL of medical grade nitrogen and dispensed into a Tedlar bag (Jensen Inert Products, Coral Springs, FL). The bag containing the helium dose was then taken to the MR scanner room where the subject was already positioned supine within the 1.5 T whole body MR scanner (Magnetom Sonata, Siemens Medical Solutions, Malvern, PA) equipped with the multinuclear imaging option. A flexible wrap coil (IGC Medical Advances, Milwaukee, WI) tuned to the helium-3 frequency was used for imaging. Starting from maximum expiration, subjects inhaled the hyperpolarized helium/nitrogen mixture from the bag through a short segment of plastic tubing. During the subsequent breath hold, contiguous axial diffusion-weighted images that covered the whole lung volume were acquired by using a

FLASH-based pulse sequence (TR/TE, 11/6.7 ms; FA, 7°; matrix, 80 x 128; FOV, 37 x 42 cm, slice thickness, 20 mm).   Diffusion sensitization was achieved by the addition of a bipolar gradient-pulse pair in the slice-select direction.   Images were obtained corresponding to two b-values, 0, and 1.6 s/cm^2, namely b_0 and $b_{1.6}$.   The images collected at b=0 s/cm^2 depict regional ventilation. To minimize the effects of signal attenuation from T1 decay and from radio-frequency pulses, the data corresponding to both b-values were acquired for a given phase-encoding view before the phase-encoding gradient strength was incremented to its next value. Assuming mono-exponential signal loss due to diffusion, two diffusion-weighted images, each corresponding to a different b value, are the minimum number required to calculate a spatial map of the ADC values.   Since the goal was to image subjects with COPD whose breath hold capacity may be limited, we elected to use the minimum number of b-values to minimize the breath hold duration.   The length of the breath hold depended on the number of images required to cover the entire lung volume but was typically less than 15 sec.

ADC Difference Map Computation

As opposed to anatomical images, such as those obtained from CT, hyperpolarized gas images of the lung tend to present substantial variation in the apparent anatomy at different acquisition times, depending on different temporal functional characteristics, such as ventilation and diffusion. In our case, this varying apparent anatomy may also arise as the result of different areas of the lung being ventilated before and after drug administration, and give place to local functional pattern differences, which may erroneously drive the registration. Since our focus is albuterol effect on COPD, i.e., the ADC maps within emphysematous tissue, we want to disregard ventilation changes at the airways level. To this end, our first step is to compute ventilated-only versions of the measured images (and, as a consequence, of the corresponding ADC maps), which will then be used during this study. This also ensures that the registration of the pre and post-albuterol images is driven by ventilated voxels only, thus preventing erroneous deformation into unventilated areas (likely to occur to match anatomy). For the sake of simplicity, we will use *pre* and *post* for short for *pre-albuterol* and *post-albuterol* administration, respectively.

For each subject, we first compute M_{Vpre} and M_{Vpost}, the pre and post ventilation binary masks, corresponding to b_{0pre} and b_{0post}, the pre and post b_0 (b = 0 s/cm^2) images, respectively. These masks, encoding ventilation pre and post-albuterol, are obtained by thresholding b_{0pre} and b_{0post} using Otsu's algorithm (17) (given our acquisition system, we assume SNR>2 for approximate Gaussian distribution). Then, in order to obtain combined ventilation masks both in pre (M_{CVpre}) and post (M_{CVpost}) reference spaces, we coregister b_{0post} into b_{0pre}, and apply the obtained transformation to M_{Vpost} to obtain M_{Vpost_inpre}. In turn, we apply the inverse transformation to M_{Vpre} to obtain M_{Vpre_inpost}. Since this study deals with intra-patient comparisons, where the acquisition was carefully performed to replicate same conditions, we assume the combined ventilation masks, obtained through coregistration only, to be accurate enough to closely match the actual ventilated regions. Finally, the combined vetilation mask in the pre space, $M_{CVpre,}$ and its counterpart in the post space, M_{CVpost}, are obtained by:

$$M_{CVpre} = M_{Vpre} * M_{Vpost\_inpre} \qquad (1)$$

$$M_{CVpost} = M_{Vpre\_inpost} * M_{Vpost} \qquad (2)$$

where * stands for pixel-wise multiplication. M_{CVpre} and M_{CVpost} have a value of 1 wherever the lungs are *always* ventilated, and 0 otherwise. Next, the *always* ventilated versions of b_{0pre} and b_{0post}, namely b_{0Vpre} and b_{0Vpost}, respectively, are obtained by masking:

$$b_{0Vpre} = b_{0pre} * M_{CVpre} \qquad (3)$$

$$b_{0Vpost} = b_{0post} * M_{CVpost} \qquad (4)$$

Similarly, we compute $b_{1.6Vpre}$ and $b_{1.6Vpost}$ (b = 1.6 s/cm^2). Coregistration is implemented using a two-step multi-resolution paradigm, where a quaternion transform, used to take care of translation, and rotation, is followed by an affine transform, which "fine-tunes" for any possible (generally minor) shear and scaling.

Following, b_{0Vpre} and b_{0Vpost} are coregistered, using the same method described above, and normalized through a (deformable) symmetric diffeomorphic (topography preserving) algorithm (16). This last step is introduced to correct any deformation caused by imaging setup/inflation. Since our objective is to align voxels corresponding to the same anatomical location but possibly presenting different ADC values (as a consequence of the albuterol administration), a metric suitable for multi-modality registration is necessary. In our case, we use mutual information (MI) as the similarity metric both for the non-deformable and deformable steps, already proven to be very efficient in multi-modality registration, and very suitable for small deformations (18, 19, 20, 21, 22). Finally, we warp b_{0Vpost} and its corresponding $b_{1.6Vpost}$ into pre space, to match b_{0Vpre} and $b_{1.6Vpre}$, respectively. Pre and post ADC maps are then computed using the general equation:

$$ADC(x) = \frac{1}{1.6} \log \frac{b_0(x)}{b_{1.6}(x)} \qquad (5)$$

where $b_0(x)$ is the signal intensity of the voxel in b_0 and $b_{1.6}(x)$ s the signal intensity of the corresponding voxel in $b_{1.6}$. We compute ADC using (1) for ventilated voxels only (Figure 1); unventilated voxels are set to zero in the ADC map. Notice the ventilation difference between pre and post-albuterol observed particularly in the upper right and upper left lobe in the b_0 images. Another advantage of using the combined ventilation masks is background noise removal, limiting the map to meaningful ADC values only, and, as a consequence, to more robust statistics.
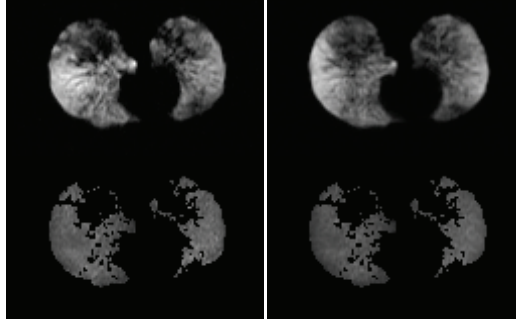
Figure 1. Pre (left) and post (right) albuterol b0 images (top) and corresponding ventilated ADC maps (bottom). ADC values are non-zero for ventilated voxels only. Ventilation differences between pre and post-albuterol is mainly observed in the upper right and upper left lobe in the b_0 images (top row).

Finally the ADC difference map, ΔADC, is computed as:

$$\Delta ADC = ADC_{post} - ADC_{pre} \qquad (6)$$

where ADC_{pre} and ADC_{post} are the values of ADC before and after albuterol administration. All our algorithms are implemented in C++ using The Insight Toolkit (ITK) libraries (18).

Validation
We compare global metrics commonly used in clinical assessment and tracking of COPD to assess whether warping affects the (warped) ADC map. These metrics, namely mean, rate of ADC above 0.35, 50% (median) and 90% percentile, are computed for the warped and non-warped (original) ADC post albuterol maps for each subject. As well, we establish basal acquisition variability, i.e., our control, by comparison of intra-subject images taken on Day1 and Day2 (Day1+1), without albuterol administration (pre).

Results
Warping Effect on ADC values
We are interested in assessing whether warping b_0 and $b_{1.6}$ images affects the corresponding ADC map computation in the new reference space. Table 1 shows the $\overline{ADC}$, % voxels with ADC $\geq$ 0.35, 50^{th} (median) and, 90^{th} percentiles for the post albuterol administration images, at Day1, before and after warping, along with the corresponding relative change (the last column). In general, it can be said that the global measures for the warped ADC map are not significantly different from the ones in the one computed in the acquisition (original) space. This is a desirable outcome if

one wants to compare ADC maps in a voxel-wise fashion. Relative high variation, though, is seen for the % voxels with ADC ≥ 0.35 measures, corresponding to subjects #1 and #3, with -53.19% and -66.62%, respectively. Despite these relative large changes, the absolute change is actually small, and can be explained as a consequence of interpolation limitations of the warped images; the negative change trend can be explained as interpolation towards background, given the general lack of signal in these images. Furthermore, the warped results still agree with the fact that these are normal subjects (low ADC signal).

#	$\overline{ADC}$ unwarped	warped	%Δ	#	50th Percentile (median) unwarped	warped	%Δ
1	0.2362	0.2347	-0.63	1	0.2331	0.2327	-0.17
2	0.3439	0.3367	-2.08	2	0.3361	0.3332	-0.88
3	0.2286	0.2286	0.03	3	0.2247	0.2267	0.91
4	0.4625	0.4595	-0.66	4	0.4554	0.4571	0.36
5	0.3374	0.3347	-0.79	5	0.2960	0.2942	-0.61
6	0.5272	0.5307	0.66	6	0.5296	0.5338	0.80

#	% voxels with ADC ≥ 0.35 unwarped	warped	%Δ	#	90th Percentile Unwarped	warped	%Δ
1	1.56	0.73	-53.19	1	0.2863	0.2707	-5.45
2	44.40	41.86	-5.72	2	0.4763	0.4456	-6.45
3	1.65	0.55	-66.62	3	0.2870	0.2688	-6.34
4	78.61	79.46	1.08	4	0.6470	0.6266	-3.16
5	34.10	34.13	0.08	5	0.5155	0.5064	-1.76
6	93.06	94.67	1.73	6	0.6700	0.6598	-1.51

Table 1. Unwarped, warped, and percentage change for $\overline{ADC}$, % voxels with ADC ≥ 0.35, and 50th (median) and, 90th percentiles. Day1, post-albuterol. The warping process shows little effect on most of the metrics. ADC change seen is less than 0.8%, 0.92% for the median, and 1.77% for the 90th percentile. Relatively elevated change is seen for % voxels with ADC ≥ 0.35 for the normal subjects (53.19% for subject #1 and 66.62% for subject #3). Still, the corresponding values for these subjects are very low, still being negligible after such a change.

Basal Variability
Table 2 shows the basal variability expected at two acquisition times, Day1 and Day2 = Day1+1, without albuterol administration. In general, no significant change is seen from Day1 to Day2. This result suggests that basal variability is negligible.

#	$\overline{ADC}_{Day1}$	$\overline{ADC}_{Day2}$	RMSE
1	0.2334	0.2286	0.0380
2	0.3610	0.3473	0.0713
3	0.2144	0.2141	0.0461
4	0.4858	0.4603	0.1072
5	0.3115	0.3489	0.0995
6	0.5312	0.5390	0.0746

Table 2. Basal variability. Measurements taken at Day1 and Day2. No albuterol administered (pre images). RMSE show no significant change, suggesting negligible basal variability, and, therefore, suitability for treatment assesment and tracking.

Albuterol Effect

Table 3 and Figure 2 show the albuterol effect on the ADC maps. Table 3 presents similar values to those obtained for the basal variability, Table 2. This suggests, as it was expected, that albuterol has no significant effect neither on the normal subjects (#1 and #3) nor on the COPD patients. This result tends to confirm the hypothesis that, in the short term, albuterol does not change the emphysematous tissue characteristics with respect to the size of the voids created by the degeneration process.

#	$\overline{ADC}_{Pre}$	$\overline{ADC}_{Post}$	RMSE
1	0.2334	0.2362	0.0431
2	0.3610	0.3439	0.0709
3	0.2144	0.2286	0.0462
4	0.4858	0.4625	0.1001
5	0.3115	0.3374	0.0625
6	0.5312	0.5272	0.0798

Table 3. Albuterol effect. Measurements taken at Day1, pre and post albuterol administration. In general, RMSE show little change, suggesting that albuterol has no effect on subjects.
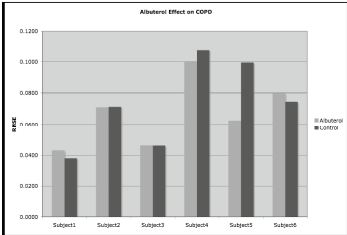
Figure 2. RMSE values for Albuterol administration effect are similar to those of the baseline (control) for all the subjects (subject #5 even shows higher degree of change for control than for albuterol administration). This suggests that there is no significant effect of albuterol.

Conclusions

We developed a method for voxel-wise comparison of hyperpolarized helium-3 ADC maps obtained in the same patient at different times. This method includes in the analysis only those regions of the lung that are ventilated at both time points (*always ventilated*) since the ADC cannot be measured in non-ventilated regions of the lung. This was achieved through the creation of combined ventilation masks in the pre and post-albuterol spaces. To test this algorithm, we imaged normal volunteers and patients with COPD before and after the administration of albuterol. As a bronchodilator, albuterol would be expected to change the pattern of ventilation but not the ADC values in ventilated regions of the lung. As expected, the patients with COPD present elevated ADC values as compared to the normal subjects, even after albuterol administration (while ventilation pattern changes were indeed observed in the patients). We found that the registration method, non-deformable followed by symmetric diffeomorphic algorithm with MI as the similarity metric, renders a good normalization of the pre and post-albuterol (as well as Day 1/Day2) images despite the ventilation changes while it does not significantly affect the global metrics. In addition, basal variability at acquisitions one day apart (Day 1/Day 2) was computed, providing the baseline for change detection. Finally, we found that there is little change in the ADC values before and after albuterol administration; i.e., RMSE of ADC difference maps support the hypothesis that no significant local change in ADC values, hence no significant change in the state of the emphysematous tissue, occurs as the result of albuterol administration.

References

1.   Middleton H, Black RD, Saam B, et al. MR imaging with hyperpolarized 3He gas. Magn Reson Med 1995; 33:271-275.
2.   Kauczor HU, Hofmann D, Kreitner KF, et al. Normal and abnormal pulmonary ventilation: visualization at hyperpolarized He-3 MR imaging. Radiology 1996; 201:564-568.
3.   Altes TA, Rehm PK, Harrell F, Salerno M, Daniel TM, De Lange EE. Ventilation imaging of the lung: comparison of hyperpolarized helium-3 MR imaging with Xe-133 scintigraphy. Acad Radiol 2004; 11:729-734.
4.   Altes TA, Powers PL, Knight-Scott J, et al. Hyperpolarized 3He MR lung ventilation imaging in asthmatics: preliminary findings. J Magn Reson Imaging 2001; 13:378-384.
5.   de Lange EE, Mugler JP, 3rd, Brookeman JR, et al. Lung air spaces: MR imaging evaluation with hyperpolarized 3He gas. Radiology 1999; 210:851-857.
6.   Donnelly LF, MacFall JR, McAdams HP, et al. Cystic fibrosis: combined hyperpolarized 3He-enhanced and conventional proton MR imaging in the lung--preliminary observations. Radiology 1999; 212:885-889.
7.   Mentore K, Froh DK, de Lange EE, Brookeman JR, Paget-Brown AO, Altes TA. Hyperpolarized HHe 3 MRI of the lung in cystic fibrosis: assessment at baseline and after bronchodilator and airway clearance treatment. Acad Radiol 2005; 12:1423-1429.

8.   Saam BT, Yablonskiy DA, Kodibagkar VD, et al. MR imaging of diffusion
        of (3)He gas in healthy and diseased lungs. Magn Reson Med 2000; 44:174-
        179.
9.   Mugler JP, 3rd, Brookeman JR, Knight-Scott J, Maier T, de Lange EE,
        Bogorad PL. Regional measurement of the 3He diffusion coefficient in the
        human lung. In:In: Proceedings of the 6th Annual Meeting of ISMRM.
        Sydney, Australia, 1998; 1906.
10.   Salerno M, Altes TA, Brookeman JR, de Lange EE, Mugler JP, 3rd. Rapid
        hyperpolarized 3He diffusion MRI of healthy and emphysematous human
        lungs using an optimized interleaved-spiral pulse sequence. J Magn Reson
        Imaging 2003; 17:581-588.
11.   Chen XJ, Hedlund LW, Moller HE, Chawla MS, Maronpot RR, Johnson
        GA. Detection of emphysema in rat lungs by using magnetic resonance
        measurements of 3He diffusion. Proc Natl Acad Sci U S A 2000; 97:11478-
        11481.
12.   Mata JF, Altes TA, Cai J, et al. Evaluation of emphysema severity and
        progression in a rabbit model: comparison of hyperpolarized 3He and 129Xe
        diffusion MRI with lung morphometry. J Appl Physiol 2007; 102:1273-
        1280.
13.   Fain SB, Panth SR, Evans MD, Wentland AL, Holmes JH, Korosec FR,
        O'Brien MJ, Fountaine H, Grist TM. Early emphysematous changes in
        asymptomatic smokers: detection with 3He MR imaging. Radiology. 2006
        Jun; 239(3):875-83.
14.   Salerno M, de Lange EE, Altes TA, Truwit JD, Brookeman JR, Mugler JP,
        3rd. Emphysema: hyperpolarized helium 3 diffusion MR imaging of the
        lungs compared with spirometric indexes--initial experience. Radiology
        2002; 222:252-260.
15.   Samee S, Altes T, Powers P, et al. Imaging the lungs in asthmatic patients by
        using hyperpolarized helium-3 magnetic resonance: assessment of response
        to methacholine and exercise challenge. J Allergy Clin Immunol 2003;
        111:1205-1211.
16.   Avants BB, Grossman M, Gee JC, Symmetric Diffeomorphic Image
        Registration: Evaluating   Automated   Labeling   of   Elderly   and
        Neurodegenerative Cortex and Frontal Lobe. WBIR 2006: 50-57
17.   Otsu, N. A thresholding selection method for grey-level histograms. IEEE
        Trans Systems, Man and Cybernetics SMC-9(1), 1979; 62-66.
18.   Insight Toolkit. www.itk.org
19.   Maes F, Collignon A, Vandermeulen D, Marchal G, and Suetens P.
        Multimodality Image Registration by Maximization of Mutual Information.
        IEEE Trans. Med. Imag., vol. 16, pp. 187–198, Apr. 1997.
20.   Brown LG. A Survey of Image Registration Techniques. ACM Comput.
        Surv. vol. 24, pp. 325–376, 1992.
21.   Viola P and Wells W. Alignment by Maximization of Mutual Information.
        Proc. 5th Int. Conf. Computer Vision, 1995, pp. 16–23
22.   Karaçali B. Information Theoretic Deformable Registration Using Local
        Image Information. Int. Jour. of Computer Vision, vol. 72 issue 3

An Automated Visualization System for Pulmonary Blood Flow Assessment using Non-Contrast Enhanced and Contrast Enhanced CT Images

Shoji Kido[1], Hidenori Shikata[2], Yoshitaka Tamura[3], Kazuo Awai[3], and Yasuyuki Yamashita[3]

[1] Applied Medical Engineering Science, Yamaguchi University
Graduate School of Medicine, Japan, kido@ai.csse.yamaguchi-u.ac.jp
[2] Ziosoft Inc., CA, USA,
[3] Diagnostic Radiology, Graduate school of Medical Sciences, Kumamoto
University, Japan

Abstract. Contrast enhanced (CE) CT images are commonly used for pulmonary embolism (PE) detection and blood flow assessment. The enhancement by the contrast medium introduces an intensity difference in the regions with and without sufficient blood flow. By the difference, thrombus can be visually recognized as a region with no enhancement in CE CT images. In parenchyma region, the enhancement can also be observed in the regions with normal blood flow. However, it is only subtle, causing difficulties in differentiating the regions with low blood flow. In order to assist blood flow evaluation in parenchyma region as well as in vessel, we have developed an automated visualization system that generates color-coded subtraction images between non-CE CT images and CE CT images via non-rigid registration. The subtraction images produced by the system resemble to the image obtained from blood flow scintigraphy SPECT scan, which shows clinical importance of this system by its potential ability to enable blood flow assessment with only CT images.

Keywords: Registration-2D/3D, Computer Assisted Diagnosis, Functional Imaging, Quantitative Image Analysis

1   Introduction

Blood transports substances to support cellular activities through blood vessels. Any blockages in a vessel lead to the lack of blood supply to the region covered by the subsequent vessel subtrees, which may cause severe symptoms. Such obstructions must be detected and treated in their early stage. Pulmonary embolism (PE) is a typical disease that prevents blood flow. The number of PE patients are increasing annually and the mortality rate is high [1]. For PE detection, contrast enhanced (CE) CT images are commonly used in clinical practice. Intensity values in CE CT images increase significantly compared with non-CE CT images in vessels with normal blood flow. The enhancement makes noticeable difference between a vessel and a thrombus,

which enables the thrombus to be visually recognized in the image. There are several systems that support PE detection such as in [2]. These systems are potentially able to detect subsegmental PEs those are difficult to identify by visual inspection. However, the systems are not usable for regional blood flow assessment in parenchyma.

The radii of the pulmonary vessels decrease along bifurcations to the periphery. In the capillary bed, blood flows from the pulmonary artery into vein. Though the peripheral vessels are well under CT resolution, the function can be noticed by the increase of CT value in CE CT images. Therefore, if there is any blockage in a vessel and blood flow decreases, low or no enhancement in the parenchyma region covered by the subsequent subtree is likely to be observed. It is clinically important to identify the region with low blood flow both for understanding severity and treatment planning. In parenchyma region, intensity increases by approximately 30-50 HU in normal case, whereas 100 HU or more enhancements can be observed in vessel. Due to the low contrast, it is relatively difficult to identify parenchyma region with low blood flow if only CE CT images were available.

In order to evaluate blood flow in parenchyma region, blood flow scintigraphy with SPECT image is useful. In blood flow scintigraphy, radioisotope (RI) is injected as maker and the SPECT scanner visualizes its distribution. Since pulmonary vessels spread in the entire lung, the distribution of the maker can be interpreted as the distribution of blood flow. Though the SPECT image cannot visualize anatomical structures by itself, CT-SPECT image fusion augments the missing information and the fusion image can improve diagnosis based on SPECT images [3, 4]. However, it is not well suited for evaluating local blood flow primarily due to the low spatial resolution.

In order to assist blood flow assessment in parenchyma region as well as in vessel, we have developed an automated visualization system that generates color-coded subtraction images between non-CE CT images and CE CT images via non-rigid registration. Advantages of this system are listed as follows;

a. The system can be used for the assessment of blood flow distribution in parenchyma region as well as PE detection.
b. The system uses only CT images that have high spatial resolution and ability to depict anatomical structures.
c. The system uses only CT images that are more commonly used than SPECT images. In addition, the management of the contrast medium is easier than RI, which may contribute to establish cost-effective blood flow evaluation method.

2   Method

The proposed system consists of three steps, preprocessing, registration and subtraction. In the preprocessing step, pulmonary vessels and lung region are extracted to form masks to be used in the following registration and subtraction steps. Registration is performed with a free-form deformation model based on B-spline. Subtraction images are generated using vessel segmentation result. The subtraction

images are color-coded and overlaid onto original CT images. Details can be found in the following subsections.

2.1 Preprocessing

Pulmonary vascular tree and lung segmentation are performed in this step. Vessels are automatically extracted from both non-CE and CE CT images based on the method in [5]. Segmentation process is as follows; Segmentation of the pulmonary vascular trees is performed by applying a combination of an adaptive Gaussian filtering and a vessel traversal. With cylindrical vessel model, the intensity values at the points inside the model decrease after applying Gaussian smoothing. Using the model, a voxel is marked as vessel point when the CT value at the voxel lowered after convolving Gaussian function with an appropriate standard deviation. A set of Hessian matrices and their eigenvalues / eigenvectors are calculated at the same time in order to determine if the voxel is close to the vessel centerline. A set of points that are close to the centerline is used as seed points for the following vessel traversal to construct connected vascular trees and to eliminate noise elements. Starting from a seed point, the traversal front position is updated by progressing toward a tangent direction approximated by an eigenvector of the Hessian matrix at the point. This process extracts pulmonary vascular trees without user interaction.

Lung region is extracted by a combination of global threshold and morphological operations. After the lung segmentation, two volumes of interest (VOIs) that tightly cover the left and the right lung are set to extract each lung. The volume of each VOI is approximately 1/4 of the original volume.

2.2 Registration

The registration for this specific application can be described as an intra-subject, mono-modality registration with non-rigid deformation and intensity alteration between the images. Lung surface moves smoothly on pleura during the respiration and that causes abrupt non-rigid deformation between two different CT scans in terms of the relative spatial distance between the structures in body, namely ribs, and the internal lung structures. The deformation often lowers registration accuracy in lung periphery when using a smooth deformation model. Therefore, the lung mask is applied to exclude structures outside the lung. The masked CT images that contain only lung region are used as input for the registration.

The height of an adult human lung with full inspiration is about 25-30 cm and a whole lung scan typically includes 300-350 slices with 1 mm slice thickness. Using whole volume for registration requires massive computational cost. In many applications, subsampled images are used to avoid the problem. Though registration with subsampled images provides visually good results, higher registration accuracy will improve the quality of the subtraction images. We therefore employed a multi-resolution registration approach [6] using automatically extracted two VOIs that include the left and the right lung separately. The first registration is performed using subsampled non-CE and CE CT images by factor of 2. Then the secondary

registration is applied independently to the left and the right lung with the original resolution. This reduces memory consumption per process and the secondary registration can be performed in parallel.

In the registration, a free-form deformation based on B-spline [7] is used as deformation model and mutual information [8] is used as similarity measure. In each resolution level, CT values are converted into 64 bins and the maximum of 50 iterations are allowed for the massive computational cost. We defined 64 bins by dividing equally the difference between maximum and minimum CT values of the lung region into 64 steps. The registration program was written using Insight Segmentation and Registration Toolkit (ITK).

2.3 Subtraction

While the registration in the previous step yields highly accurate registration results, the registered image will have some extent of misalignment. Simple subtraction of the intensities between the corresponding points includes errors introduced both by registration and interpolation. Smoothing is a way to possibly reduce the direct effect of the errors. However, taking mean value of all neighboring voxels without distinguishing vessel and parenchyma is likely to introduce additional error since the extent of the enhancement is different between these two regions. Therefore, the source values of the subtraction are calculated as mean of neighboring vessel points or non-vessel points depending on the location. If a voxel is on a vessel point, the mean intensity of all vessel points among the neighboring voxels is used for subtraction. Similarly, if a voxel is on a non-vessel point, the mean intensity of all non-vessel points among the neighboring voxels is calculated. When a point is on vessel (non-vessel) point and no vessel (non-vessel) points were available in the neighbor, the point is excluded from taking subtraction. The vessel segmentation result for both non-CE CT image and CE CT image are available to distinguish the non-vessel points and the vessel points. This process possibly relaxes registration errors and contributes to provide smooth subtraction images. We empirically used a cubic VOI of $7 \times 7 \times 7$ voxels to define neighboring voxels.

3 Results

The proposed system was applied to 10 sets of non-CE and CE CT scans from 10 patients and the color-coded images were generated for each case. All scans were acquired by a 40-slice multi-detector row CT scanner. Each image consisted of 512×512 pixels and the pixel dimension was about $0.6 \ mm^2$. Slice thickness was 1mm and approximately 300-350 slices were available per scan.

In the preprocessing step, it is visually confirmed that the vessel and lung segmentation were successfully performed for all cases ($10 \times 20 = 20$ scans), and the separation into the left and the right lung were also performed without failure.

Followed by the preprocessing step, the multi-resolution registration was performed. The registration takes approximately 15 minutes in each level with a PC

that equips a Pentium 4 3.0 GHz and 2 GB RAM, resulting in 45 minutes for total registration. After the registration, simple subtraction from the deformed CE CT image to non-CE CT image was taken for evaluating registration results. Please note that this simple subtraction was taken for evaluation purpose only and is different from the way to produce final subtraction images as explained in Section 2.3. Since the intensities are to increase in most of the locations in CE CT images, negative values after the subtraction, especially at vessel point, imply misalignment. Therefore, we counted the number of vessel points whose subtraction value is negative as a measure to evaluate the registration.

Table 1 shows the ratio between the number of negative points and the total number of vessel points. For comparison, two additional registration results are also shown in the table. Method 1 used subsampled, nopreprocessed CT images as input and up to 500 iterations were allowed. Method 2 used subsampled, lung masked CT images and up to 500 iterations were allowed. Method 3 is the multi-resolution registration used in our experiments. It took approximately 45 minutes per case for all three methods. The figures in the table show that the multi-resolution approach significantly reduced the number of negative points for all cases. The result indicates better registration with the multi-resolution method.

Table 1   Percentage of the negative voxels in the vessel mask

	Method 1 (%)	Method 2 (%)	Method 3 (%)
case 1	41.1	21.7	11.6
case 2	27.8	15.3	9.3
case 3	18.9	12.0	6.5
case 4	21.3	10.6	5.4
case 5	22.2	12.0	7.8
case 6	29.5	21.8	13.6
case 7	23.7	17.3	11.8
case 8	28.3	20.1	11.2
case 9	33.9	21.0	13.4
case 10	22.9	15.7	14.5

In order to visualize the voxels with negative subtraction value, volume rendering images obtained from case 2 are shown as an example in Figure 1. In the Figure 1(a) and (c), a low opacity value was assigned to the voxels whose subtraction value is in the range from 100 to 250, which infer well aligned vessels. In (b) and (d), voxels with negative value were overlaid on (a) and (c) with a brighter color. The figure shows that the negative value were observed most around the lobular fissures and the

diaphragm pointed by the white arrows, indicating that special care may be needed in these regions to obtain more registration accuracy.


(a) (b) (c) (d)
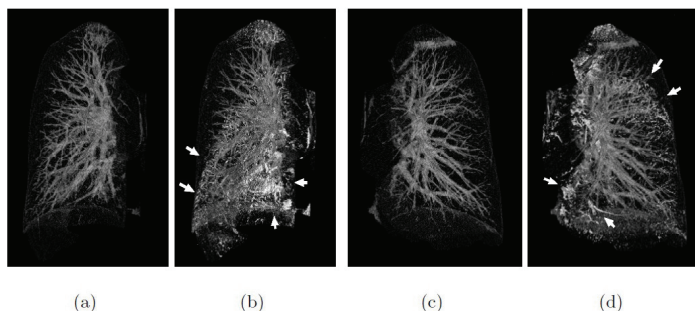
Fig. 1 Volume rendered images of the simple subtraction images after the registration. (a) and (c) visualize voxels whose subtraction values ranging from 100 to 250 in grey color, indicating well aligned vessels.(b) and (d) visualize negative voxels after subtraction in a brighter color, indicating potential misalignment.

Figure 2 shows two examples of the final output, color-coded subtraction images and corresponding RI-SPECT images. The RI-SPECT images available for this study had $64 \times 64 \times 64$ voxels in size and voxel dimension was $6.4 \, mm^3$. The RI-SPECT image was manually registered and overlaid onto the corresponding CT image with a colormap shown in each figure. Blue and red colors denote low and high blood flow, respectively. In Figure 2, (a) and (c) are RI-SPECT images and (b) and (d) show the subtraction images generated by the system. In Figure 2, (a) and (b) show a case which shows a defect region of pulmonary blood flow in right lung (case 1). (c) and (d) show a case which shows a defect region of pulmonary blood flow in right upper lung (case 2).

In case 1, a defect region of pulmonary blood flow in the right lung was indicated by a red arrow in RI-SEPCT image (a). The subtraction image shown in (b) also rendered the similar region with the same color. However, the defect region of pulmonary blood flow in subtraction image is smaller than that of RI-SPECT image. In RI-SPECT image, the defect region of pulmonary blood flow is visualized by obstruction of peripheral pulmonary arteries caused by Tc-99mMAA. However, subtraction image suggests that pulmonary blood flow exists in the defect region of RI-SPECT image. This indicates some pulmonary arteries are not obstructed completely. In case 2, (c) shows pulmonary blood flow is extremely low in most of the upper lobe of the right lung. However, the subtraction image in (d) suggests that pulmonary blood flow exists in same region. Also this image indicates defect region of pulmonary blood flow exists in left lung. This is not indicated by RI-SPECT image.

Figure 3 shows two examples that illustrate an airway and a vessel appeared in blue color. In (a), an occluded airway and a vessel run in parallel near a large bronchogenic carcinoma pointed by a yellow arrow. They can be clearly distinguished by the difference of the overlapped color. We can also evaluate activity of obstructive pneumonia caused by bronchogenic carcinoma. In (b), a vessel with nearly no blood flow was shown in blue. So, we can easily evaluate the existence or nonexistence of blood flow.

The Figure 2 and 3 show that the subtraction images produced by the system can support blood flow assessment in both parenchyma and vessels. The availability of the high spatial resolution is an advantage over RI-SPECT image for evaluating local blood flow.

(a)   (b)

(c)   (d)
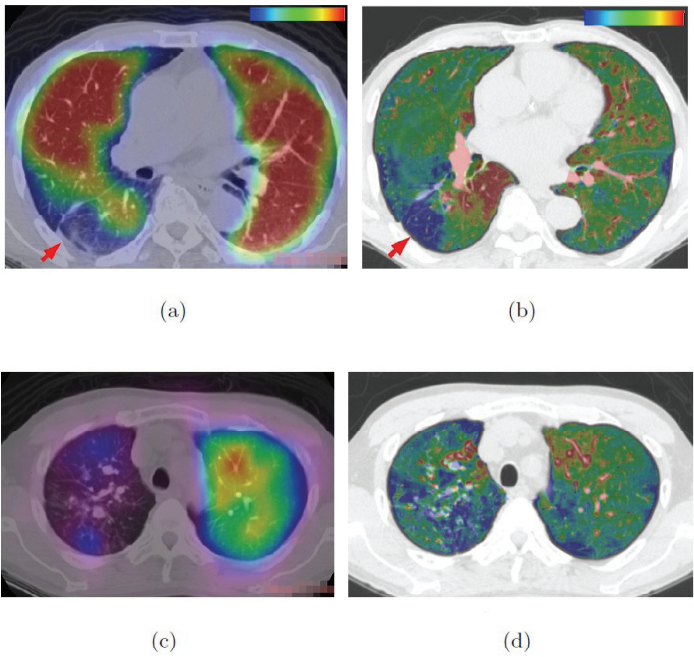
Fig. 2 Examples of the color-coded subtraction images that show low blood flow in parenchyma region. (a) and (c) are RI-SPECT images manually registered on the corresponding CT images. (b) and (d) are the subtraction images produced by our system.

(a)   (b)
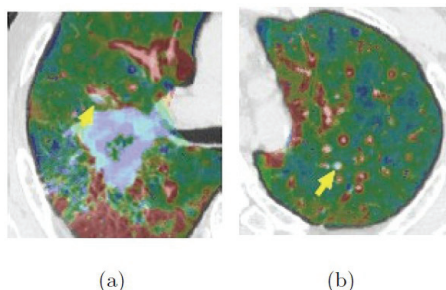
Fig. 3 Examples of the color-coded subtraction images that depict an airway and a vessel with no blood flow. (a) shows an occluded airway and (b) shows a vessel with nearly no enhancement.

3   Concluison

In this paper, we presented an automated visualization system for spatial assessment of pulmonary blood flow using non-CE and CE CT images. In the clinical practice, blood flow in vessel is evaluated with CE CT images and RI-SPECT images are used for blood flow assessment in parenchyma region. The proposed system can be used for both purposes. In many hospitals in Japan, we can not obtain RI-SPECT images easily, and costs of RI-SPECT study are high. However, our system uses only CT images, which may reduce the overall diagnosis time and cost. Moreover, our system may add more information about pulmonary blood flow to that of RI-SPECT study. The future task includes the development of an algorithm to detect the low blood flow region and PE from the images generated by our system.

Acknowledgements

References

1. Schoepf, U.J.: Diagnosing pulmonary embolism: time to rewrite the textbooks. The International Journal of Cardiovascular Imaging 21, 155-163 (2005)
2. Masutani, Y., MachMahon, H., Doi, K.: Computerized detection of pulmonary embolism in spiral CT angiography based on volumetric image analysis. IEEE Transactions on Medical Imaging 21, 1517-1523 (2002)
3. Gutman, F., Hangard, G., Gardin, I., Varmenot, N., Pattyn, J., Clement, J.F., Dubray, B., Vera, P.: Evaluation of a rigid registration method of lung perfusion SPECT and thoracic CT. American Journal of Roentgen logy 185, 1516-1524 (2005)
4. Suga, K., Kawakami, Y., Iwanaga, H., Seto, A., Matsunaga, N.: Comprehensive assessment of lung CT attenuation alteration at perfusion defects of acute pulmonary thromboembolism with breath-hold SPECT-CT fusion images 185, 1516-1523 (2005)
5. Shikata, H., Hoffman, E.A., Sonka, M.: Automated segmentation of pulmonary vascular trees from 3D CT images. In:Proc. SPIE, San Diego, USA (2004)
6. Schnabel, T., Rueckert, D., Quist, M., Blackall, J.M., Castellano-Smith, A.D., Hartkens, T., Penney, G.P., Hall, W.A., Liu, H., Truwit, C.I., Gerritsen, F.A., Hill, D.L.G., Hawkes, D.J.: A generic framework for nonrigid registration based on nonuniform multi-level free-form deformations. In: Lecture Notes in Computer Science (MICCAI 2001) 2208, 573-581 (2001).
7. Rueckert, D., Sonoda, L.I., Hayes, C., Hill, D.L.G., Leach, M.O., Hawkers, D.J.: Nonrigid registration using free-form deformations: Application to breast MR images. IEEE Transactions on Medical Imaging 18, 712-721 (1999)
8. Mattes, D., Haynor, D.R., Vesselle, H., Lewellen, T.K., Eubank, W.: PET-CT image registration in the chest using free-form deformations. IEEE Transactions on Medical Imaging 22, 120-128 (2003)

Fuzzy Pulmonary Vessel Segmentation Using Optimized Vessel Enhancement Filtering

Jens N. Kaftan[1,2], Atilla P. Kiraly[3], Marius Erdt[4],
Michael Sühling[2], and Til Aach[1]

[1] Institute of Imaging and Computer Vision, RWTH Aachen University,
52056 Aachen, Germany,
`jens.kaftan@lfb.rwth-aachen.de`
[2] Siemens Healthcare Sector, Computed Tomography, 91301 Forchheim, Germany
[3] Siemens Corporate Research, Princeton, NJ 08540, USA
[4] Fraunhofer Institute for Computer Graphics, Cognitive Computing &
Medical Imaging, 64283 Darmstadt, Germany

Abstract. Vessel segmentation within pulmonary images serves as a basis for a variety of applications, including PE detection and visualization, lung nodule detection, assistance in bronchoscopic navigation, lobe segmentation, and surgical planning. Although applications have different segmentation requirements, speed and accuracy is a clear benefit. A new approach combining a single parameter vessel enhancement filter and fuzzy connectedness is presented. The advantages of vessel filtering are brought to bear with a minimal impact on time by limiting the scales. Vesselness and intensity features are combined within a fuzzy segmentation framework, reducing the number of required scales and avoiding some of the drawbacks of each feature alone. Validation was performed on five datasets and Dice Similarity Coefficients (DSC) demonstrate an improvement of 9% (from 81% to 90%) on average for small vessels without influencing the accuracy for large vessels (95%) compared to an intensity-based method alone.

1 Introduction

Vascular tree segmentation in pulmonary computed tomography (CT) images is a core component of a variety of applications. Both the computer-aided diagnosis (CAD) and visualization of pulmonary emboli (PE) require vessel segmentation [1–3]. Although smaller vessels may have less direct clinical relevance, segmentation of these vessels can provide important information for tree hierarchy, lobar lung segmentation, and lung region assessment [4]. In PE CAD, the segmentation is used to reduce false positives while in visualization the segmentation is used as a basis for display. The same holds true for lung nodule CAD [5]. Navigation through the airways can make use of the vessel segmentation as well [6]. Finally, scoring and determining bronchio-arterial ratios within the airways also makes use of vessel segmentation [7].

Several challenges face vessel segmentation methods. Variations in attenuation due to partial volumes effects and the presence of PE can limit the segmentation extent. Nearby high density structures such as airway walls or connective

tissue can be confused with vessels. Finally, arteries and veins in close proximity can cause their segmentations to overlap [4, 8]. Although we do not address this last issue, we believe that an improved segmentation is also advantageous for artery-vein separation methods.

Intensity-based segmentation approaches using either threshold-based methods [4, 9, 10], front propagation techniques [8], or fuzzy techniques [11] have proven to be very efficient. These methods use an intensity model that is directly or indirectly utilized to detect and delineate vessels. However, due to variations in intensity as explained above, these methods tend to leak into surrounding non-vessel structures in some areas while missing smaller vessels.

Vessel enhancement filters, such as the Hessian-based filters [12, 13] have proven to be very capable in vessel segmentation. These methods compute eigenvalues of the Hessian matrix and combine them into a final value. These filters need to be applied at multiple scales to capture vessels of varying diameters. Approaches such as Zhou et al. [1] use up to 12 different scales. Speed quickly becomes an issue when multiple scales are used. Additionally, bifurcations usually do not have high responses to line-filters resulting in disconnected pieces.

We propose to use vessel enhancement filtering (Section 2) in combination with our intensity-based fuzzy approach [11]. Compared to other enhancement techniques our filter response function has been analytically designed to combine the eigenvalues of the Hessian matrix in such a way, that the output becomes maximal for tubular structures with Gaussian intensity profile. In particular it does not rely on any other parameter than the scale factor [14]. We demonstrate that this combination (Section 3) allows for accurate vessel segmentation with only a few filter scales, which would not be possible with vesselness filters alone. In addition, the intensity information is capable of steering the fuzzy-connectedness region growing process in areas such as bifurcations for which the filter output is lacking. The result is a synergy of the benefits of the individual methods compensating their individual drawbacks. The speed of intensity-based methods is preserved with the specificity benefits of vessel filtering (Section 5).

2   Optimized Vessel Enhancement

Vessel enhancement filters are typically based on the eigenvalues of the Hessian matrix H [12, 13] and have already been successfully applied to pulmonary vessel segmentation [1, 15]. In practice, the image signal is convolved with the six second order derivatives of the Gaussian g and the eigenvalues of the resulting Hessian matrix are combined into a filter output that enhances tubular-like structures.

Assuming that a vessel $v(\mathbf{r})(\mathbf{r} = (x, y, z))$ can be modeled locally by a cylinder with a radial Gaussian intensity profile with, e.g., a, b being the intensity values at the vessel center and boundary, respectively, which is oriented along the x-axis for the following considerations,

$$v(\mathbf{r}) = b + (a - b) \exp{-\frac{y^2 + z^2}{2\sigma_v^2}} \tag{1}$$

we compose the filter $h(\mathbf{r})$ as a linear combination of second order derivatives of the Gaussian with the standard deviation σ_h chosen to be the same as the standard deviation of the vessel model σ_v such that its convolution with this ideal vessel signal $v(\mathbf{r})$ is maximized [14]

$$S = v * h = \int_{\mathbb{R}^3} v(\mathbf{r})h(-\mathbf{r})d\mathbf{r} \qquad (2)$$

This can be solved analytically using the mathematical framework of Lagrange multipliers yielding

$$h(\mathbf{r}) = c \cdot \left(\frac{2}{3}g_{xx}(\mathbf{r}) - g_{yy}(\mathbf{r}) - g_{zz}(\mathbf{r}) \right) \qquad (3)$$

with $c = \sqrt{\frac{3}{5\pi^{3/2}}} \cdot \sqrt{\sigma_v}$.

In reality, the orientation of the vessel is not known, but the filter has to be oriented along the vessel to obtain maximum response. This direction is equal to the eigenvector corresponding to the eigenvalue with the smallest magnitude of the Hessian matrix. Let the eigenvalues of H be $\lambda_{1,2,3}$ with $|\lambda_1| \leq |\lambda_2| \leq |\lambda_3|$, then the optimal filter output can be computed as:

$$S_{\mathrm{opt}} = v * h_{\mathrm{rot}} = \frac{2}{3}\lambda_1 - \lambda_2 - \lambda_3 \qquad (4)$$

Since vessel structures are brighter than the background we expect $\lambda_{2,3} < 0$. Unfortunately, very bright, plate-like structures for which the cross-sectional profile is Gaussian only in one direction, i.e., the intensity decreases rapidly in one direction but not in the orthogonal one, the filter output might still become large because of either $|\lambda_2|$ or $|\lambda_3|$ being large. Hence, we multiply the filter output (4) by an isotropy factor

$$\kappa = \begin{cases} 0 & \text{if } \lambda_2 > 0 \text{ or } \lambda_3 > 0 \\ 1 - \frac{|\lambda_2| - |\lambda_3|}{|\lambda_2| + |\lambda_3|} & \text{else} \end{cases} \qquad (5)$$

As the pulmonary vascular structure consists of vessels with varying diameters, multiscale results are therefore obtained combining the filter output at different scales σ_{h_i}. It can be shown that the filter output with an idealized vessel is proportional to $\sigma_h^{3/2}$, which needs to be accounted for

$$S_{\mathrm{opt}} = \max \left(\kappa \cdot \sigma_{h_i}^{-3/2} \cdot S_{\mathrm{opt}}(\sigma_{h_i}) \right) \ \forall \ i \qquad (6)$$

Results of the optimized vesselness filter, shown in Figure 1, allow for a clear distinction between vessel- and non vessel-like structures. Fig. 1a shows the original HU values while in Fig. 1b the lung regions have been replaced by the filter output. Additionally we have recently shown its applicability to liver vessel enhancement in CT data [14]. Note that compared to other Hessian-based techniques such as [12], our filter function requires only a single parameter (scale).

(a) Original (b) Filter output
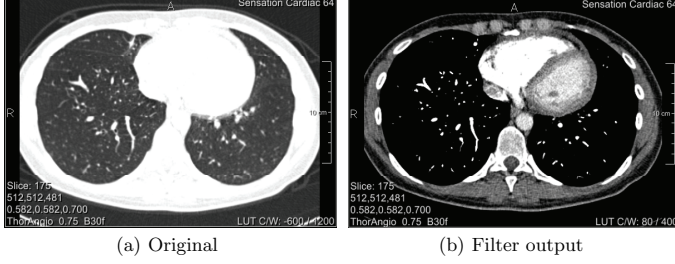
Fig. 1. Contrasted CT dataset of the lung. (a) Original (lung window) and (b) visualization for which the lung areas have been replaced by the filter output using $\sigma_h = [1, 2, 4]$ mm while the body regions are shown as original value (vessel window).

3 Combined Fuzzy Segmentation

For the sake of completeness we will first briefly review our general fuzzy segmentation framework [11] before detailing the incorporation of the vesselness filter output into the intensity-based system and discussing the different parameters.

3.1 Fuzzy Vessel Segmentation

Assuming a segmentation of the left and right lung [16] is already given, we first detect multiple seed points throughout the whole lung. To this end we segment the major vessels with a high specificity using an intensity-based method [4]. Hence we apply a lower threshold $T = 150$ HU and eliminate components smaller than $V_{min} = 500$ voxels in size. Next each resulting component is converted into one or more seed-points by locating and clustering vessel points with locally maximal distance to the vessel surface utilizing a 3D distance transformation. The cluster representatives $\mathbf{s}_i$ are selected with preference to large distance values, i.e., large radius estimates. Assigning to the so detected seed-points a probability measure of $P_{\text{vessel}}(\mathbf{s}_i) = 1$, we calculate the probability measure for the remaining voxels using the fuzzy connectedness algorithm [17].

The likelihood that two neighboring voxels $\mathbf{c}, \mathbf{d}$ belong to the same class, here to the vascular tree, is described by the local affinity $\mu_\kappa(\mathbf{c}, \mathbf{d})$. Using, e.g., an intensity-based probability function, the affinity can be defined as

$$\mu_\kappa(\mathbf{c}, \mathbf{d}) = \begin{cases} e^{-\frac{1}{2\sigma^2}\left(\frac{f(\mathbf{c})+f(\mathbf{d})}{2}-\mu\right)^2} & \text{if } \frac{f(\mathbf{c})+f(\mathbf{d})}{2} < \mu \\ 1 & \text{else} \end{cases} \qquad (7)$$

with $f(\mathbf{c})$ being the intensity value at position $\mathbf{c}$ and μ, σ^2 being the expected intensity value and variance of the used Gaussian function. The "strength of connectedness" μ_N of two distant voxels $\mathbf{c}, \mathbf{d}$ along a certain path $p_{\mathbf{c}, \mathbf{d}}$ is simply

the smallest pairwise fuzzy affinity along this path. A path $p_{\mathbf{c},\mathbf{d}}$ from $\mathbf{c}$ to $\mathbf{d}$ is a sequence of $m > 2$ neighboring voxels $\langle \mathbf{c}^{(1)} = \mathbf{c}, \mathbf{c}^{(2)}, ..., \mathbf{c}^{(m)} = \mathbf{d} \rangle$, i.e.,

$$\mu_N(p_{\mathbf{c},\mathbf{d}}) = \min \left[\mu_\kappa(\mathbf{c}^{(1)}, \mathbf{c}^{(2)}), ..., \mu_\kappa(\mathbf{c}^{(m-1)}, \mathbf{c}^{(m)}) \right] \tag{8}$$

The global connectivity $\mu_K(\mathbf{c}, \mathbf{d})$ is then the largest of the strength of connectedness of all possible paths $\mathbb{P}_{\mathbf{c},\mathbf{d}}$ between $\mathbf{c}, \mathbf{d}$:

$$\mu_K(\mathbf{c}, \mathbf{d}) = \max_{p_j \in \mathbb{P}_{\mathbf{c},\mathbf{d}}} \left[\mu_N(p_j) \right] \ \forall j \tag{9}$$

The probability measure that a voxel $\mathbf{x}$ belongs to a vessel is hence:

$$P_{\text{vessel}}(\mathbf{x}) = \max_{\mathbf{s}_i} \left[\mu_K(\mathbf{x}, \mathbf{s}_i) \right] \quad \text{with} \quad P_{\text{vessel}}(\mathbf{s}_i) = 1 \ \forall i \tag{10}$$

Note that even if P_{vessel} drops below 0.5 for a voxel, this voxel can still most likely belong to the vascular tree. In fact, an appropriate threshold has to be applied for binarization.

3.2 Feature Combination

Using an intensity-based function alone, one can observe that small vessels that are significantly darker than large vessels are often missed while already leaking into non-vessel structures. Hence it is reasonable to focus on small vessels when applying the vessel enhancement filter while vessels of larger scale are segmented well using an intensity-based function alone. Additionally to minimize the computational complexity we prefer to use as few filter scales as possible. A combination of $\sigma_h = [1, 2, 4]$ mm have been experimentally determined to be a good choice (see also Fig. 1). Having now the original and filter output values, the local affinity becomes a multimodal function with $\mu_{\kappa_{\{1,2\}}}(\mathbf{c}, \mathbf{d})$ being dependent on the parameters $\mu_{\{1,2\}}, \sigma^2_{\{1,2\}}$ and $f_{\{1,2\}}(\mathbf{c})$ being the intensity value and filter output at $\mathbf{c}$, respectively. The scalar output of μ_κ has been chosen to be the maximum value of $\mu_{\kappa_{\{1,2\}}}$. The parameters $\mu_{\{1,2\}}$ are the average intensity and filter output values of the seed points $\mathbf{s}_i$ while $\sigma^2_{\{1,2\}}$ have to be chosen (Sec. 5).

Consequently, the fuzzy connectedness region growing targets large vessels via the intensity-based function and smaller vessels via the filter output. However, in areas where the filter output is locally low, such as it is the case for some bifurcations, the intensity-based function will still prevent the growing process to terminate early. That is, locally low filter responses will not cause additional vessels to be missed while also having no impact on the false positive rate.

4 Validation

The validation of pulmonary vessel segmentation systems is especially difficult because of the complexity and size of the vascular tree structure. Often such methods are only qualitatively validated because of a missing ground truth

for clinical data. Additionally some authors estimate the accuracy of their approaches by counting the number of manually placed centerline points that are included into the segmentation output [1], which however allows no quantification of false positive regions. Wu et al. [9] compare their segmentation results of clinical data with additive artificial noise to the result of the original data. Such evaluation feature indicates the robustness of the segmentation method against noise but does not quantify the segmentation accuracy itself.

We use a fast, semi-automatic method to create reference segmentation in sub-volumes based on the random walker algorithm [18, 19]. The user first defines a 3D region of interest (ROI) and threshold-based fore- and background seed points are added within this subvolume. These seeds are then used as initialization for the random walker algorithm, a graph-based segmentation approach that can be used for interactive segmentation purposes. To this end, each voxel is represented by a node and neighboring voxels are connected by weighted edges using a weighting function, such as

$$w_{ij} = \exp\left(-\beta \cdot \Delta d_{ij} \cdot |I(\mathbf{x}_i) - I(\mathbf{x}_j)|\right) \tag{11}$$

Here $I(\mathbf{x}_i)$ is the intensity value of voxel i, Δd_{ij} the distance between voxel i and j and β a free parameter (here: $\beta = 150$). Next, the random walker algorithm is consecutively applied along with user interaction until the segmentation output is sufficiently accurate. Although strictly spoken one would have to validate the semi-automatically created segmentation results itself, which is as already discussed difficult due to a missing ground truth, the characteristics of the described method indicates that it is very well suited for validation purposes:

- The random walker algorithm allows the creation of any arbitrary segmentation given enough seed points, i.e., the user is not bounded by the system [18].
- Only a comparibly small number of seed points is required to produce a segmentation that is almost identical to a manual segmentation [19].
- Variations of the exact position of the seed points results only in small differences in the final result, i.e., the method allows the creation of results with a higher reproducability than manual methods [19].

Using the so created reference segmentations for quantitative validation and parameter optimization we use sensitivity ($\frac{TP}{TP+FN}$), specificity ($\frac{TN}{TN+FP}$), and the Dice similarity coefficient (DSC)

$$DSC = \frac{2 \cdot TP}{2 \cdot TP + FP + FN} \tag{12}$$

with TP, FP, and FN being the true positive, false positive, and false negative voxel count as validation features. Note that each threshold for P_{vessel} will result in one set of features.

Although the semi-automatically created reference segmentations have a high accuracy, they are still limited by the time the radiologist (or any other medical expert) can spend on their creation. We avoid superficial differences between the automatic segmentation and the ground truth from influencing the results

(a) DSC for large vessels   (b) DSC for small vessels
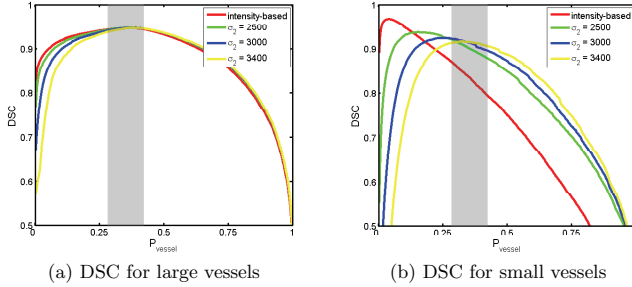
Fig. 2. DSC for all datasets with varying parameters. The red curves show the performance using the intensity-based method only. For larger vessels (a) the variation of parameters does not have any significant influence on the accuracy within the optimal band of thresholds highlighted by the gray bar. In (b) one can see that the inclusion of the vesselness criteria shifts the DSC curve for small vessels towards higher thresholds.

by not counting voxels within a margin of one voxel along the outer vessel surface as neither positive nor negative. This methodology also reduces a potential validation bias since the interior segmentation should be more independent on the semi-automatic method than the contour.

5 Results

The proposed method has been evaluated on five different randomly selected contrast enhanced chest CT scans from clinical routine of patients referred for PE. These data have been acquired using Siemens Sensation 16/64 scanners with voxel sizes ranging from 0.55-0.7 mm in x, y and 0.6-0.7 mm in z-direction. For quantitative validation 30 manually selected ROIs of size 50^3 voxel have been semi-automatically segmented. For each patient, two ROIs have been randomly placed in regions of large vessels and four in regions of small vessels within the periphery of the lung. Although the intensity-based method alone provides in general good results, the evaluation of large and small vessels individually (Fig. 2, red curves) reveals that for thresholds of the probability measure P_{vessel} that are very well suited for large vessels, the accuracy for small vessels is not optimal.

Adding the output of the vesselness filter to the fuzzy segmentation method does especially increase the segmentation accuracy of smaller vessels for high and medium thresholds without significantly affecting, as expected, the segmentation output for large vessels, since we apply the filter for smaller scales only. The comparison for different parameter settings of σ_1 and σ_2 are shown in Figure 2. In fact, only the ratio between both parameters is of importance when combining both affinity features, hence we keep $\sigma_1 = 200$ HU constant while varying $\sigma_2 \in [2500, 3400]$. One can observe that an increased value of σ_2 will shift the DSC

(a) DSC incl. large and small vessels   (b) ROC incl. large and small vessels
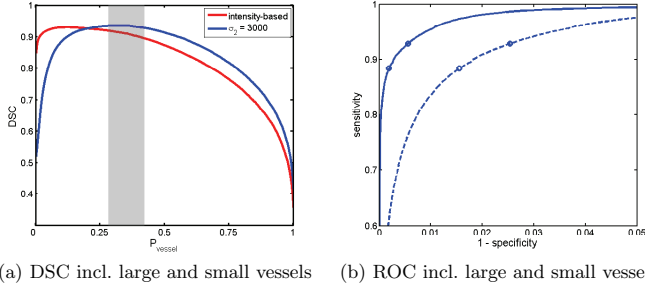
Fig. 3. Dice similarity coefficient and ROC curve for all datasets including large and small vessels with varying parameters. (a) The red curve shows the DSC using the intensity-based method only and the blue curve using a combination with $\sigma_2 = 3000$. Note that the blue curve has its optimum in the middle of the preferred threshold range. (b) ROC curves with (solid line) and without the one voxel security margin (dashed line) used for evaluation. The circles represent the operating points of the left and right boundary of the band of optimal thresholds, respectively.

curve for small vessels towards higher threshold. Unfortunately, the maximum accuracy also decreases slightly with increasing σ_2. This is most likely caused by introducing also some false positive regions by the vesselness feature, where the filter output responses relatively high to non-vessel structures. Note, however, that the optimal threshold for small vessels using the intensity feature alone has practically no relevance since it would cause severe leakage in other regions. Taking this into account a parameter of $\sigma_2 = 3000$ (blue curve) turns out to be a good compromise as its DSC just turns to its maximal value at the left boundary of the specified optimal range of thresholds. Using, e.g., a threshold for P_{vessel} of 0.4 that segments the vessels relevant for PE applications consistently well, the DSC increases from 81% to 90% for small vessels while being 95% for large vessels. Comparing the intensity-based function only with the combination using $\sigma_2 = 3000$, the overall DSC curves are shown in Figure 3a. The overall sensitivity and specificity within the band of optimal thresholds varies between $88.4-92.9\%$ and $99.8-99.4\%$ with security margin and between $88.4-92.9\%$ and $98.4-97.5\%$ without the security margin (Fig 3b).

Segmentation examples are shown in Figure 4 using a threshold for P_{vessel} of 0.4. The top row shows that the proposed method (right) is able to capture more vessels than the purely intensity-based method (left). For the example shown in the lower row even a reduction of the threshold for P_{vessel} to 0.2 for the intensity-based method would only cause severe leakage into non-vessel structures but not the inclusion of the missed vessel. The vessel segmentation without vesselness filters typically requires 30 seconds. The addition of the filter responses adds about 20 seconds (per scale) to this time on a common PC. However, more efficient implementations involving the GPU reduce this time drastically [14].
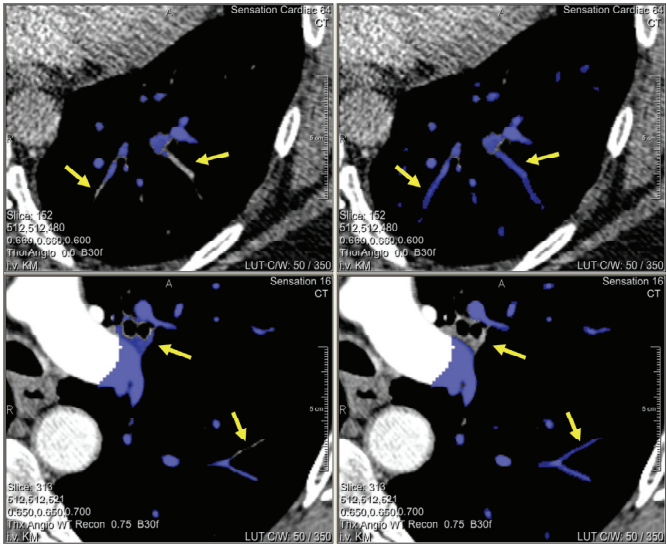
Fig. 4. Segmentation result using the proposed method (right) in comparison to the intensity-based method (left) for two different patients (best viewed in color). One can observe that the new segmentation extends further to the periphery than the purely intensity-based one.

6   Summary and Conclusions

We have presented a vessel segmentation method combining a minimal parameter vesselness filter with fuzzy connectedness. The method combines the filter output with an intensity model in our fully automatic fuzzy approach to pulmonary vessel segmentation in contrast enhanced CT data. One can observe that using intensity features alone the optimal thresholds band suited for large vessels does not overlap with the band for small vessels. Hence such methods tend to leak into surrounding non-vessel structures in close proximity to larger vessels while missing smaller vessels. The results show that using the combination improves the DSC by 9% (from 81% to 90%) for small vessels without influencing the accuracy for large vessels (95%) compared to the intensity-based method alone. Additionally we were able to limit the number of required vesselness filter scales to three, resulting in a computationally efficient method. The result is a method with the benefits of the vesselness filter, i.e., accuracy, and the intensity model, i.e., speed, without the drawback of each method individually. In the future we plan to extend the validation to more datasets, especially also to patients with other lung deseases.

References

1. Zhou, C., Chan, H.P., et al.: Automatic multiscale enhancement and segmentation of pulmonary vessels in CT pulmonary angiography images for CAD applications. Med. Phys. **34**(12) (2007) 4567–4577
2. Kiraly, A.P., Novak, C.L., et al.: A comparison of 2D and 3D evaluation methods for pulmonary embolism detection in CT images. In: SPIE MI. Vol. 6146. (2006) 132–140
3. Masutani, Y., MacMahon, H., Doi, K.: Computerized detection of pulmonary embolism in spiral ct angiography based on volumetric image analysis. IEEE TMI **21**(12) (2002) 1517–1523
4. Kiraly, A.P., Pichon, E., et al.: Analysis of arterial subtrees affected by pulmonary emboli. In: SPIE MI. Vol. 5370. (2004) 1720–1729
5. Agam, G., Armato III, S.G., Wu, C.: Vessel tree reconstruction in thoracic CT scans with application to nodule detection. IEEE TMI **24**(4) (2005) 486–499
6. Geiger, B., Kiraly, A.P., et al.: Virtual bronchoscopy of peripheral nodules using arteries as surrogate pathways. In: SPIE MI. Vol. 5746. (2005) 352–360
7. Kiraly, A.P., Odry, B.L., et al.: Computer-aided diagnosis of the airways: Beyond nodule detection. Journal of Thoracic Imaging **23**(2) (2008) 105–113
8. Buelow, T., Wiemker, R., et al.: Automatic extraction of the pulmonary artery tree from multi-slice CT data. In: SPIE MI. Vol. 5746. (2005) 730–740
9. Wu, C., Agam, G., et al.: Regulated morphology approach to fuzzy shape analysis with application to blood vessel extraction in thoracic CT scans. In: SPIE MI. Vol. 5370. (2004) 1262–1270
10. Masutani, Y., MacMahon, H., Doi, K.: Automated segmentation and visualization of the pulmonary vascular tree in spiral CT angiography: An anatomy-oriented approach based on three-dimensional image analysis. Journal of Computer Assisted Tomography **25**(4) (2001) 587–597
11. Kaftan, J.N., Kiraly, A.P., et al.: Fuzzy pulmonary vessel segmentation in contrast enhanced CT data. In: SPIE MI. Vol. 6914, 69141Q. (2008)
12. Frangi, A.F., Niessen, W.J., et al.: Multiscale vessel enhancement filtering. Lecture Notes in Computer Science **1496** (1998) 130–137
13. Sato, Y., Nakajima, S., et al.: Three-dimensional multi-scale line filter for segmentation and visualization of curvilinear structures in medical images. Medical Image Analysis **2**(2) (1998) 143–168
14. Erdt, M., Raspe, M., Sühling, M.: Automatic hepatic vessel segmentation using graphics hardware. In: 4th Int Workshop on Medical Imaging and Augmented Reality (MIAR). (2008) to appear.
15. Shikata, H., Hoffman, E.A., Sonka, M.: Automated segmentation of pulmonary vascular tree from 3D CT images. In: SPIE MI. Vol. 5369. (2004) 107–116
16. Hu, S., Hoffman, E.A., Reinhardt, J.M.: Automatic lung segmentation for accurate quantitation of volumetric x-ray ct images. IEEE TMI **20**(6) (2001) 490–498
17. Udupa, J.K., Samarasekera, S.: Fuzzy connectedness and object definition: theory, algorithms, and applications in image segmentation. Graph. Models Image Process. **58**(3) (1996) 246–261
18. Grady, L.: Random walks for image segmentation. IEEE Trans on PAMI **28**(11) (2006) 1768–1783
19. Grady, L., Schiwietz, T., et al.: Random walks for interactive organ segmentation in two and three dimensions: Implementation and validation. In: Proc of MICCAI. LNCS 3750 (2005) 773–780

Analysis of 4D CT cine images for the characterization of organ motion due to breathing

Maria Francesca Spadea[1,2], Marta Peroni[2], Marco Riboldi[2], Guido Baroni[2],
George TY Chen[3], Gregory Sharp[3]

[1] Department of Experimental and Clinical Medicine, University of Magna Graecia,
Catanzaro, Italy, [2] Bioengineering Department, Politecnico di Milano University, Milano, Italy,
[3] Massachusetts General Hospital – Harvard Medical School, Boston, MA, USA

Abstract. A semi-automatic procedure to correlate the motion of lung tumor points in 4D-CT images with a respiratory surrogate signal is presented. Data analysis was performed to characterize of the robustness of external/internal correlation properties in the clinical framework of gated radiotherapy treatments. A cross-correlation based algorithm was implemented to perform template matching for tracking the spatial movement of tumor's points in 6 patients. A graphical interface was developed to allow users to navigate through un-binned CT images. The detected internal movement of features in 3D was then retrospectively synchronized with the RPM signal, and the correlation index R^2 was computed. Results also include the range of motion of selected points, and the prediction error. The developed procedure allowed a fast analysis for external/internal correlation of lung anatomy. The study is generally relevant for all the treatments in which organ motion compensation and control is an issue.

Keywords: organ motion, 4D-CT, lung tumor, external/internal correlation.

1 Introduction

Organ motion due to breathing is an issue in different medical treatments such as radiotherapy [1], [2] and robotic assisted surgery [3], [4], where the main goal is to improve the accuracy of therapeutic procedures while being minimally invasive. Problems occur at different levels: first of all, image artifacts on diagnostic images do not provide consistent information for treatment planning; secondly, the know how transfer from treatment plan to intra-operative environment can be not obvious if the treated position and the expected position of an internal target differ.

For instance, in the radiation therapy framework, the actual delivered dose to the patient might diverge from the planned dose because of anatomical changes during the beam delivery in a treatment fraction (intra-fraction) or in between fractions (inter-fraction), as compared to the time of treatment planning.

Recently, conventional radiotherapy and surgery have moved in the direction of Image Guided Radiotherapy (IGRT) [5] and Image Guided Surgery (IGS) [6], for which on-line tools have been developed by many companies to follow the moving organs during a treatment and from one fraction to the next one. Time-resolved imaging techniques, such as 4-Dimensional Computer Tomography (4D-CT), are increasingly emerging as optimal strategy to overcome motion artifacts and limitations in diagnostic imaging and treatment planning/delivering (4D imaging, 4D treatment planning, 4D dosimetry). In this imaging modality, multiple images are acquired during the respiratory cycle and retrospectively sorted into volumetric image sets corresponding to different breathing phases. A respiratory signal (usually the motion of external surface) must also be acquired and synchronized with the image formation process. Hence, 4D data can be analyzed to determine the mean tumor position, tumor range of motion for treatment planning, the relation of tumor trajectory to other organs and to extend the static treatment plan to the all phases of the respiration cycle. However, a limitation of 4D CT is that it is affected by variations in respiratory patterns during acquisition. A second issue arises in the correlation between external fiducial movement and tumor/organ motion. The relation of between the motion of internal and external anatomy need to be assed to exploit the benefits of both IGRT and IGS techniques.

This work follows a previous pilot study [7], aiming at assessing the correlation between the motion of external fiducials and internal features in lung tumors on the basis of 4-Dimensional Computer Tomography (4D-CT) data. Previously, vessel bifurcations were selected as inner targets to be correlated with RPM signal. In this case, the motion of tumor points were tracked. Data analysis was conducted to extract the Spearman correlation coefficient, the range of motion of the tumor, the prediction error in tracking the internal motion from the external fiducial monitoring.


2 Material and Methods

The 4DT-CT data of a group of 6 lung patients (pt) were used in this study. Images were acquired through a 4-slice scanner (LightSpeed QX/i, GE Medical Systems, Milwaukee, WI) operating in axial cine-mode. Patients' breathing signal was provided by the Real-Time Position Management system (RPM, Varian® Medical System, Palo Alto CA), in which a single surface surrogate (infra-red marker block) is tracked in real-time with a video camera positioned at the foot of the CT couch [8]. The RPM block was placed abdominal surface of the patient. According to protocol in use at our Institute, described in detail in Rietzel *et al* [8], a pre-determined number N of 4-slice chunks (Ch_i with $i=1…N$) was acquired at each couch position (CP). Images were reconstructed from 360° projections, which require either 0.8 or 1.0 seconds to acquire (T_r). For analysis purposes, we assume the acquisition occurs instantaneously at the mid-scan time, half way through the full rotation. The CP time duration (cine duration, T_{cine}) was set on the basis of the observed subject's breathing period plus the time required for a tube rotation. Therefore, the number N varied among patients as a function of the T_{cine}, T_r and of the midscan time delay (ΔMidscan time) between 2 contiguous images in a CP, according to the following equation :

$$N = \frac{\left(T_{cine} - T_r\right)}{\Delta Midscan} + 1 \qquad \qquad (1)$$

The acquisition of projections at each CP starts independently of the current respiratory state and therefore no correspondence exists between Ch_i at different couch positions. At each couch position, the coverage in superior-inferior (SI) direction was 1 cm (2.5 mm thickness per 4 slices). The resolution both in latero-lateral (LL) and anterior-posterior (AP) directions ranged from 0.76 mm to 0.98 mm.

2.2 Images processing: automatic feature matching algorithm

The time stamp of each image (midscan time) was time synchronized with the respiratory signal acquired by RPM system. Un-binned images were processed and sorted on the basis of acquisition time, and matched with the RPM. Anatomical landmarks in the lungs were selected in 4D-CT cine-data, and correlated with the respiratory signal.
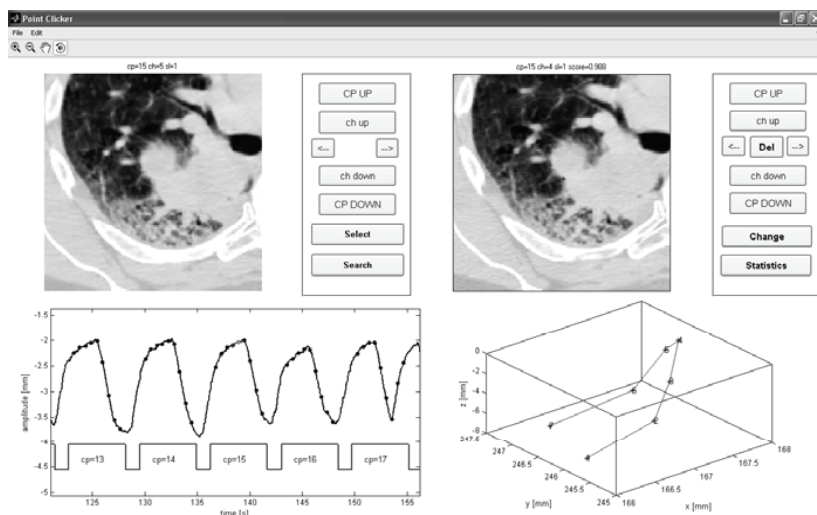

Fig. 1. View of the graphical user interface of the application. Upper left quarter, navigation panel and image for point selection. In the bottom left quarter, RPM trace with indication of the instant of time of the selected point keeps track of the position of the automatic searched point while navigating in the volume. Upper right quarter, navigation panel and image for verification of the semi-automatic procedure. In the bottom right quarter, 3D trajectory of the tracked point is updated in real-time according to navigation throughout the volume and user manual corrections.

A semi-automatic procedure for inner feature tracking in 4D-CT was developed in Matlab environment (MatLab® version 7.0, the MathWorks, Natick, MA). The user first chooses a point on a selected image (reference image) and then decides which

slices to include in the tracking process. A graphical interface assists the user in navigation within the 4D-CT volume data (see Fig. 1).

Through a virtual navigation keyboard one can select the CP, the Ch_i, and the slice within the couch position (sl_i with i=1...4, indicated above the panel). The selected point is marked with a green dot on the image, and its corresponding external position within the RPM respiratory trace is shown below, thus helping the user in choosing the couch positions to be included in the feature searching. This is shown on the left of Fig. 1, where a point is selected on the third slice of a 4-slice set acquired at the maximum inhale condition, and the green dot in the RPM plot below shows the external amplitude. The recommended automatic search region would therefore consist of the two previous couch positions plus the current one, since it is expected that the diaphragm will move superiorly. Feature tracking was based on a 2-D template matching image process. The template was defined as a matrix of 62x62 pixels centered in the point selected on the reference slice. The cross-correlation matrix between the template and the searching window (92x92 pixels) was computed for each slice encompassed in the selected volume. The pixel location (in terms of row and column indexes of the image searching window) in each processed image corresponding to the maximum value of the correlation coefficient (score) was selected as result of the template matching. Hence, a correlation score (ranging from 0 to 1) was obtained for each slice at every respiratory phase in correspondence of the three couch positions. A second image (on the right in figure 2) is produced for the verification of the semi-automatic procedure, which is again assisted by interactive navigation. For each chunk the slice with maximum score (among the 4 ones) is presented. The user may accept the result, change the location of the feature position along the LL and AP directions and/or within the 4 slices or delete the selected feature if it is not present. The 3D trajectory of the point is showed below the verification CT slice. To increase the resolution along the superior-inferior (SI) direction, a 2^{nd}-order spline, interpolating the best score values in the 4 slices of each phase, was computed. An example of the interpolation procedure is showed in Fig. 2. Feature SI-location was selected according to the maximum of the interpolant function.
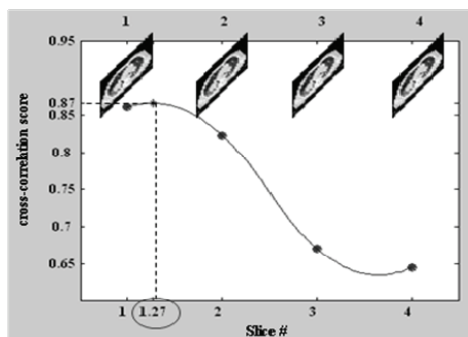

Fig. 2. Example of spline interpolation between slices. In this case scores were 0.864, 0.829, 0.675, 0.650 for slice 1, 2, 3, 4 respectively. Maximum value of the interpolant was found between slices 1 and 2

2.3 Data analysis

For each patient 3 points (each one in a different couch positions, from head to feet) on the edge of the cancer mass were tracked. The detected internal movement of tumor's points was correlated with the RPM signal and analyzed for linear correlation. A dedicated section in the algorithm was implemented to analyze data in terms of the motion amplitude and the correlation index R2 of the linear fitting over the entire patient population. In order to measure the prediction error of the internal motion based on the external signal, a cross validation leave-one-out was implemented. This analysis was enriched by a 3-dimentional linear fitting to understand if a 2 parameter model based on external motion amplitude and its gradient can better fit the internal movement (see Fig. 3).

The results were supported by non parametric statistical analysis computed by means of the software Statistica 6.0 (StatSoft Inc, Tulsa OK, USA).
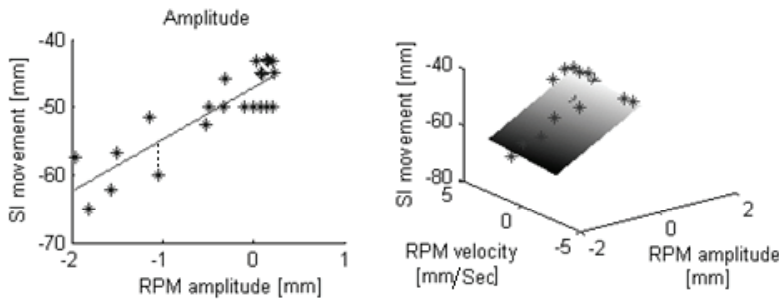
Fig. 3. Example of the line and plane fitting to extract internal external correlation.

4 Results

In Figure 4, the correlation index R^2 between the external and the internal motion in each direction is shown (p-value<0.01). Maximum values of R^2 (95% of confidence interval) are in SI direction, being the median±quartile 0.93±0.09 . Internal motion along LL and AP direction featured low correlation with the movement of the external point (0.51±0.52 and 0.66±0.32 in LL and AP respectively).
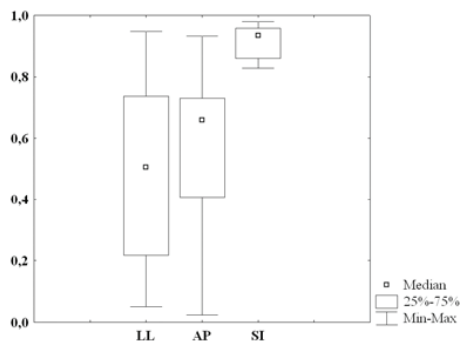
Fig. 4. Median±quartile and min-max range of R2 in each direction.

Statistical difference between R2-SI and R2-LL/AP was proved through Kruskall-Wallis test. In table 1, the range of motion averaged on the set of three selected points for each subject is presented. Patient 3 exhibited largest tumor motion in AP and SI directions, being up to 8 mm and 17 mm respectively. LL direction was most stable except for Patient 2 (up to 6 mm about).

Table 1. Estimated range of motion of the tumor obtained through the analysis of the movement of 3 points. Mean (standard deviation) values over three points are reported.

Patient #	LL (mm)	AP (mm)	SI (mm)
P1	1.94 (0.22)	2.21 (0.98)	6,08 (1.33)
P2	5.45 (0.85)	2.16 (1.12)	8,17 (2.37)
P3	2.67 (0.79)	7.16 (0.69)	16,28 (0.62)
P4	1.16 (0.58)	0.71 (0.38)	8,31 (1.40)
P5	0.26 (0.23)	3.32 (0.52)	6,05 (1.18)
P6	0.91 (0.49)	3.26 (1.13)	6,54 (1.73)

In a effort to understand repeatability of tumor motion over different breathing cycles, we compared the trajectory of the three points of the tumor. Patient 3 is reported in Fig. 5 as an example: in this case, the three points were tracked on couch position number 15, 16 an 18.
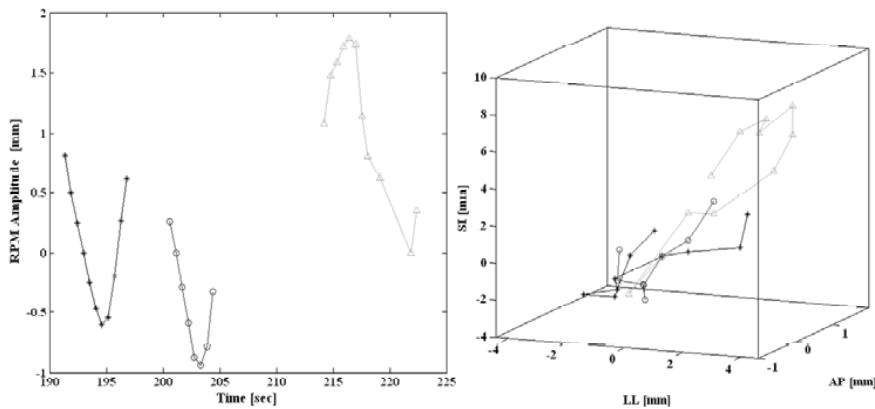
Fig. 5. Trajectory of three points selected at the beginning (*), middle (○) and towards the end (⅄) of the lesion of Patient 3.

Points 1 and point 2 were visible only in few phases for the selected respiration cycle, as is observable in the left panel of the figure. Extensions over contiguous CP were not considered in this analisys Qualitative examination demonstrates a quite good repeatability of the trajectory of points belonging to the same anatomical structure in different breathings, especially in the SI direction.

Considering the prediction error, the RPM amplitude-based (2D) prediction was compared to the RPM amplitude and velocity-based prediction (3D) and results are reported in table 2. Wilcoxon matched pair test revealed no statistical difference between the two methods. Maximum values were found in SI direction of Patient 3 (about 2 mm).

Table 2. RPM amplitude-based (2D) prediction and RPM amplitude and velocity-based prediction (3D)

	Prediction Error 2D			Prediction Error 3D		
	LL (mm)	AP (mm)	SI (mm)	LL (mm)	AP (mm)	SI (mm)
25%	0,26	0,30	0,50	0,25	0,29	0,43
Median	0,34	0,48	0,74	0,31	0,43	0,60
75%	0,37	0,73	1,14	0,34	0,69	1,03

We put forward the hypothesis that there is a relation between the size of the prediction error and the size of the range of motion. This is supported by results displayed in Fig. 6 (2D prediction error vs. range of motion). Although a good linear relation cannot be demonstrated, the increasing trend of prediction error as a function range of motion is evident (p value<0.001).
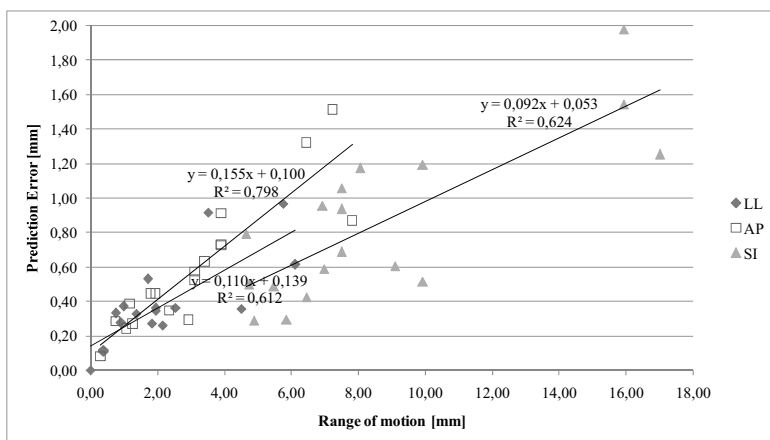
Fig. 6. Analysis of the relation between range of motion and prediction error.

Discussion and Conclusion

In this work, we present a semi-automatic and fast method for characterizing the correlation between the external surface movement and internal tumor motion using the same data and images that are routinely acquired for 4D-CT. Most of literature about this topic in the radiotherapy field is based on data acquired by using fluoroscopic imaging [9], [10], [11]. Although these studies sufficiently describe the tumor motion and the quality of external/internal correlation, they are not based on a protocol widely spread into the clinical practice. Furthermore, the monitoring of internal movements by using fluoroscopic images requires the implantation of seeds, which represents a serious hazard for lung patients. Instead, 4D-CT based planning and treatment procedures are currently well-accepted, and RPM is the most widespread device used as surrogate of the breathing signal. For this reason it is crucial to understand the sensitivity of this instrument, and to evaluate how well it can estimate internal organ motion. The advantages of using cine-mode protocol, especially for investigations purposes, are widely described [8]. Because this analysis is based on the synchronization of cine mode images with the respiratory signal, no inaccuracies in phase detection and binning were introduced.

Our results revealed a good correlation between the internal SI motion and the external respiratory surrogate. Low correlation in the other directions is probably due to small range of motion in LL and AP. The size of range of motion also influence the size of the prediction error as shown in Fig. 6. However, the maximum value was about 2 mm that might be considered acceptable as extreme limit for gated radiotherapy.

In conclusion, we have developed a automatic procedure and analysis tools for studying the correlation between the internal and external motion in lung tumors, using data already routinely acquired for 4D-CT. This tool will aid in predicting internal motion through external surrogates. Further research is needed to improve the accuracy and spatial coherence of internal anatomy in 4D CT data acquisition

References

1. Goitein, M. Organ and tumor motion: an overview. Semin. Radiat. Oncol. 14(1), 2-9 (2004)
2. Keall, P.J., Mageras, G.S., Balter, J.M. et al. The management of respiratory motion in radiation oncology report of AAPM Task Group 76. Med. Phys. 33(10),
3. Müller, S.A., Maier-Hein, L., Mehrabi, A., et al. Creation and establishment of a respiratory liver motion simulator for liver interventions. Conf. Proc. IEEE Eng. Med. Biol. Soc. 2614-2617 (2007)
4. Lesniak, J., Tokuda, J., Kikinis, R., Burghart, C., Hata, N. A device guidance method for organ motion compensation in MRI-guided therapy. Phys. Med. Biol. 52(21), 6427-6438 (2007)
5. Xing L, Thorndyke B, Schreibmann E, Yang Y, Li TF, Kim GY, Luxton G, Koong A. Overview of image-guided radiation therapy. Med. Dosim. 31(2):91-112 (2006).
6. Marescaux, J., Solerc, L. Image-guided robotic surgery. Semin Laparosc Surg. 11(2), 113-122 (2004)
7. Spadea, M.F., Riboldi, M., Baroni, G., Chen, G.T.Y., Sharp, G. A Feature matching approach for the automatic correlation of internal and external motion in lung tumors. Med. Phys. 34 (Suppl. 1): 2390-2391 (2007)
8. Rietzel E, Pan T, Chen CTY. Four-dimensional computed tomography: Image formation and clinical protocol. Med. Phys. 32(4), 874-888 (2005)
9. Gierga, D.P., Brewer, J., Sharp, G.C., Betke, M., Willett, C.G., Chen, G.T. The correlation between internal and external markers for abdominal tumors: implications for respiratory gating. Int J Radiat Oncol Biol Phys. 61(5), 1551-1558 (2005)
10. Hoisak, J.D.P., Sixel, K.E., Tirona, R., Cheung, P.C.F., Pignol, J.P. Correlation of lung tumor motion with external surrogate indicators of respiration. Int J Radiat. Oncol. Biol. Phys., 60(4):1298:1306 (2004).
11. Yan H, Yin FF, Zhu GP, Ajlouni M, Kim JH. The correlation evaluation of a tumor tracking system using multiple external markers. Med Phys. 2006 Nov;33(11):4073-4084.

Robust segmentation of pulmonary nodules of various densities: from ground-glass opacities to solid nodules

Toshiro Kubota[1], Anna Jerebko[2], Marcos Salganicoff[2], Maneesh Dewan[2], and Arun Krishnan[2]

[1] Department of Mathematics and Computer Science, Susquehanna University
[2] CAD and Knowledge Solutions, Siemens Medical Solutions, Inc. USA

Abstract. Accurate segmentation of a pulmonary nodule is an important and active area of research in medical image processing. Although many algorithms have been reported in literature for this problem, those that are applicable to various density types have not been available until recently. In this paper, we propose a new algorithm that is applicable to solid, non-solid and part-solid types and solitary, vascularized, and juxtapleural types. The algorithm works on the distance map computed from a foreground image. It first locates the core of a nodule in a manner that is robust against the presence of attached structures such as pleura and vessels, and then detaches the nodule from the attached structures by a variant of region growing and region partition.

The performance of the technique is evaluated using 23 data sets with manual segmentation and 1521 data sets with manual diameter measurements. The experiments show that the algorithm is highly reliable in segmenting convex nodules of various types.

1 Introduction

Pulmonary nodules are potential manifestations of lung cancer, and their detection and inspection are essential for screening and diagnosis of the disease. The growth of a nodule is considered the most important cue for assessing its malignancy.

Various segmentation methods targeted for pulmonary nodules have been developed, and some have been deployed in commercial applications. Many technical issues still remain, including accuracy and handling of non-solid and part-solid nodules. Most of existing segmentation algorithms are targeted toward nodules composed only with solid components, although studies have shown that nodules of non-solid and part-solid types are frequent and have higher risks of being malignant than solid ones [1]. More recent segmentation methods reported in literature claim to handle these types of nodules[2–4]. However, we believe that the field is relatively new and requires further investigation.

The goal of our work is to develop a semi-automated nodule segmentation algorithm that is applicable to not only the solid type but also the non-solid

and part-solid types. Figure 1 shows examples of pulmonary nodules of different
density types. As exemplified in the figure, the appearance varies among different
nodule types. In particular, non-solid nodules are extremely subtle with fuzzy
boundaries, and part-solid nodules exhibit highly irregular intensity variations
and boundary shapes. Thus, handling them under a single framework presents
a great challenge to the segmentation problem. The inputs to the algorithm are
a thin-slice thoracic CT volume and a click point. The output is a segmentation
map of a nodule found in the vicinity of the click point. One assumption we make
is that the nodule is convex. Thus, we are interested in segmenting a convex part
of the nodule.

(a)   (b)   (c)   (d)   (e)   (f)
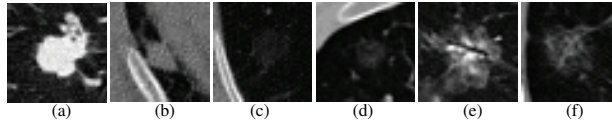
Fig. 1. Examples of pulmonary nodules of various densities. (a)-(b): Solid nodules,
(c)-(d): Non-solid nodules, (e)-(f): Part-solid nodules.

The important feature of our algorithm is twofold. First, it locates the core
of a nodule using a new approach that is highly robust against the presence of
attached structures such as pleura and vessels and is invariant to the size of
the nodule. The problem of segmenting a juxtapleural nodule lies in detecting
a boundary between the nodule and the lung wall. We take the problem in two
steps: locating the core of the nodule and growing a region from the core. Insen-
sitivity to attached structures is a key in achieving the first step. Many previous
techniques search both location and size of a nodule either in iterative manners
or in scale-space representations [5, 6]. The size-invariance property of our ap-
proach makes it unnecessary to estimate the size of the nodule, thus renders the
localization process computationally more efficient. Second, it employs a two-
step region growing process as described in Sections 2.3 and 2.4, where the first
step over-segments the nodule and the second step refines the segmentation via
a competition process. The approach extracts a convex nodule from attached
structures such as lung walls and vessels.

Throughout the paper, the following notations are used. A lower bold letter
($\mathbf{x}$) is used for a vector, an upper bold letter ($\mathbf{X}$) is used for a 3D volume, and a
non-bold letter (x) is used for a scalar. Without ambiguity, we also use an upper
bold letter associated with a binary volume (eg. segmentation) to indicate those
voxels whose binary values are non-zero. For example, we use $\mathbf{L}$ to denote a set
of voxels that are foreground. Furthermore, we use $\neg\mathbf{L}$ to denote those voxels
whose binary values are zero in $\mathbf{L}$.

2   Algorithm

2.1   Overview of the algorithm

The segmentation algorithm consists of six stages: pre-processing, figure-ground separation, localization of a nodule core, region growing, region partition, and post-processing. The first two stages are briefly discussed in this section. The other four are discussed, respectively, in the following sub-sections.

At the pre-processing module, a sub-volume that is large enough to contain a nodule of interest is extracted. The sub-volume is then resampled at a fixed interval to produce an isotropic sub-volume, called a region of interest (ROI). We denote a set of voxels in ROI as Ω. The figure-ground separation applies a bi-class segmentation algorithm previously developed by the authors to the ROI[3]. The result of the operator is a binary map denoted as $\mathbf{L} : \Omega \to \{0, 1\}$, where 0 and 1 voxels indicate backgrounds and foregrounds, respectively.
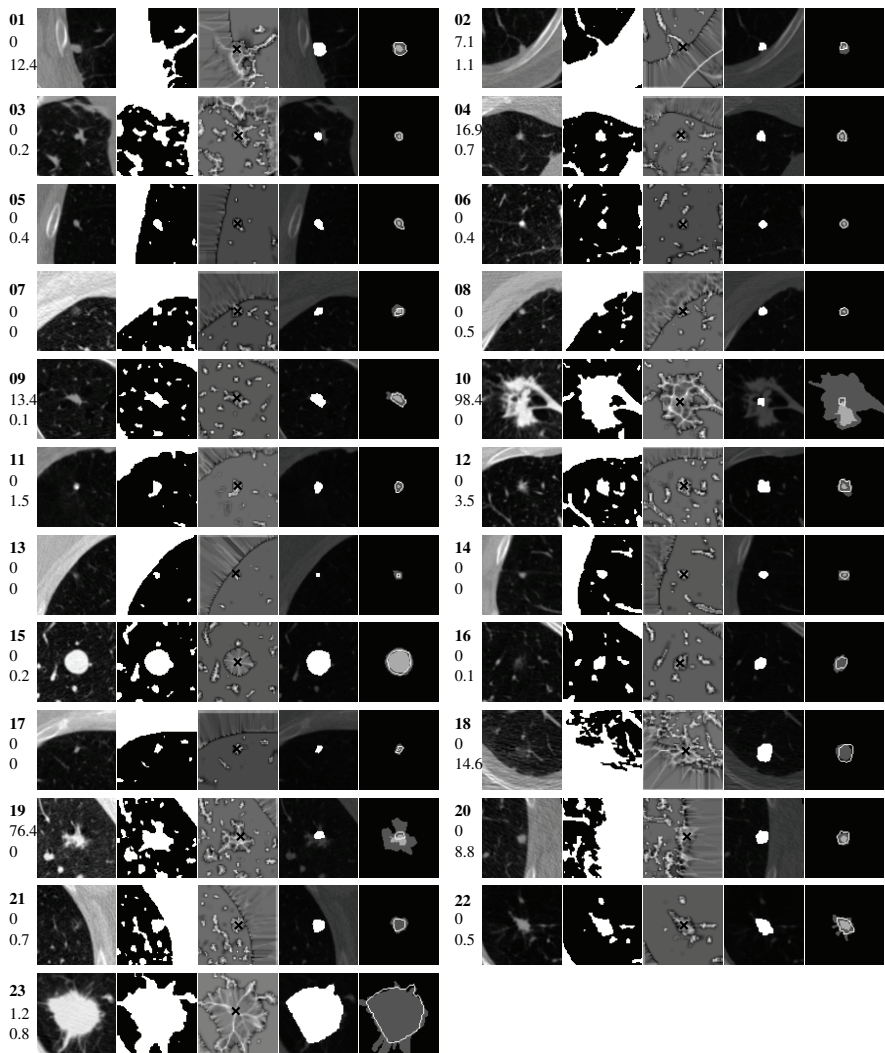
Figure 2 shows results of the figure-ground separation process applied to 23 data sets provided by the NIH Lung Imaging Database Consortium (LIDC)[7]. For each data, five images are shown. The left image is a representative axial slice of the nodule in the original CT data. The second image from the left is the result of the figure-ground separation. The other three images are results of seed point localization and segmentation processes, which are discussed in detail below. A bold number accompanying each set of images is a unique number for identifying the data, and two non-bold numbers indicate the accuracy of the segmentation. We will describe the accuracy measures in Section 3. Note that the images are shown at the resolution of the original CT data. Thus, all results shown are interpolated back to the resolution of the original CT data.

2.2   Localization of Nodule Center

After the figure-ground separation, our next task is to locate a core of a nodule inside the ROI. We first apply Euclidean distance transform to $\mathbf{L}$ to compute for each location in $\mathbf{L}$ the shortest Euclidean distance to the background. Call the resulting distance map $\mathbf{D}$. Our motivation for working on $\mathbf{D}$ is twofold. First, $\mathbf{D}$ is free from signal noise present in the original intensity volume. Thus, no noise removal and signal restoration processes are needed in the subsequent processes. Furthermore, it renders the rest of the algorithm insensitive to intensity variations of images. This is important as solid, non-solid and part-solid nodules present significantly different intensity distributions. Second, $\mathbf{D}$ compactly encodes information regarding the shape of the foreground. For example, $\mathbf{D}$ tells how close each voxel is to the background. It provides, by means of the gradient vector computed on $\mathbf{D}$, a rough direction to the background.

A limitation of $\mathbf{D}$ is that it is extremely sensitive to a small hole present in the foreground and partially sensitive to an attached structure. The foreground extraction procedure described above effectively removes small holes in $\mathbf{L}$. We will discuss how to handle the attachment next.

---

[3] For double blind review, details of our previous work are intentionally omitted.

Fig. 2. Segmentation results on LIDC data set. For each data set, five images are shown. They are from left to right: original slice, figure-ground separation map, sphericity map, segmentation result, and comparison with manual segmentation. In the rightmost image, a dark gray color shows the area selected by at least one segmentation, a light gray color shows the area selected by 90% of the methods, and a white line shows the boundary of our segmentation result. If no light gray area is shown, the slice did not contain an area agreed by 90% of the methods. Three numbers on the side of each image set are a unique identification number, under-segmentation fraction percentage, and over-segmentation fraction percentage.

When the nodule is modeled as a sphere and more than a half of the nodule is buried in the lung wall, the local maximum of $\mathbf{D}$ no longer resides inside the nodule. To mitigate the problem, we propose another transformation built on top of $\mathbf{D}$ as defined below.

$$\mathbf{S}(\mathbf{x}) = \int_{\mathbf{y} \in N(\mathbf{x})} \chi \left(\mathbf{D}(\mathbf{x}) - \mathbf{D}(\mathbf{y}) \right) d\mathbf{y} \tag{1}$$

where $N(\mathbf{x})$ is a ball of radius δ centered at $\mathbf{x}$ and $\chi(x) = 1$ if $x > 0$ and 0 otherwise. We call $\mathbf{S}$ *sphericity*.

It is convenient to introduce the following definition of $\mathbf{C}(d)$:

$$\mathbf{C}(d) = \{ \mathbf{x} \in \Omega | \mathbf{D}(\mathbf{x}) \leq d \}. \tag{2}$$

We call $\mathbf{C}(d)$ the *level core* of $\mathbf{D}$ at d. The next lemmas indicate that $\mathbf{S}$ can replace $\mathbf{D}$ while providing more robust protection against a pleural attachment.

Lemma 1. *For a convex foreground, D has a single local maximum component, and for every d, $C(d)$ is convex.*

Lemma 2. *For a convex foreground, S has a single local maximum component, whose position is identical to the local maximum of D. Furthermore, with a sufficiently small δ, x is a local maximum of S if x is a local maximum of D.*

Lemma 3. *For a spherical nodule, its center remains a local maximum of S under any degree of partial occlusion by a half plane.*

The first lemma justifies our approach in segmenting a nodule by delineating every level set around the local maximum. The second lemma indicates that $\mathbf{S}$ can be used instead of $\mathbf{D}$ to locate the core of a nodule. The third lemma suggests that the local maximum of $\mathbf{S}$ is more stable against partial occlusion by the half plane than that of $\mathbf{D}$.

Figure 3 compares $\mathbf{D}$ and $\mathbf{S}$ of a sphere under various degrees of occlusion by a half plane. The radius of the sphere is 100 voxel unit. The top row shows instances of $\mathbf{D}$ and the bottom row shows instances of $\mathbf{S}$. For computational simplicity, $N(\mathbf{x})$ is set as a 3x3x3 cube instead of a sphere. From left to right, columns (a-e) show a solid plane penetrating the sphere at (a) 30 voxels above, (b) 0 voxel above, (c) 30 voxels below, (d) 60 voxels below, and (e) 90 voxels below the sphere center, respectively. When the plane penetrates at below the center, the local maximum of $\mathbf{D}$ shifts to inside the wall while the local maximum of $\mathbf{S}$ retains its position regardless of the degree of partial occlusion. Although Lemma 3 is limited to a sphere, robustness of $\mathbf{S}$ in detecting the nodule center has been observed for other more general shapes.

Note that the sphericity value at the nodule center is 1 regardless of the size of the nodule when $N(\mathbf{x})$ is completely contained inside the nodule. With our setting of 3x3x3 voxels for $N(\mathbf{x})$, this implies that the sphericity value is invariant to the nodule size no smaller than 3x3x3 voxels.

To locate a nodule center, we compute **S** and locate a local maximum of **S** in a neighborhood of the click point. In our experiment, a 7x7x7 sub-volume centered at the click point constitutes the neighborhood. Figure 2 shows sphericity maps and seed points resulted from applying the above procedure to 23 LIDC data sets. The middle image in each set of five images is the axial slice of **S** at the detected seed point. The seed point is shown by a black cross.

(a)   (b)   (c)   (d)   (e)
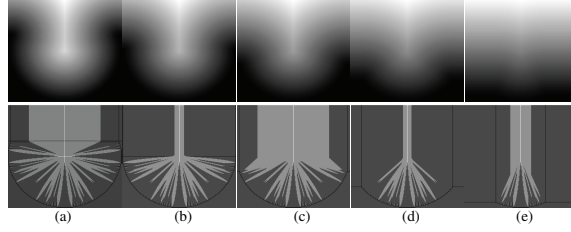
Fig. 3. Distance (top) and sphericity (bottom) maps at various degrees of occlusion.

2.3   Region Growing

Once we locate the core of a nodule, the next task is to extract the nodule from the foreground. We apply region growing on **D** starting from the seed point detected in Section 2.2 and incrementally include each level set to the growing region. The approach can be implemented in the following algorithm. Note that **R** denotes a resulting segmentation map, which is initialized to constant zero.

Input: **D**: Distance map, **s**: seed point
Output: **R**: Segmentation map
$d \Leftarrow \mathbf{D}(\mathbf{s})$;
while $d > 0$ **do**
  foreach ***z*** *that is adjacent to the current region in* ***R*** **do**
    if ***D(z)=d*** **then**
      $\mathbf{R}(\mathbf{z}) \Leftarrow 1$;
  decrement d to the next possible distance value
              Algorithm 1: Region-growing algorithm

The approach is mainly justified by Lemma 1; For a convex isolated nodule, there is a single local maximum component, from which each level set can be included into the region incrementally. For a non-isolated nodule, the foreground may not be convex. As exemplified by Lemma 3, the nodule center can be located by the local maximum of **S**, and we can extract the core of the nodule by Algorithm 1. However, at some point of d, $\mathbf{C}(d)$ becomes non-convex. Including the entire $\mathbf{C}(d)$ into the region results in significant over-segmentation. By

approaching the inclusion of each level set in a region growing manner (that is to include only a single layer around the current region into the region), we can limit the over-segmentation.

Figures 4(a) and (d) show results of Algorithm 1 applied to two artificial data which illustrate juxtapleural and vascularized cases. In both cases, the nodule is a sphere, and the foreground is non-convex due to the attached structure. The black, gray, and white regions are $\neg\mathbf{L}$, $\mathbf{L}$ and $\neg\mathbf{R}$, and $\mathbf{R}$, respectively. Although the algorithm successfully segmented the entire nodule, it also included a portion of the attached structure near the nodule. We leave it to the next step to remove the oversegmented portion.
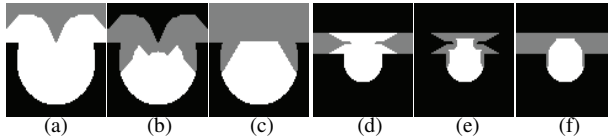

(a)   (b)   (c)   (d)   (e)   (f)

Fig. 4. Segmentation of simple artificial data.

2.4 Region Partition

The objective of this stage is to remove over-segmented parts of the region-growing segmentation. We treat the problem as follows. The segmentation of Algorithm 1 consists of multiple convex regions, among which the one including the seed point is the nodule of interest. Thus, the objective is to isolate the nodule from other convex regions. We take the same strategy as in Algorithm 1, as each convex part has a local maximum of $\mathbf{D}$ and a convex core around it. We need to compute a new $\mathbf{D}$ using $\mathbf{R}$, which we denote as $\tilde{\mathbf{D}}$. Since $\mathbf{R}$ is already isolated from a pleural wall, it is not necessary to carry out computation of $\mathbf{S}$ and find its local maxima; $\tilde{\mathbf{D}}$ and its local maxima are sufficient.

Local maxima in $\tilde{\mathbf{D}}$ are treated as a new set of seed points, which are either *positive* or *negative*. The positive seed points are in the vicinity to the original seed point and contribute toward delineating the nodule part. The negative seed points are those that are not positive, and contribute toward delineating non-nodule components. We use the following steps to select positive seed points. First, the location of the original seed point is moved to the local maximum of $\tilde{\mathbf{D}}$ by tracing the gradient of $\tilde{\mathbf{D}}$. Call this relocated seed point $\tilde{\mathbf{s}}$. Then, those local maxima of $\tilde{\mathbf{D}}$ that are within $\tilde{\mathbf{D}}(\tilde{\mathbf{s}})$ away from $\tilde{\mathbf{s}}$ are classified positive.

Once we collected both positive and negative seed points, we perform a region growing on $\tilde{\mathbf{D}}$ from each seed point in down-hill directions. A region grown from a positive seed point is marked as nodule, while a region grown from a negative seed point is marked non-nodule. The process continues until no further growth is possible. Algorithm 2 summarizes the process. $\mathbf{P}$ denotes the result, which is initialized to constant zero.

Input: $\tilde{\mathbf{D}}$: distance map, $\mathbf{s}^+$: positive seed points, $\mathbf{s}^-$: nagative seed points
Output: $\mathbf{P}$:Partition map
$\mathbf{P}(\mathbf{s}^+)$=1; $\mathbf{P}(\mathbf{s}^-)$=-1;
repeat
    foreach $\boldsymbol{x}$ *in* $\{\boldsymbol{z}|P(\boldsymbol{z}) = 0\}$ **do**
        if $\boldsymbol{x}$ *has a positive neighbor* $\boldsymbol{y}$ & $\tilde{\boldsymbol{D}}(\boldsymbol{x}) \leq \tilde{\boldsymbol{D}}(\boldsymbol{y})$ **then**
            $\mathbf{P}(\mathbf{x})$=1;
        if $\boldsymbol{x}$ *has a negative neighbor* $\boldsymbol{y}$ & $\tilde{\boldsymbol{D}}(\boldsymbol{x}) \leq \tilde{\boldsymbol{D}}(\boldsymbol{y})$ **then**
            $\mathbf{P}(\mathbf{x})$=-1;
until *No change to $\boldsymbol{P}$* ;
Set all negative voxels in $\mathbf{P}$ to 0;

Algorithm 2: Region-partition algorithm

Results of Algorithm 2 applied to the artificial data of Figure 4(a) and (d) are shown in Figure 4(b) and (e), in which black, gray, and white regions are $\neg\mathbf{R}$, $\mathbf{R}$ and $\neg\mathbf{P}$, and $\mathbf{P}$, respectively. In both examples, most of the nodule is included in $\mathbf{P}$ while non-nodule components are successfully eliminated.

2.5 Post-processing

Segmentation after Algorithm 2 is not necessarily convex. At this final stage, $\mathbf{P}$ is first made convex by convex hull of $\mathbf{P}$, and then taken intersection with $\mathbf{L}$. Figure 4(c) and (f) show results of this stage applied to the artificial data of Figure 4(a) and (d). We call the segmentation after this stage $\mathbf{F}$.

3 Experiments

3.1 LIDC Segmentation

First, we apply the segmentation algorithm to the LIDC data set. The study came from two cancer cases with collected segmentation done by six radiologists using three methods: one fully manual and two with automated programs. Thus, it collected a total of 18 segmentation maps. A sub-volume of 81x81xL voxels is extracted around the nodule where L is the number of slices in the data set. A click point is set at the center of the sub-volume. We also derive two different segmentation maps from the manual segmentation. One is a set of voxels included in at least one out of 18 methods. The other is a set of voxels included in 90% of the methods. The former is denoted as $\mathbf{G}_1$ and the latter $\mathbf{G}_2$.

Figure 2 shows the result of our segmentation for each LIDC nodule. As stated in Section 2.1, each nodule corresponds to a set of horizontally aligned five images where the first is a representative axial slice of the original sub-volume, the second shows $\mathbf{L}$ at the slice, the third shows $\mathbf{S}$ and $\mathbf{s}$. The fourth image shows $\mathbf{F}$ super-imposed on top of the original sub-volume. The fifth is made by first painting $\mathbf{G}_1$ in dark gray, then painting $\mathbf{G}_2$ in light gray, and finally painting the boundary of $\mathbf{F}$ in white. Shown to the left of the image set are

three numbers, which are, from top to the bottom, the unique data identification number (in bold), under-segmentation fraction (p_U) and over-segmentation fraction (p_O). The under/over segmentation fraction percentages are defined as $p_U = 100|\neg\mathbf{F} \cap \mathbf{G}_2||\mathbf{G}_2|^{-1}$ and $p_O = 100|\mathbf{F} \cap \neg\mathbf{G}_1||\mathbf{G}_1|^{-1}$, respectively, where $|\cdot|$ is the cardinality of the set.

3.2 Automated Diameter Measurement

We apply our segmentation algorithm to a larger set of CT data. Since it is difficult to obtain a reliable segmentation on a large set of data, we use the diameter of a nodule to test the accuracy of our segmentation. The test data consist of 1521 nodules from 253 cases, where 1237 are solid, 206 are non-solid, and 77 are part-solid. Each sub-volume after isotropic resampling is 41x41x41 voxels. The ELCAP protocol ([8]) is used to measure the nodule diameter both manually and automatically from the segmentation obtained by our algorithm.

Figure 5(a) shows a scatter plot of the manual and automated diameter measurements. The horizontal axis is the manual diameter measurement and the vertical axis is the automated diameter measurement. Figure 5(b) shows a histogram of normalized errors ($\bar{\epsilon}$) defined as $\bar{\epsilon} = |D_m - D_a|D_m^{-1}$ where D_m and D_a are manual and automated diameter measurements, respectively. Table 3.2 shows descriptive statistics of estimates.
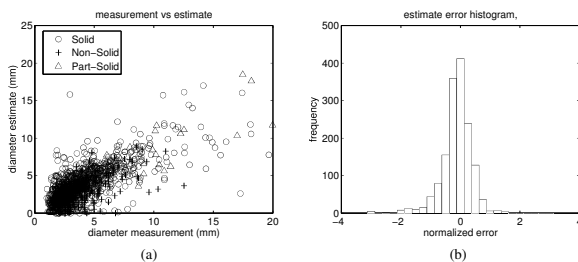
Fig. 5. Diameter measurement results on 1521 nodule data.

4 Discussion and Conclusion

The paper presented a new general purpose segmentation algorithm for a pulmonary nodule. The algorithm is applicable to solid, non-solid and part-solid types and handles juxtapleural and vascularized ones without having the lung separated from the pleural walls. The results on 23 LIDC data and 1521 data with manual diameter measurements are highly encouraging.

A limitation of the approach is that it cannot handle spiculated nodules due to its underlying assumption of nodule convexity. For a spiculated nodule, we

Table 1. Descriptive statistics of the diameter measurement experiment

	ρ^1	μ^2	m^3	σ^4	τ^5
All	0.755	0.359	0.254	0.397	0.607
Solid	0.759	0.372	0.256	0.422	0.603
Non-solid	0.624	0.334	0.285	0.266	0.611
Part-solid	0.864	0.214	0.177	0.163	0.664

[1] ρ: Pearson correlation coefficient.
[2] μ: mean normalized absolute error.
[3] m: median normalized absolute error.
[4] σ: std normalized absolute error.
[5] τ: mean computation time (sec) on 2.8GHz PC with 1G memory.

may be able to treat it as one comprised of multiple convex parts, and extract each part in succession. Another possible solution is to perform segmentation at a coarse resolution where the degree of spicularity is reduced. Another limitation is that it tends to over-segment non-solid nodules. The limitation can be circumvented by referring to the ROI intensity volume and carefully trimming voxels that are likely to be partial volumes.

References

1. Henschke, C., Yankelevitz, D.F., Mirtcheva, R., McGuinness, G., McCauley, D., Miettinen, O.S.: CT screening for lung cancer: Frequency and significance of part-solid and nonsolid nodules. Am. J. of Roent. **178** (2002) 1053–1057
2. Kauczor, H.U., Heitmann, K., Heussel, C.P., Marwede, D., Uthmann, T., Thelen, M.: Automatic detection and quantification of ground-glass opacities on high-resolution ct using multiple neural networks: Comparison with a density mask. Am. J. of Roent. **175** (2000) 1329–1334
3. v. Ginneken, B.: Supervised probabilistic segmentation of pulmonary nodules in CT scans. In: MICCAI. (2006) 912–919
4. Wang, J., Engelmann, R., Li, Q.: Segmentation of pulmonary nodules in three-dimensional CT images by use of a spiral-scanning technique. Med. Phys. **34**(1) (2007) 4678–89
5. Diciotti, S., Picozzi, G., Falchini, M., Mascalchi, M., Villari, N., Valli, G.: 3-d segmentation algorithm of small lung nodules in spiral CT images. IEEE Trans. Inf. Technol. Biomed. **12**(1) (2008) 7–19
6. Okada, K., Comaniciu, D., Krishnan, A.: Robust anisotropic gaussian fitting for volumetric characterization of pulmonary nodules in multislice CT. IEEE Trans. Med. Imaging **24**(3) (Mar. 2005) 409–423
7. Institute, N.C.: Lidc: Datasets as a public resource. http://imaging.cancer.gov/reportsandpublications/reportsandpresentations/firstdataset
8. Henschke, C.: International early lung cancer action program: Protocol. [Online], http://www.ielcap.org/ielcap.pdf (Dec. 2007)

Local Characteristic Features for Computer Aided Detection of Pulmonary Embolism in CT Angiography

Jianming Liang and Jinbo Bi
Siemens Medical Solutions, CKS IKM
Malvern, PA 19355, USA

Abstract. An automated detection system is constructed for detecting pulmonary embolism from computed tomographic pulmonary angiographic images. Our previous work has presented novel effective algorithms to identify suspicious PE regions from images and reduce false detections by designing powerful classifiers. However, these techniques have to take effects in conjunction with discriminative features used to characterize each identified PE candidate. This paper investigates three sets of novel features: 1) features based on local candidate co-occurrence matrices to remove false detections induced by noise and poorly mixed contrast; 2) features characterizing vessel properties to eliminate candidates outside of vessel; 3) features discriminating between arteries and veins to remove candidates from veins. We tested these features in our multiple instance learning classification setting, and they constantly improved the detection accuracy when the 3 sets of features are included sequentially. The resulted PE CAD system has been deployed in clinical settings with capabilities of incrementally reporting any detection immediately once becoming evident during searching, offering real-time support and achieving 85% sensitivity at 5 false positives.

1 Introduction

Computer aided detection of pulmonary embolism (PE) in computed tomographic pulmonary angiographic (CTPA) images has received growing attention in recent years. Pulmonary embolism (PE) is the third most common cause of death in the US, with at least 650,000 cases occurring annually, although treatment with anti-clotting medications is highly effective. Each year more than 400,000 PE cases are not correctly diagnosed, and approximately 100,000 patients die who would have survived with the proper diagnosis and treatment. CTPA has emerged as the first-line diagnostic tool for PE. In CTPA, an embolus appears as dark regions residing in bright vessel lumen as illustrated in Fig. 1. A CTPA volume consists of hundreds slices of the lung. The accurate and efficient interpretation of such a large image volume is complicated by various PE look-alikes and also limited by human factors, such as attention span and eye fatigue. It is highly desirable to have an automated detection system to assist radiologists in detecting and characterizing emboli in an accurate, efficient and reproducible way. Such a CAD system must satisfy stringent real-time requirement since PE cases often occur in emergency room. It must achieve high detection sensitivity with as few false positives as possible to acquire clinical acceptance, because inaccurate use of anti-clotting medications can lead to subsequent hemorrhage and bleeding.

PE detection has been attacked by several groups [1–4]. All the existing methods in the literature require sophisticated vessel segmentation. Vessel segmentation is computationally time-consuming, has been problematic in small vasculature where
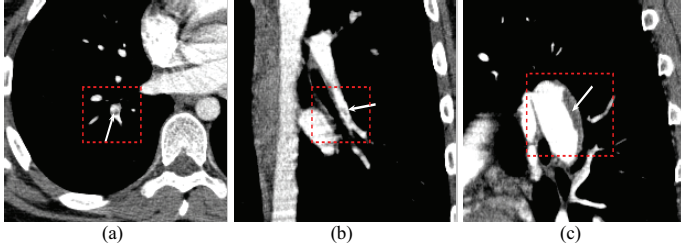
Fig. 1. Acute (a,b) and chronic (c) pulmonary emboli.

subsegmental PEs often occur, and further complicated by PE if exists. None of these approaches can meet the real-time requirement. To circumvent the problem, our early work developed an effective system to search for any suspicious PE regions in the entire lung without segmenting the vessels, and then reduce non-PE regions by distinguishing them from true PE candidates with local characteristic features. Since the number of suspicious PE regions (called candidates) is rather limited (about 50) in each case, and it is much more efficient to compute the features descriptive to the local vascular structure of the candidates than segmenting the whole pulmonary vascular tree.

We investigate three sets of novel local characteristic features in this paper. First, features based on local candidate co-occurrence matrices are calculated to remove false detections due to noise and poorly mixed contrast. Second, Hessian-based features are designed to characterize vessel properties to eliminate candidates outside of vessel. Last, image features discriminating between arteries and veins are computed to remove candidates from veins. These features were tested in the multiple instance classification setting. Together with our candidate generation and classification algorithms, the overall approach forms a comprehensive and clinically usable PE CAD system which is capable of reporting the first detection if any within 30 seconds and achieving 85% sensitivity under 5 false positives on average across a large number of cases. For completion, we first review our concentration oriented tobogganing algorithm for quickly generating suspicious regions for PE incrementally (candidate generation) and our multiple instance learning (MIL) framework for false positive reduction (classification) in Section 2 before presenting our three sets of novel features in Secs. 3, 4 and 5, respectively.

2   Retrospect of candidate generation and classification

2.1   Candidate generation

When the vascular structure is not segmented, a major challenge for automatic PE detection is to effectively separate the emboli from the vessel wall and to quickly remove partial volume effects around the vessel boundaries while correctly preserving the PE pixels, since all the voxels in those areas have the same original CT values as those in the PE regions. To this end, we first reported an approach based on basic tobogganing in [5] with a simple operation called "sliding": A pixel v with intensity $P(v)$ and

Table 1. False positive distributions based a meticulous analysis of training cases.

FP types	in vein	in artery	on lymph	on artery wall	between vessels	vein wall	as Nodule	at bifurcation	others
percentage	39.06%	20.24%	12.00%	9.88%	4.71%	4.47%	3.29%	2.82%	3.53%

neighbors $N(v)$ slides down to pixel $g = \arg\min_{t \in N(v) \cup \{v\}} P(t)$. A pixel that cannot slide to any of its neighbors is called a *concentration*. All the pixels that slide down to the same concentration form a toboggan *cluster*. The basic idea is to treat the image as a landscape and let pixels slide within the landscape. Almost all the pixels along the vessel wall in the same intensity range as true PEs will merge into air regions and the true PE regions will stand out automatically. However, this basic tobogganing approach requires to scan the whole volume two times and cannot meet the real-time requirement for PE detection. To accelerate the tobogganing process, we developed a concentration-oriented tobogganing algorithm, achieving the same sensitivity as basic tobogganing but reducing computation time from originally 2 minutes to 27 seconds.

Readers can consult with [6] for a complete description of our concentration-oriented tobogganing algorithm. Briefly, the algorithm has two steps. It first searches for concentration c from the given pixel s and then expands from the found concentration c to extract the whole toboggan cluster C. The expansion includes a base step and an iterative step. In the base step, it includes the concentration c as the first pixel in the cluster and pushes all its neighbors with CT values between -50 HU and 100 HU into an active list A. In the iterative step, it selects pixel q with the minimal CT value from the active list A, if the selected pixel toboggans to an already clustered pixel, then conditionally pushes its neighbors to the active list A to ensure the uniqueness of the pixels in the active list A, otherwise, the selected pixel belongs to the cluster's external boundary B. The iterative step is repeated till the active list A is empty. This concentration oriented tobogganing algorithm is repeatedly applied on all those external boundary pixels, until a whole PE candidate has been extracted. This process is illustrated in Fig. 2.

The concentration oriented tobogganing algorithm is then iteratively applied on each of the external boundary pixels with CT value in the PE HU range. Any additional extracted toboggan cluster is merged into the previously extracted toboggan clusters, and any additional external boundary pixels are also merged. Once no external boundary pixel is left, all the toboggan clusters are extracted and merged, automatically forming a connected component – a PE candidate.

We validated the tobogganing algorithm through a clinical study of 177 cases (collected from multiple medical institutions) with 872 clots marked by expert chest radiologists. These cases were divided into two sets: training (45 cases with 156 clots) and test (132 cases with 716 clots). This algorithm successfully detected 90.38% (141/156) of the PE in the training set and 90.1%(645/716) of the PE in the test set. However, it also produces false positives—candidates that do not intersect with any PEs. On average, 47.5 and 40.3 false positives for each case were generated for the training set and the test set, respectively. Our false positive distribution is shown in Tab. 1.
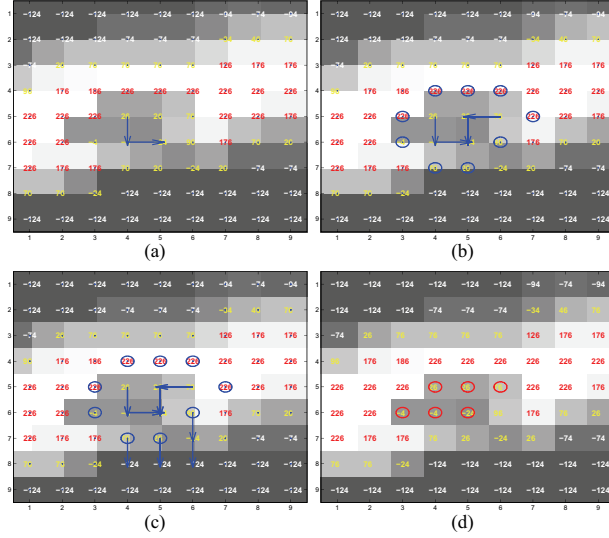
Fig. 2. Using the concentration oriented toboggan algorithm for detecting PE in a small vessel. During the scan in row by row, the first pixel which does merge into dark regions ($< $ -50 HU) is (4,5) and regarded as a PE pixel, from which our concentration oriented toboggan algorithm is applied to extract a whole PE candidate. (a) Step A of the algorithm: Finding the concentration. (b) Step B: Expanding from the concentration to cover a whole toboggan cluster and providing all the external boundary pixels (circled). (c) Repeatedly apply the algorithm on all those external boundary pixels with CT value between -50 HU and 100 HU to form a PE candidate (d).

2.2 Classification

A system that "cries wolf" too often will be rejected out of hand by radiologists. Therefore, to reduce the false positives, an effective multiple instance learning algorithm was designed [7] based on the 1-norm support vector machine (SVM). The 1-norm SVM constructs a linear decision function $f(\mathbf{x}) = \mathbf{w}^T \mathbf{x} + b$ by minimizing the regularized empirical error $\gamma \|\mathbf{w}\|_1 + \sum_{i \in S_+} \xi_i + \sum_{i \in S_-} \xi_i$ where γ is a tuning parameter and $\xi = \max\{0, 1 - y(\mathbf{w}^T \mathbf{x} + b)\}$ and $y = 1$ for candidates in S_+ (PEs), $y = -1$ for candidates in S_- (non-PEs). The multiple-instance learning approach makes use of the fact that some candidates can belong to the same PE, so that a PE is missed only if all the candidates belonging to it are classified as non-PE. Let the i^{th} PE contain ℓ_i candidates, represented as a set of feature vectors $\{\mathbf{x}_{ij}\}_{j=1}^{\ell_i}$. Let $S_i, i = 1, \cdots, m$ be the index set of all candidates that belong to the i^{th} PE assuming there exist totally m PEs. The goal of our classification algorithm is to determine a decision boundary that separates, with high accuracy, at least one candidate from each set S_i, $i = 1, \cdots, m$ on one side and as many as possible negative detections on the other side. It implies that the candidate in S_i with the smallest loss ξ should be correctly classified for each i. Mathematically, it means we

merely need to minimize the smallest error occurred on each set S_i and hence the objective function is to $\min_{\mathbf{w},\xi} \gamma||\mathbf{w}||_1 + \sum_{i=1}^{m} \min\{\xi_{ij}, \ j \in S_i\} + \sum_{i\in S_-} \xi_i$ subject to constraints (1) $\mathbf{w}^T\mathbf{x}_i + b \geq 1 - \xi_i, \ i \in U_{i=1,\cdots,m}S_i$, (2) $\mathbf{w}^T\mathbf{x}_i + b \leq -1 + \xi_i, \ i \in S_-$, and (3) $\xi_i \geq 0, \ i = 1, \cdots, \ell$.

We resort to aggregation of multiple classifiers trained over various sample patient sets to reduce the variance of the constructed classifier. We carry out T trials, and in each trial, two third of the training cases are randomly sampled and used for training. Linear decision functions are constructed in different trials. The final classifier is obtained by averaging the weight vectors of these decision functions. Features with very small weights in the final model will be removed to reduce noise level. This classification algorithm can reduce much more false positives naturally due to its design. However, the power of the classification algorithm relies on the discriminativeness of the features. We therefore develop three sets of features described in the following sections.

3   Local candidate characteristic features

As shown in Tab. 1, there are many false positives generated from vessels which are often due to noise or poorly mixed contrast material. One way to distinguish these false positives from true PE candidates is to design features based on the spatial distributions of intensity values computed within the candidate cluster. Therefore, in addition to the basic intensity distribution features such as mean, variance, skewness and kurtosis, we compute a number of features according to gray level co-occurrence matrices.

The co-occurrence matrices are widely used in texture analysis to capture the spatial dependence of intensity values within an image using second-order statistics. For a given 3D volume I of $n_1 \times n_2 \times n_3$ and a number of intensity levels N, a co-occurrence matrix $\mathbf{C}$ is an $N \times N$ matrix and acts as an accumulator over the image I. Specifically, a co-occurrence matrix along a direction (d_x, d_y, d_z) is defined as follows:

$$C_d(i,j) = \sum_{x=1}^{n1}\sum_{y=1}^{n2}\sum_{z=1}^{n3}\begin{cases} 1, & \text{if} I(x,y,z) = i \ \& \ I(x+d_x, y+d_y, z+d_z) = j, \\ 0, & \text{otherwise}. \end{cases} \quad (1)$$

Haralick [8] proposed fourteen statistical features computed from the co-occurrence matrices. Based on our experiments, the following four features have the most discriminative power in capturing the spatial dependence of intensity values within a PE candidate:

- **Entropy**, $\mathcal{E} = -\sum_i^N \sum_j^N C_d(i,j) \log C_d(i,j)$, measures the randomness of a gray-level distribution and is expected to be high if the gray levels are distributed randomly within a PE candidate.
- **Energy**, $\mathcal{G} = -\sum_i^N \sum_j^N C_d^2(i,j)$, measures the number of repeated pairs of gray levels, and is expected to be high if co-occurrence concentrates on certain pairs (i,j).
- **Contrast**, $\mathcal{C} = -\sum_i^N \sum_j^N (i-j)^2 C_d(i,j)$, measures the amount of local variations within a PE candidate and is expected to be low if the gray levels of each pixel pair are similar.

– **Homogeneity**, $\mathcal{O} = -\sum_i^N \sum_j^N C_d(i,j)/(1 + |i - j|)$, measures the smoothness of a PE candidate and is expected to be large if the gray levels of all pixel pairs are similar.

Based on these definitions, we expect that true PE candidates have lower entropy, higher energy, lower contrast and larger homogeneity, so as to be distinguished from those false positives due to noise or poorly mixed contrast material. These features serve as the basic set of image features. We validate these features by evaluating the PE detection performance using a recently-developed classification algorithm [7]. Figure 3 shows test performance of 60% sensitivity at 5 false positives per volume with these features, and hence this set of features is not yet sufficient to meet our goal.
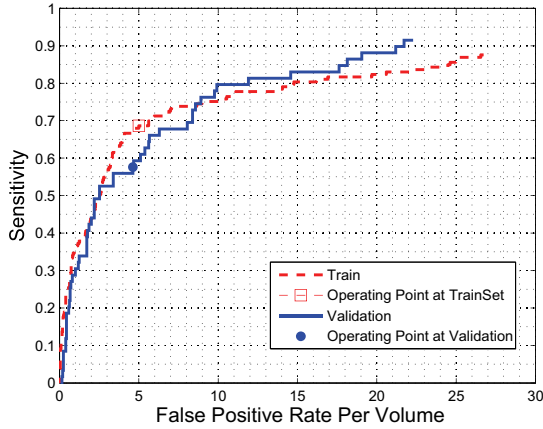

Fig. 3. ROC plot using local candidate features.

4 Vesselness characteristic features

Again as shown in Tab. 1, we have 12% of false positives generated from the connective (lymph) tissues located outside of the arteries. To eliminate this type of false positives, we design a set of features to determine if a candidate is located inside of a vessel based on the multi-scale Hessian vessel-likelihoods. Ideally, if a voxel is within a vessel, the Hessian matrix **H** formed from the second derivatives of intensity in the neighborhood of this voxel will have two negative eigenvalues and a third one is close to zero. The eigenvectors corresponding to the first two eigenvalues lie in a plane orthogonal to the central axis of the vessel, and the eigenvector corresponding to the third eigenvalue is in a line with this central axis. Based on the eigenvalues, vessel-likelihoods can be

computed based on the Frangi [9], Sato [10] and Lorenz [11] methods. However, this idea works only if the vessel does not have PEs. So, we precede this procedure by filling the detected regions (candidates) with high intensity values, to simulate the inside of a PE-ridden artery. To accommodate the different vessel radii, three vessel-likelihoods are computed at five different scales and the maximal response for each vessel-likelihood is taken across the five scales. In summary, our Hessian-based vesselness features are computed as follows:

1. Fill the detected candidates with high intensity values
2. Compute vessel-likelihood features for each voxel $\mathbf{x}$ in the candidate by:
   (a) Constructing the Hessian matrix $\mathbf{H}(\mathbf{x}, s)$ based on the second derivatives in scale s at voxel $\mathbf{x}$
   (b) Decomposing $\mathbf{H}$ into eigenvalues λ_1, λ_2, and λ_3 with $|\lambda_1| \leq |\lambda_2| \leq |\lambda_3|$ in scale s
   (c) Computing three vessel-likelihoods in scale s at voxel $\mathbf{x}$:

$$\mathcal{F}(\mathbf{x}, s) = \left[1 - e^{-\frac{1}{2}(\frac{\lambda_2}{\alpha\lambda_3})^2}\right] e^{-\frac{1}{2|\lambda_2\lambda_3|}(\frac{\lambda_1}{\beta})^2} \left[1 - e^{-\frac{\lambda_1^2 + \lambda_2^2 + \lambda_3^2}{2\gamma^2}}\right] \quad (2)$$

$$\mathcal{S}(\mathbf{x}, s) = \begin{cases} \sigma^2|\lambda_3| \left(\frac{\lambda_2}{\lambda_3}\right)^\xi \left(1 + \frac{\lambda_1}{|\lambda_2|}\right)^\tau, & \lambda_3 \leq \lambda_2 \leq \lambda_1 \leq 0 \\ \sigma^2|\lambda_3| \left(\frac{\lambda_2}{\lambda_3}\right)^\xi \left(1 - \rho\frac{\lambda_1}{|\lambda_2|}\right)^\tau, & \lambda_3 \leq \lambda_2 \leq 0 \leq \lambda_1 \leq \frac{|\lambda_2|}{\rho} \end{cases} \quad (3)$$

$$\mathcal{L}(\mathbf{x}, s) = \sigma^\eta \left|\frac{\lambda_2 + \lambda_3}{2}\right| \quad (4)$$

   (d) Maximizing vessel-likelihoods at voxel $\mathbf{x}$ across scale s:

$$\mathcal{F}(\mathbf{x}) = \max_s \mathcal{F}(\mathbf{x}, s), \quad \mathcal{S}(\mathbf{x}) = \max_s \mathcal{S}(\mathbf{x}, s), \quad \mathcal{L}(\mathbf{x}) = \max_s \mathcal{L}(\mathbf{x}, s) \quad (5)$$

3. Compute statistical features (*i.e.*, maximum, minimum, mean, median and standard deviation) based on the computed three vessel-likelihoods for all the voxels in the candidate.

True PE candidates have much higher vessel-likelihoods than those generated outside of vessels (*e.g.*, lymph false positives). We plot the training and test detection performance of our classification algorithm in Fig. 4 with these vesselness features in addition to the basic features. It clearly shows the improvement of around 20% increase on sensitivity at 5 false positives when vesselness features are included.

5   Vein/Artery characteristic features

We have nearly 40% false positives generated in veins. In CTPA images, we found that the veins have different intensity distribution and texture patterns from arteries. Furthermore, those false positives within arteries were generated mainly due to poorly mixed contrast material, demonstrating distinct intensity and texture patterns. Therefore, we design a set of features to capture the intensity and texture patterns of veins and arteries to eliminate false detections from veins. To do so, we extract a segment of vessel from a given candidate and then compute a set of features based on the extracted vessel segment. While taking advantage of the benefits of vessel segmentation, our approach is to only segment the relevant areas around the candidate, thereby avoiding an expensive complete segmentation of the complete pulmonary vessel tree.
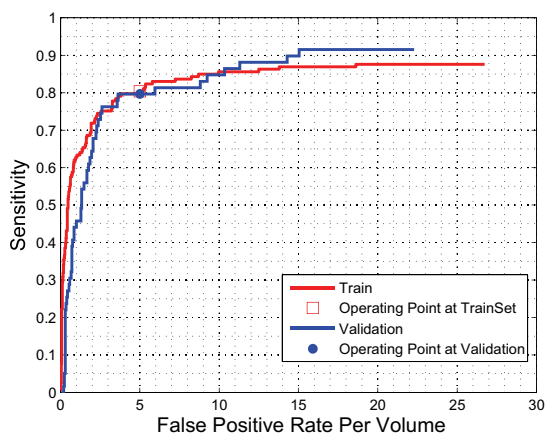
Fig. 4. ROC plot using vesselness features in addition to local candidate features.

5.1 Vessel segment extraction

Each PE candidate generated by our CG is a connected component, consisting of a set of voxels. In this step, we first extract the candidate boundary surface and then obtain the vessel segment with a modified Dijkstra's algorithm. The PE candidate boundary surface is extracted by checking the neighbors of each voxel of the candidate. If any of its neighbors does not belong to the candidate, then the voxel is on the boundary surface and is recorded. The vessel segment containing the candidate component is extracted by a graph-searching based iterative region-growing process by taking the extracted boundary points of the PE candidate as initial seeds. This process creates a minimum cumulative cost path map. The cost is set to the Euclidean distance from the boundary of the candidate component. The growth occurs from the voxel with the lowest cumulative distance cost on the boundary of the growing region. In the growing process, only those voxels with intensity value above a threshold (100 HU) are considered, so that the growth is along the contrast enhanced vessels. As a result, the extracted region grows in layers (equal distance to the boundary of the candidate component) and stops if the vessel length is larger than a given value (30 mm in our case).

5.2 Vessel feature computation

Once the vessel segment is extracted, similarly we can compute the basic intensity distribution features (mean, variance, skewness and kurtosis) and texture features based on gray level co-occurrence matrices from the extracted vessel segment as described in Section 3. In addition, we compute wavelet-based texture features. In the case of volumetric images, the discrete wavelet decomposition is obtained by applying a pair of
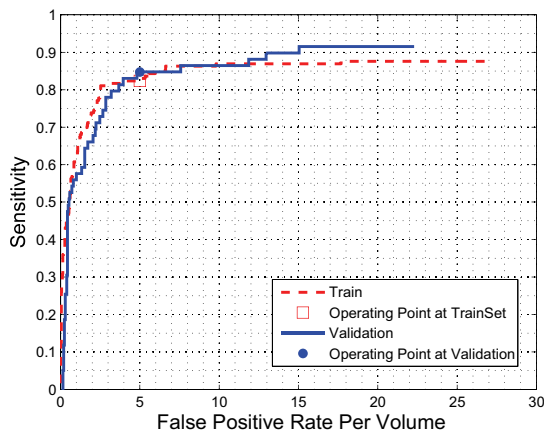
Fig. 5. ROC plot using the entire set of features.

wavelet filters: a lowpass filter h and a highpass filter g, along the x, y, and z direction of the volumetric image. In a one-level decomposition of the 3D discrete wavelet transform (DWT), the particular pair of filters h and g corresponds to a particular type of wavelet used and $\downarrow_x 2$ stands for downsampling by 2 along x direction, generating nine subvolumes from the input volumetric image. For example, subvolume LLL corresponds to the lowest frequencies and subvolume LHH gives the low frequency along x direction and high frequencies along y and z directions. A multi-level DWT decomposition is obtained by repeating the same procedure to the subvolume LLL until the desired level is reached. In our application, we used Daubechies 8-tap filters and a 2-level decomposition of 3D DWT to compute the wavelet subvolumes. Then for each subvolume, the energy feature is calculated from its wavelet coefficients,

$$\mathcal{W} = \frac{1}{N_x N_y N_z} \sum_x \sum_y \sum_z |w(x, y, z)| \tag{6}$$

where N_x, N_y, and N_z are dimensions of the subvolume and $w(...)$ is a wavelet coefficient within the subvolume. For each candidate, we have nine energy features computed from the nine subvolumes.

We validate the third set of features by incrementally including them into our training process. Hence the final set of features include co-occurrence based features in Section 3 and features described in Sections 4 and 5. The final system achieves 85% test sensitivity at 5 false positives per volume.

6   Conclusion

We have developed three sets of new local characteristic features for eliminating false positives in automated PE detection. The developed features have been fully integrated

into our PE CAD system and validated on over 130 cases, showing that the PE CAD system performance has been incrementally improved by more than 20% at the ROC operating point with our second and third set of the features, respectively, in our multiple instance learning framework. Our approach has a set of distinguished features, requiring no vessel segmentation, reporting any detection incrementally in real time, and detecting both acute and chronic pulmonary emboli, achieving a sensitivity of 85% at 5 false positives, resulting in a clinically usable PE CAD system, which has been deployed in many clinical sites around the world.

References

1. Y. Masutani, H. MacMahon, and K. Doi, "Computerized detection of pulmonary embolism in spiral CT angiography based on volumetric image analysis," *IEEE Transactions on Medical Imaging*, vol. 21, no. 12, pp. 1517–1523, 2002.

2. M. Quist, H. Bouma, C. van Kuijk, O. van Delden, and F. Gerritsen, "Computer aided detection of pulmonary embolism on multi-detector CT," in *RSNA*, Chicago, USA, Nov 2004.

3. C. Zhou, H. Chan, S. Patel, P. Cascade, B. Sahiner, L. Hadjiiski, and E. Kazerooni, "Preliminary investigation of computer-aided detection of pulmonary embolism in three-dimensional computed tomography pulmonary angiography images," *Academic Radiology*, vol. 12, no. 6, pp. 782–792, 2005.

4. S. Park, C. Bajaj, G. Gladish, and D. Cody, "Automatic pulmonary embolos detection and visualization," poster at http://www.ices.utexas.edu/~smpark, 2004.

5. J. Liang, M. Wolf, and M. Salganicoff, "A fast toboggan-based method for automatic detection and segmentation of pulmonary embolism in ct angiography," in *The Eighth International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI 2005)*.

6. J. Liang and J. Bi, "Computer aided detection of pulmonary embolism with tobogganing and multiple instance classification in CT pulmonary angiography," in *Proceedings of the 20th International Conference on Information Processing in Medical Imaging, 20th International Conference, IPMI 2007*, Kerkrade, The Netherlands, 2007, pp. 630–641.

7. J. Bi and J. Liang, "Multiple instance learning of pulmonary embolism detection with geodesic distance along vascular structure," in *Proceedings of IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'07)*, Minneapolis, Minnesota, USA, 2007.

8. R. M. Haralick, K. Shanmugam, and I. Dinstein, "Textural features for image classification," *IEEE Transactions on Systems, Man and Cybernetics*, vol. 3, no. 6, pp. 610–621, 1973.

9. A. F. Frangi, W. J. Niessen, K. L. Vincken, and M. A. Viergever, "Multiscale vessel enhancement filtering," in *Medical Image Computing and Computer-Assisted Intervention*, 1998, pp. 130–137.

10. Y. Sato, S. Nakajima, H. Atsumi, T. Koller, G. Gerig, S. Yoshida, and R. Kikinis, "3d multiscale line filter for segmentation and visualization of curvilinear structures in medical images," in *CVRMed-MRCAS'97: Proceedings of the First Joint Conference on Computer Vision, Virtual Reality and Robotics in Medicine and Medial Robotics and Computer-Assisted Surgery*.   London, UK: Springer-Verlag, 1997, pp. 213–222.

11. C. Lorenz, I.-C. Carlsen, T. M. Buzug, C. Fassnacht, and J. Weese, "Multi-scale line segmentation with automatic estimation of width, contrast and tangential direction in 2d and 3d medical images," in *CVRMed-MRCAS'97: Proceedings of the First Joint Conference on Computer Vision, Virtual Reality and Robotics in Medicine and Medial Robotics and Computer-Assisted Surgery*.   London, UK: Springer-Verlag, 1997, pp. 233–242.

A Multi-Scale Topo-Morphologic Approach for Separating Arteries and Veins in Pulmonary CT Images

Punam Kumar Saha[1,2], Milan Sonka[1], Zhiyun Gao[1], Eric Hoffman[2]

[1]Depts. of Electrical & Computer Engineering, and [2]Radiology
The University of Iowa, Iowa City, IA 52242

Abstract. In this paper, we present a new method for separating two iso-intensity objects attached to each other at different locations at various scales and apply the method to separating arteries and veins in pulmonary CT images. The method starts with two sets of seeds — one for arteries and another for veins. Initialized with seeds, arteries and veins grow iteratively while maintaining their spatial separation and eventually forming two disjoint objects at convergence. The method combines fuzzy distance transform, a morphologic feature, with a topologic connectivity property to iteratively separate finer and finer details starting at a large scale and progressing towards smaller scales. The method has been validated in mathematically generated tubular objects with different levels of fuzziness, scale and noise. Also, it has been successfully applied to *in vivo* clinical CT pulmonary data for separating arteries and veins. Results have demonstrated the method's ability to resolving multi-scale adherence of two iso-intensity objects even when there are no sign of intensity variation at conjoining locations.

1 Introduction

Image Segmentation [8] – a method of producing a spatio-temporal object definition in an image – has remained a salient task in most medical imaging applications. Different segmentation methods [3] focus on different image features and properties and often, the design of an effective segmentation algorithm in a limited SNR and resolution regime is highly challenging and application dependent. Here, we have selected a specific segmentation task of separating arteries and veins (*A/V*) via *in vivo* pulmonary CT imaging with no blood pool enhancing contrast. Although, such a method of separating A/V is very useful, the challenges are multi-folded including — (1) A/V are indistinguishable by their intensity values in non-contrast pulmonary CT images, (2) often, there is no trace of intensity variation at locations of adherence between A/V, (3) complex and tight coupling between A/V with arbitrary and multi-scale geometry, especially, at branching locations and (4) limited SNR and resolution of *in vivo* imaging. Patient-specific structural abnormalities of vascular trees further complicate the task. Several works have been reported in literature toward solving the problems of separating arteries and veins using improvised image acquisition techniques; a thorough discussion on difficulties of such approaches, especially, for smaller vessels has been presented by Bemmel *et al* [10]. As far as our knowledge goes, only a few post-processing methods have been published on separating arteries and veins [2,10]. The previous methods have been applied to MR data only and do not use morphological scale information.

These methods primarily rely on intensity variation or edge information at the adherence locations between A/V and may not work for *in vivo* CT images where no intensity variations are visible at locations of adherence between A/V. Recently, two approaches for artery/vein classification from pulmonary CT images appeared [1,11].

Here, we introduce the use of morphologic scale features which enables separating two structures even when there is no edge feature at locations of adherence between the two which often happens between A/V in CT data. Specifically, we develop a new topo-morphologic method that iteratively separates A/V starting from two sets of seeds at high scale and then progressively solves the problem at lower scales. The method combines fuzzy topological approaches [4,9,5] with fuzzy distance transform [6] — a morphological feature — without requiring a parameter. Separated A/V trees may significantly contribute to our understanding of pulmonary structure and function and has immediate clinically applicable applications, e.g., for assessment of pulmonary emboli. The knowledge of separated A/V may significantly boost performance of airway segmentation methods.
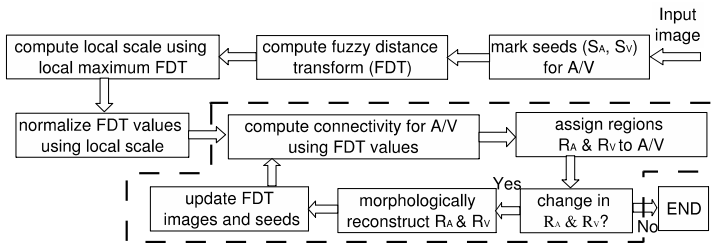

Fig. 1. A schematic description of the overall method.

2 Theory and Methods

The overall work flow diagram of the method is presented in Figure 1 that separates two iso-intensity objects using multi-scale morphological features. As input, it receives fuzzy segmentation of the assembly of two iso-intensity objects and two sets of seeds — one for each object — and it outputs separated objects. Although the method immediately extends to multiple objects, here we formulate a solution for two objects only and we will refer one object as the rival of the other. Let us consider an image consisting of two iso-intensity fuzzy objects with significant noise and overlapping as shown in Figure 2(a); a few cross sectional images are shown in Figure 2(b). The two cylinders with gradually reducing radii are running in parallel. The diameter of one cylinder is significantly larger than other; a sinusoidal swing is added to both in the xy-coordinate direction so that central lines of both cylinders lie on an xy-plane, say
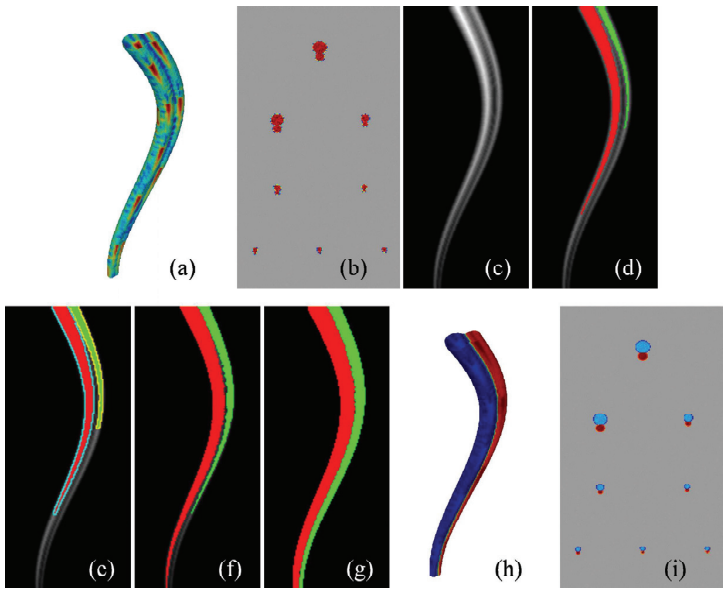
Fig. 2. (a) 3D rendition of a phantom image. (b) A few cross-sectional images. (c) FDT image on the central xy-plane. (d) Results of separation of the two cylinders after the first iteration using FDT-based connectivity. (e) Morphologic reconstruction based on the results of (d). (f,g) Same as (d) after second (f) and terminal (g) iterations. (h,i) 3D rendition and cross-sectional images of the result.

the *central xy-plane*. Usefulness of the method may be better understood in three and
higher dimensions. On the other hand, it is always easier to illustrate a method in 2D.
Therefore, we illustrate results of different steps of the method on the central xy-plane.
One seed is manually placed through a graphical interface on the approximate center
of each cylinder at the top-most cross-section. As shown in Figure 1, the process is it-
erative and begins with marking the seeds followed by fuzzy distance transform (FDT)
computation [6] of the assembly of two cylinders prior to iterations; see Figure 2(c)
for the results of FDT computation. It is visually notable from the FDT image that we
cannot find a single FDT threshold to separate the two cylinders from their head to tail.
Therefore, no straightforward morphological opening operator can solve the problem.
On the other hand, over a small region, the two cylinders seem separable using their
FDT values. In other words, the problem demands regional selection of an optimum
opening structure to separate the two cylinders.

The above observation motivated us using *local scale* to normalize FDT values that
reduces the effect of spatial scale variations. Local scale at a point p is defined as the
FDT value of the locally-deepest point (a point with locally maximum FDT value) that
is nearest to p. With the normalized FDT map, the method adopts an iterative strategy
that first separates the cylinders over large-scale regions using FDT-based relative con-
nectivity [5] where a point is grabbed by an object if its connectivity to the point is
strictly greater than that of its rival. A separator is built between the two objects using
a morphological reconstruction method that simultaneously and radially dilates each
currently segmented region until blocked by its rival (maximum radius of the dilating
structure is determined by FDT values). Figures 2d & e show the results of initial sep-
aration and morphologic reconstruction of two cylinders after first iteration. In the next
iteration, the FDT-connectivity paths of one object are not allowed to enter into the
region assigned to its rival. This strategy facilitate resolving fusion at smaller scale re-
gions (Figure 2f) and this iterative process is continued as long as there is any change.
For this phantom image, the method stopped after 12 iterations (see Figure 2g–i for
final separation).

3   Results and Discussion

Effectiveness of the method has been examined by applying it to mathematically gen-
erated phantoms and to clinical CT pulmonary images. Five mathematical phantoms
were computer-generated as tubular objects running across the slice direction with dif-
ferent levels of fuzziness, overlap, scale and noise. Initially, the phantom images were
generated at high resolution and then downsampled using 3×3×3, 4×4×4 and 5×5×5 win-
dows. Each downsampled image was degraded with additive noise at SNR of 12. Using
a graphical interface, exactly one seed point was manually selected for each tubular
object in a phantom near its center on the slice at the lower-most level. Although there
is no thoretical restriction on positioning the seed points, the seeds are were positioned
close to the vessel centerlines. Phantoms and results are depicted in Figure 3 at $4{\times}4{\times}4$
and $5{\times}5{\times}5$ downsampling only, as the method has always successfully separated the
two cylinders at $3 \times 3 \times 3$ downsampling. At extremely low resolution and high noise,
the morphological information may be entirely lost, leading to a failure of separation

in Figure 3t,v. Specifying another pair of seeds would be help achieving correct separation. Cross sectional images of the first phantom are illustrated at both resolutions to depict the 3D phantom image quality used and the complexity of separating the two objects. The smallest radius of the phantom of Figure 3a is 1.87 pixels at the downsampling by $4 \times 4 \times 4$ pixels; the largest radius in the same example at same downsampling resolution is 18.75 pixels. Except for the examples of Figures 3a and f, the radii of the two cylinders differed significantly.
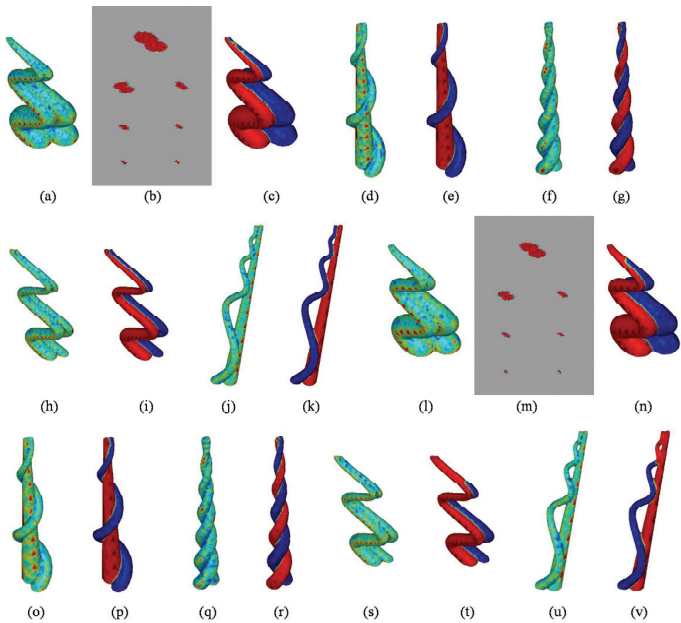
Fig. 3. (a-k) Results of our method's application to several 3D phantom images downsampled by $4 \times 4 \times 4$. (a,b) 3D rendition and cross-sectional images of one phantom image. (c) Separated cylinders. (d-k) Results for other four phantoms. (l-v) Results of application of the method at downsampling by $5 \times 5 \times 5$.

The method has been applied to separate arteries and veins in pulmonary CT images (see Figure 4). Thoracic region of two females of age 22 Y and 27 Y were scanned using a Siemens Sensation 64 MDCT scanner at 120 kVp and 200 mAs. One subject was scanned in feet first supine while the other was scanned in head first supine position.

Images were acquired at 0.75 mm slice thickness and were reconstructed with 0.5 mm slice-thickness and 0.55×0.55 mm^2 in-plane resolution. Entire vascular tree of both lungs was segmented in each CT image data using the method of Shikata *et al.* [7] that uses tree-connectivity on the CT intensity image enhanced by the output of a tubular enhancement filter. Resulting A/V separation outcomes are shown in Figure 4. Seed points were manually selected using a 2D slice-display graphical interface followed by the application of the morphological separator algorithm. The subsequent automated process requires 2 to 3 mins. to complete the A/V separation for the right or left lung. Finally, the separated A/V were displayed with color coding (red for veins and blue for arteries; see Figures 4b,c) through a 3D surface rendition software tool developed in-house using VTK-based 3D visualization filter classes.
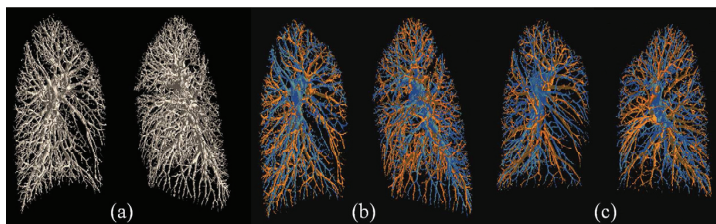
Fig. 4. Results of application to pulmonary CT images. (a) 3D surface renditions of vessel trees in left and right lungs. (b) Color-coded 3D rendition of separated artery/vein trees. (c) Same as (b) but for another dataset. Blue – arteries (carrying non-oxygenated blood), red – veins.

The method has demonstrated feasibility of separating two iso-intensity structures with multi-scale adherences even when there is no sign of intensity-based separation at joining locations. The method seeks morphological identities of each object at a specific scale and separates them without using any gradient- or edge-like features. Introduction of the ideas of morphological reconstruction and separator allows the method to seal the joining border at current scale and then seek morphological features identifying different objects at finer scales. The mathematical phantom demonstrates the ability of the method when the geometry of coupling of two objects are known and its performance with just one seed for each object is encouraging. For pulmonary CT images, the geometry of coupling between arteries and veins are far more challenging and quite unknown. However, the method has shown acceptable performance with a reasonable number of seeds. Approximately 25–35 seeds were manually selected on each of the A/V subtrees. It may be pointed out that the seeds were selected using a 2D graphical tool. Often multiple seeds were placed within the same locality of an object and therefore, not all of the seeds contributed to true A/V separation. We have found that seed selection is an important task in the entire process and the current seed-selection tool is

far from optimal. An effective seed selection process must be performed through a 3D
graphical interface coupled with the morphological separator providing a transparent
interactive environment to the user. Currently, we are developing such a seed-selection
graphical system and we believe that, once it is developed, the number of seeds required
will significantly decrease.

References

1. T. Buelow, R. Wiemker, T. Blaffert, C. Lorenz, and S. Renisch. Automatic extraction of the
   pulmonary artery tree from multi-slice ct data. In *Proceedings of SPIE: Medical Imaging*,
   volume 5746, pages 730–740, San Diego, CA, 2005. 2
2. T. Lei, J. K. Udupa, P. K. Saha, and D. Odhner. Artery-vein separation via MRA - an image
   processing approach. *IEEE Transactions on Medical Imaging*, 20:689–703, 2001. 1
3. D. Pham, X. Chenyiang, and J. L. Prince. A survey of current methods in medical image
   segmentation. *Annual Review of Biomedical Engineering*, 2:315–337, 2000. 1
4. A. Rosenfeld. Fuzzy digital topology. *Information and Control*, 40:76–87, 1979. 2
5. P. K. Saha and J. K. Udupa. Iterative relative fuzzy connectedness and object definition:
   theory, algorithms, and applications in image segmentation. In *Proceedings of IEEE Work-
   shop on Mathematical Methods in Biomedical Image Analysis*, Hilton Head, South Carolina,
   2000. 2, 4
6. P. K. Saha, F. W. Wehrli, and B. R. Gomberg. Fuzzy distance transform - theory, algorithms,
   and applications. *Computer Vision and Image Understanding*, 86:171–190, 2002. 2, 4
7. H. Shikata, E. Hoffman, and M. Sonka. Automated segmentation of pulmonary vascular
   tree from 3D CT images. In *Proceedings of SPIE: Medical Imaging*, volume 5369, pages
   107–116, San Diego, CA, 2004. 6
8. M. Sonka, V. Hlavac, and R. Boyle. *Image Processing, Analysis, and Machine Vision*. PWS
   Publishing, Pacific Grove, CA, 1999. 1
9. J. K. Udupa and S. Samarasekera. Fuzzy connectedness and object definition: theory, algo-
   rithms, and applications in image segmentation. *Graphical Models and Image Processing*,
   58:246–261, 1996. 2
10. C. M. van Bemmel, L. J. Spreeuwers, M. A. Viergever, and W. J. Niessen. Level-set-based
    arteryvein separation in blood pool agent CE-MR angiograms. *IEEE Transactions on Medi-
    cal Imaging*, 22:1224–1234, 2003. 1
11. T. Yonekura, M. Matsuhiro, S. Saita, M. Kubo, Y. Kawata, N. Niki, H. Nishitani, H. Ohmatsu,
    R. Kakinuma, and N. Moriyama. Classification algorithm of pulmonary vein and artery based
    on multi-slice ct image. In *Proceedings of SPIE: Medical Imaging*, volume 6514, 65142E,
    San Diego, CA, 2007. 2

Software Demonstrations

Pulmonary Analysis Software Suite 9.0: Integrating Quantitative Measures of Function with Structural Analyses

Junfeng Guo[1], Matthew K Fuld[1,2], Sara K Alford[1,2], Joseph M. Reinhardt[2], Eric A Hoffman[1,2]

[1] University of Iowa, Department of Radiology, Iowa City, IA, United States
[2] University of Iowa, Department of Biomedical Engineering, Iowa City, IA, United States

{junfeng-guo, matthew-fuld, sara-alford, joe-reinhardt, eric-hoffman}@uiowa.edu

Abstract. With the recent advances in multidetector-row CT (MDCT) to dynamically assess the lung with a z axis coverage now including up to 70% of the whole lung, we have integrated quantitative measures of regional pulmonary perfusion and ventilation into the Pulmonary Analysis Software Suite (PASS), allowing for detailed assessments of structure-to-function relationships and are using this integrated system to identify early smokers at risk of chronic obstructive pulmonary disease (COPD). Previously we developed methods to segment lungs, lobes, airways, and blood vessels and extract histogram and texture-based measures of the lung parenchyma. These segmentation, analysis and display tasks were integrated into a comprehensive software package: PASS.

Keywords: MDCT, Medical Image Analysis, Texture Analysis, Functional Imaging, Ventilation, Perfusion

1 Introduction

Quantitative assessment of lung structure along with indices of parenchymal pathology are taking on increased roles in the detection and tracking of pulmonary disease. To date the focus has largely been on airway morphometry and indices of parenchymal destruction, and air trapping. The parenchymal analysis has, in large part, focused on the use of the density histogram within the lung field to identify voxels falling below a given density threshold to define volumes of emphysema-like lung or air trapping. Some work has shown that texture measures can provide more accurate detection and quantification of pathology not limited to enlargement of peripheral air spaces [1, 2]. To date, our quantitative tools for the assessment of the lung parenchyma have been integrated into a software package which we have dubbed the Pulmonary Analysis Software Suite or PASS, and PASS has been used in a number of large multi-center studies including the NIH sponsored National Emphysema Treatment Trial [3], the Lung Imaging Database Consortium [4], and as

clinical applications for PASS have expanded, this software has been commercialized and integrated with airway analysis and guidance software (Pulmonary Workstation Plus, VIDA Diagnostics, Coralville, IA). In addition to structural information, dynamic CT imaging is capable of delivering regional measures of ventilation and perfusion by following the accumulation of Xenon gas over multiple breaths [5, 6, 7] or the first pass kinetics of a sharp bolus of iodinated contrast agent as it passes through the lungs [8, 9]. To date, the functional measures have been limited largely to research studies because of the limited z-axis coverage during axial scanning protocols. With the recent advances in multidetector-row CT (MDCT) to dynamically assess the lung with a z axis coverage now including up to 70% of the whole lung, we have been motivated to integrate our quantitative measures of regional pulmonary perfusion and ventilation into PASS, allowing for detailed assessments of structure-to-function relationships. In an early application of integrated system we are identifying smokers at risk of chronic obstructive pulmonary disease (COPD) but with normal pulmonary function tests to assess correlations between early structural (histogram-based analysis) and functional (perfusion measures) changes. In this paper, we present an overview of the newly expanded PASS that now provides the link between structure and function for the inclusion of regional distribution of ventilation and perfusion as part of a comprehensive phenotype determination [10].

2 Software Components

2.1 PASS Core Functionality

2.1.1 Task Concept

The PASS framework is organized with a task-document-frame-view hierarchical structure. Depending on the task, one or more documents can be opened simultaneously, each document can open one or more frames, and each frame can show multiple different kinds of views. This design schema can be separated into two layers, the underneath layer is the document-frame-view structure which is a traditional multi-document interface (MDI). The upper layer is the task manager that uses a sequence of wizard dialogs to guide the user through predefined steps to finish a particular analysis process. The goal of this schema is to make the PASS software as flexible as a traditional MDI while also being simple to use as similar wizard driven applications. **Figure 1** shows an example of the task concept in which three tasks are open simultaneously. Each task is distinguished by a unique background color for both the tab bar at the top of the window and the floating dialog box. The floating dialog box is the interface to task manager that controls the task. Users can switch tasks by pressing the colored tabs at the top of the window, and the task manager will be switched automatically to the correct task.
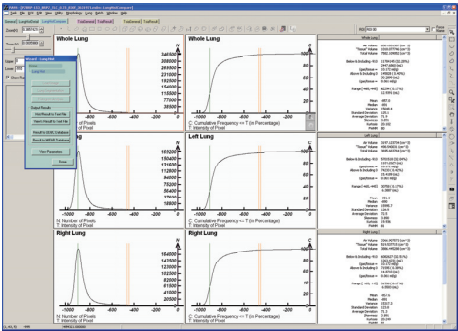
(a)
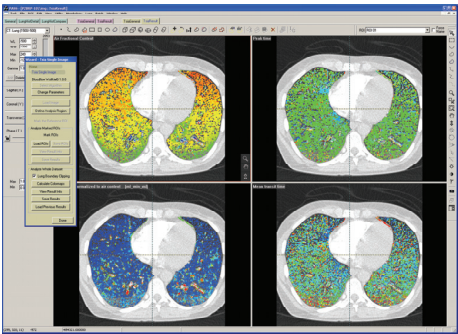
(b)

Fig. 1. Three tasks are opened in PASS. **(a)** The first task is the Histogram Analysis task (displayed in blue). This task opens one document which in turn opens three frames of result data. The third frame, which contains 9 views, is pictured above. **(b)** The third open task is the TSIA Single Image task (displayed in yellow). One document is opened which creates two frames. The second frame, which contains 4 views, is currently in display.

2.1.2 PASS Profile

In order to facilitate situational based behavior of the general class pieces of the PASS structure while under control of a wizard, a profile method has been designed. Each instance object of the above class will be assigned a corresponding type of profile that defines the behavior of that object. For example, depending on the task, document, frame, or location within a frame, an object of the transverse view class may display

different content. Similarly, parameters are saved for each algorithm, allowing them
to be retrieved based on a user's previous usage. Each task has a user-configurable
set of default profiles for the objects used in that task.
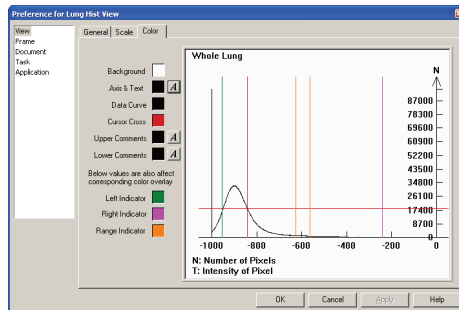
Fig. 2. PASS includes a Profile editing interface, pictured above are the settings for Lung
Histogram View. Through this interface the user can adjust the profile settings for current view,
frame, document, or task respectively. The options that can be adjusted vary depending on the
type of the object selected in the profile editor.

2.1.3 Separation of Core Functionality and Time Series Algorithms

The PASS software package is extensive, containing hundreds of GUI components
and algorithm modules. In order to manage the development in an efficient way, the
GUI framework is separated from the algorithm modules. For example, the Time
Series Image Analysis (TSIA) algorithm modules are independently developed
outside of PASS framework and are dynamically loaded into PASS at runtime using a
common memory module (CMM) interface.

2.1.4 DICOM Viewer

PASS enables the viewing of 2D, 3D, and 4D image data in a variety of ways.
Typically, the input to the PASS is a multi-dimensional MDCT data set of thorax as a
series of DICOM files. Inside a user specified directory, PASS finds all valid DICOM
files, automatically detects the scan plan, and organizes them into a single multi-
dimensional dataset. In some situations, a multi-dimensional scan is acquired in
multiple steps; with each step making up only a smaller portion of the thorax or of the
time series. The final dataset needs to be created by intelligently joining several of
these partial datasets. In the case where extra slices or overlap occurs between
datasets, a dialog prompts the user to interactively remove the extra slices. The newly
created dataset can be saved in the extended DICOM format for future use enabling
the user to skip the joining process in the future.

2.1.5 Extension to DICOM Standard

The standard DICOM file format is limited to 2D images. Additions have been made to the standard to enable it to handle higher-dimensional (3D & 4D) image data. A private section is inserted into the standard 2D DICOM tag structure that includes information regarding the higher-dimensionality of the dataset including number of phases, depth of volume, phase timing, etc. The pixel bitmap for the first slice is stored in the public area as usual; while the rest of the pixel bitmaps are stored under the private section. A special tag is generated by grouping the series-varying information from the original series into a single tag. This method enables the retrieval of 2D DICOM information from within the new higher-dimensional dataset and also retains the ability to recover the original DICOM files if necessary without data loss.

2.1.6 Segmentation

Lungs. One of the important first steps for quantitative analysis is segmentation of left and right lungs. The PASS lung segmentation processing consists of three main steps: an extraction step to identify the lungs; a separation step to separate the right and left lungs; and an optional smoothing step to smooth the lung boundaries. Complete details and validation information for the lung segmentation algorithm are given in [11].

Airways. Airway processing consists of three separate steps: (1) segmentation of the airway lumen, (2) analysis of the tree structure to define the branch centerlines and branching relationships, and (3) measurement of the airway lumen and airway walls. Complete algorithm details are given in [12, 13].

Vessels. The vessels are segmented using a variant of the method described by Shikata et al. [14]. Shikata et al. use a two-step approach: a line-filtering of the raw data based on an eigen-analysis of the Hessian matrix calculated at multiple scales, followed by an optional vessel tracking approach to extract the small vessel segments which are missed in the first step. For most purposes the second step is not necessary and is undesirable because of the increased computational cost. We threshold the line-filtered result to retain only voxels with line-filtered value less than -0.08, and then subsequently use size filtering to remove all segments less than 75 mm^3.

Currently vessel segmentation in used in PASS as a step within the lobe segmentation process (see below). Additionally, vessel segmentation allows for 3D visualizations of the vascular tree. There is ongoing work to separate the arteries and veins for further quantification of the vessel tree.

Lobes. Identification of the lobar fissures can help decompose the 3D lung into its major structural components. These components can be used as a basis for reporting measurements, and for inter-subject comparisons. The lobar segmentation is guided by the airway tree and vessel trees. After vessel segmentation, a distance map is computed based on the segmented vasculature. This distance map is analyzed using a

watershed transform. The watershed simulation proceeds and the basins are merged using markers automatically generated from the anatomically labeled airway tree. After the watershed analysis is completed, we obtain an initial, approximate segmentation of the fissures. The initial segmentation is refined using a 3D optimal surface segmentation. Full details on the lobar segmentation are given in [15].

Free Hand Drawing. For certain applications such as lung nodule analysis [4], there has been the desire to have a tool for manually identifying structure boarders so that automatically defined edge locations can be compared with radiologist defined gold standards. The free hand drawing tool includes a live wire feature that aids the user in drawing boarders. The free hand drawing can also be modified by shifting individual points along a trace, or multiple points around a boarder can be identified and a cubic spline fit will complete the boarder.

2.1.7 ROI Dictionary / Indexing

Lung segmentation assigns every pixel a region property or ROI label. Typically a mask image is used to record these labels. However, a major limitation of this approach is that the meanings of ROI labels are defined locally. Thus the mask bitmap will be unexplainable without first knowing the specific type of the mask image and having detailed documentation regarding its labeling schema. It is impossible to make global definitions for ROI labels since each pixel in the mask image has limited size (usually 8-bits), but the meaning it should be able to describe is infinite. In order to solve the problem between the limited label and infinite meaning, we proposed a ROI Dictionary/Indexing method. In this method, we build a global ROI dictionary to record all global ROI names and their relationships. For each mask image, a file called ROI Indexing is created which maps the local ROI labels to the global ROI names. With the help of Indexing file, a previous locally defined ROI image becomes a global resource, allowing the use of a mask image in a simple and uniform way.

2.2 Structural Analysis

2.2.1 Histogram Analysis

The gray level histogram has been widely used as a tool for detection of early parenchymal disease. PASS calculates the histogram and cumulative histogram for all pixels in the lung parenchyma. Further subdivisions of the histogram can be achieved using the regions defined from segmentation. Measurements can be reported for any of the defined ROIs, including separating the lung on a lobar basis. Several statistical parameters, such as gray level mean, median, standard deviation, average deviation, skewness, kurtosis, and histogram full width at half maximum (FWHM) are calculated. Total volume, air volume, and tissue volume are computed for each region. The slope and intercept value of the knee line and ankle line of the cumulative histogram curve have been identified as key parameters that may distinguish normal lungs from lung showing signs of early parenchymal disease.
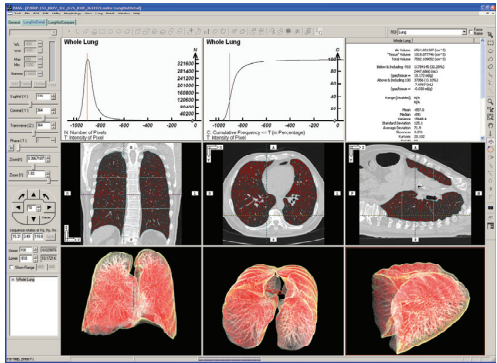
Fig. 3. Histogram analysis in PASS: Parenchymal histogram (upper row left), cumulative histogram (upper row middle), histogram results (upper row right), coronal – transverse – sagittal views (middle row), and 3D renderings (bottom row). Low attenuation areas in red.

2.2.2 Hole Analysis

The onset of emphysema is characterized by the development of poorly ventilated regions within the lungs ("holes"). We try to assess the distribution of holes within the lungs to estimate the severity of the disease process. The distribution of the holes in the normal lung is due only to cross sectional sampling of air within the airway tree. Therefore, we expect the probability distribution of hole cross-sectional area to change as new holes are created and as smaller holes join together to create larger ones during disease progression. Thus the hole distribution curve provides information that can be used to assess the disease severity.

2.2.3 AMFM Texture Analysis

The adaptive multiple feature method (AMFM) was introduced as a texture-based method which could take into account co-existing pathologies and provide a simultaneous classification of multiple simultaneous disease processes. Two kinds of tasks are made possible in the package 2D and 3D. The 2D AMFM uses only 2D features, and therefore works on any 2D images; 3D AMFM uses 3D features and can produce better results for 3D isotropic datasets [1].

2.3 Functional Analysis

2.3.1 Time-Series Image Analysis (TSIA) Task

Overall Concept. The system uses an axially acquired time-series volumetric MDCT data set of the thorax as the input. The acquisitions of these 4D functional scans

employ either ECG or respiratory image gating techniques. The TSIA single image task is selected from the wizard that then guides the user through the analysis process. First an algorithm is selected and the appropriate parameters are set based on previous analyses or current user input. Lungs are segmented and for perfusion analysis a reference ROI is selected within a pulmonary artery region. Results are output either as 3D floating point Analyze images or comma separated value files. On a typical medical image workstation, for a 5x5 grid, Ventilation analysis takes approximately 10 minutes and Perfusion 0.5 hour. These times can clearly be reduced with optimization efforts. For our current human ventilation studies, the total effective radiation exposures dose (HE) is 315/480 mrem for male/female; for perfusion studies, the total effective dose (HE) is 236.25/360 mrem for male/female.

Shared Algorithm Framework. Although the internal algorithm can be totally different from method to method, the basic programming interface for different time sequence image analysis algorithm can be made very similar by passing parameters, options, and results all in format of dynamic array. This makes it possible for us to define a shared algorithm framework and a shared GUI interface. The actual algorithms are designed as plug-in modules and will be dynamically loaded at run time.

Perfusion Algorithm. A central bolus injection of iodinated contrast agent is delivered by a power injector system during an ECG-gated axial dynamic MDCT scan. This software uses indicator dilution theory and first pass kinetics, assuming a bolus injection, residue detection model, to determine the regional Pulmonary Blood Flow (PBF) and perfusion parameters as previously described [8, 9]. Data is filtered to remove major airways and vessels. Mean Transit Time (MTT) and PBF normalized to the mean PBF of the imaged region are examined and heterogeneity can be estimated from its coefficient of variation (CV). MDCT-based perfusion measurements demonstrate significant differences in heterogeneity of perfusion parameters in subjects with imaging-only based evidence of lung pathology and show potential for detection of early inflammatory changes in the lung [10].
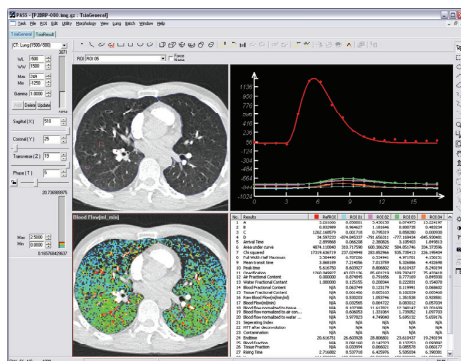
Fig. 4. This figure shows a screenshot of PASS using the TSIA Single Image Task using the perfusion algorithm on a human subject who is a smoker with early signs of emphysema on MDCT. Views: Gray scale CT image with user defined ROIs (upper left), Non-linear curve-fitted ROIs (upper right), Colormap of pulmonary blood flow (lower left) and per-ROI calculated perfusion parameters (lower right).

Ventilation Algorithm. Ventilation is evaluated from a Xenon wash-in series [5, 6, 7]. The density values for a ROI taken during the Xenon wash-in will yield a curve that can be fitted to a single exponential model with a time-constant τ. This time constant is equal to the inverse of the specific ventilation ($s\dot{V}$), the ventilation per unit volume. $\tau = 1 / s\dot{V}$.
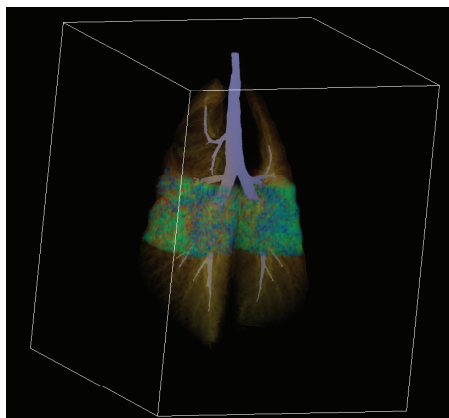
Fig. 5. Results from a ventilation study utilizing the 3D visualization component of PASS. Combination overlay of the lung parenchyma (yellow), airway tree (light blue), and 3D ventilation colormap.

3   Conclusion

Overlay functional data onto the structural image, in both 2D section slice and 3D volume rendering display, provides a unique way to include functional with structural data as efforts increase to use imaging as a means of determining disease phenotypes as part of a search for underlying genetic bases for disease susceptibility.

References

1. Xu Y., Sonka M., McLennan G., Guo J., and Hoffman E.A.: MDCT-based 3-D texture classification of emphysema and early smoking related lung pathologies. IEEE Transactions on Medical Imaging. 25(4), 464--475 (2006)
2. Uppaluri R., Hoffman E.A., Sonka M., Hartley P.G., Hunninghake G.W., McLennan G.: Computer recognition of regional lung disease patterns. Am J. Respiratory and Critical Care Medicine. 160, 648--654 (1999)
3. National Emphysema Treatment Trial Research Group: A Randomized Trial Comparing Lung-Volume Reduction Surgery with Medical Therapy for Severe Emphysema. The New England Journal of Medicine. 348(21), 2059--2073 (2003)
4. Armato S.G. 3rd, McNitt-Gray M.F., Reeves A.P., Meyer C.R., McLennan G., Aberle D.R., Kazerooni E.A., MacMahon H., van Beek E.J., Yankelevitz D., Hoffman E.A., Henschke C.I., Roberts R.Y., Brown M.S., Engelmann R.M., Pais R.C., Piker C.W., Qing D., Kocherginsky M., Croft B.Y., Clarke L.P., The Lung Image Database Consortium (LIDC): an evaluation of radiologist variability in the identification of lung nodules on CT scans. Acad Radiol. 14(11), 1409--21 (2007)
5. Simon B.A., Marcucci C., Fung M., and Lele S.R.: Parameter estimation and confidence intervals for Xe-CT ventilation studies: a Monte Carlo approach. J Appl Physiol. 84, 709--716 (1998)
6. Chon D., Simon B.A., Beck K.C., Shikata H., Saba O.I., Won C., and Hoffman E.A.: Differences in regional wash-in and wash-out time constants for xenon-CT ventilation studies. Respiratory Physiology & Neurobiology. 148(1-2), 65--83 (2005)
7. Tajik J.K., Chon D., Won C., Tran B.Q., Hoffman E.A.: Subsecond multisection CT of regional pulmonary ventilation. Acad Radiol. 9(2), 130--46 (2002)
8. Won C., Chon D., et al.: CT-based assessment of regional pulmonary microvascular blood flow parameters. J Appl Physiol. 94(6), 2483--93 (2003)
9. Chon, D., Beck K.C., et al.: Regional pulmonary blood flow in dogs by 4D-X-ray CT. J Appl Physiol. 101(5), 1451--65 (2006)
10. Alford S.K., van Beek E.J., Hudson M., Baumhauer H., McLennan G., and Hoffman E.A.: Characterization of regional alterations in pulmonary perfusion via MDCT in nonsmokers and smokers. European Congress of Radiology, March 2008. Oral Presentation B-820.
11. Hu S., Hoffman E.A., Reinhardt J.M.: Automatic lung segmentation for accurate quantitation of volumetric X-ray CT images. IEEE Transactions on Medical Imaging. 20 (6), 490—498 (2001)
12. Tschirren J., Hoffman E.A., McLennan G., Sonka M.: Intrathoracic airway trees: segmentation and airway morphology analysis from low-dose CT scans., IEEE Transactions on Medical Imaging. 24(12), 1529--1539 (2005)
13. Tschirren, J., McLennan, G., Palagyi, K. Hoffman, E.A., Sonka, M.: Matching and anatomical labeling of human airway tree. IEEE Transactions on Medical Imaging. 24(12), 1540--1547 (2005)
14. Shikata H., Hoffman E.A., and Sonka M.: Automated segmentation of pulmonary vascular tree from 3D CT images. Proc. SPIE. 5369, 107-116 (2004)
15. Ukil S., Reinhardt J.M., etc: Lung lobe segmentation in X-ray CT image using anatomic guidance. IEEE Transactions on Medical Imaging. (2008, in press)

Airway Inspector: an Open Source Application for Lung Morphometry

Raúl San José Estépar[1], George G. Washko[2], Edwin K. Silverman[2,3]
, John J. Reilly[2], Ron Kikinis[1], and Carl-Fredrik Westin[1]

[1] Surgical Planning Lab, Brigham and Women's Hospital, Boston, MA
{rjosest,westin,kikinis}@bwh.harvard.edu
[2] Pulmonary and Critical Care Division, Brigham and Women's Hospital, Boston,
MA {gwashko,jreilly}@partners.org
[3] Channing Laboratory, Brigham and Women's Hospital, Boston, MA
reeks@channing.harvard.edu *

Abstract. Quantitative analysis of computed tomographic (CT) images of the lungs is becoming increasingly useful in the medical management of subjects with Chronic Obstructive Pulmonary Disease (COPD) and other lung diseases. Airway Inspector is an open source initiative based on 3D Slicer for the morphometric analysis of airway trees and lung parenchyma to foster the quantitative needs to carry out image-based lung disease studies. Airway Inspector allows the user to navigate an airway tree and provide airway morphometric measurements at every points. The airway wall extraction can be done by several methods, namely, traditional methods like Full-Width at a Half Max (FWHM) and Zero Crossing of the Second Order Derivative (ZCSEC), and a new breed of methods based on phase congruency. Airway Inspector can also perform quantitative analysis of parenchyma diseases by means of histogram analysis of the lung parenchyma CT density.

1 Introduction

Recent investigations using quantitative image analysis of multislice CT (MSCT) data sets have provided new insights into the characterization of both emphysema and airway disease. For example, measurements of airway wall thickness on CT images are predictive of lung function; thicker airway walls tend to be found in subjects with reduced airflow [16]. This increased interest in using CT scans for the characterization of lung disease has unveiled the need for a software platform that can be used to implement reliable methods for lung CT analysis. Given the evolving nature of the field, an open platform is an optimal way to test different approaches and perform the necessary validation that is needed before those methods can be used in research networks and clinical trials.

Airway Inspector (www.airwayinspector.org) is a open source software platform for the morphometric analysis of lung CT images based on 3D Slicer. 3D
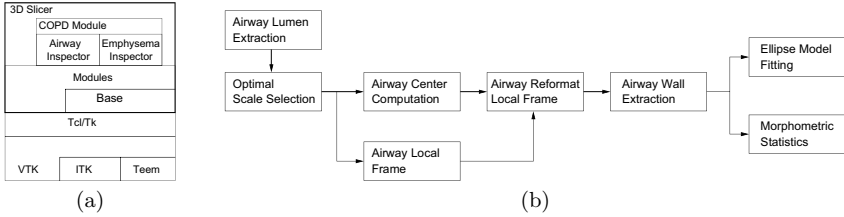
Fig. 1. (a) Airway Inspector and 3D Slicer software components and dependencies. (b) Computational components of Airway Inspector.

Slicer (www.slicer.org) is a free open source platform for medical image visualization and computing developed at the Brigham and Women's Hospital. 3D Slicer is available under a BSD-style license that has been also adopted by Airway Inspector. 3D Slicer has a modular architecture that allow the developers to extend its capabilities by adding new modules. The modules are mostly written in C++ and Tcl/Tk. While C++ is used to implement the main processing elements, Tcl/Tk is used for GUI design, interfacing and quick prototyping by means of wrappers to the methods implemented in C++. The main libraries that 3D Slicer depends on are the Visualization Toolkit (VTK) [19] and the Insight Toolkit (ITK) [23] for image visualization and image processing respectively. Additionally, 3D Slicer also depends on Teem [8], a multipurpose collection of libraries for representing, processing, and visualizing scientific multidimensional raster data. A diagram showing the dependencies of Airway Inspector is shown in Fig. 1a.

Airway Inspector is a component of the COPD module; a module for the analysis of both parenchyma and airway diseases. As such, the module implements different quantitative methods for the assessment of emphysema and the assessment of airway wall remodeling by means of Airway Inspector.

2 Airway Inspector: Computational Components

The main computational components of Airway Inspector are shown in Fig. 1b.

2.1 Airway Lumen Extraction

The entry point to Airway Inspector is a point inside the airway lumen under interrogation. This point can be either manually selected by an user or it can be automatically extracted by means of the segmentation of the airway lumen tree and the extraction of the centerline. For the automatic segmentation of the airway tree, Airway Inspector currently relies on the *Editor* module capabilities provided by 3D Slicer. The *Editor* module allows the user to perform a region growing process from a seed point placed, for example, in the trachea. After the segmentation of the airway tree, the COPD module provides a method to extract the centerline [2].

2.2 Continuous data reconstruction

An ubiquitous functionality across different computation components is the need of continuous data reconstruction, data interpolation and differentiation operations (first order and second order derivatives). To that end, Airway Inspector relies on Teem. The library inside Teem that takes care of the reconstruction is gage. Gage is an interface for continuous data reconstruction from sampled data. Gage implements different interpolation kernels; among them, the traditional cubic B-spline [13, 10] is used to compute intensity data at a continuous location and the corresponding gradient vector and Hessian matrix at that location.

2.3 Automatic scale selection

When querying the derivative operator at a given point, a critical factor is the scale at which those derivatives should be extracted. Based on the seminal work of Eberly [4], airways can be seen as valley lines of intensity. The Hessian matrix capture second-derivative information that can distinguishes ridges and valley lines, saddles, valley surfaces and ridge surfaces, and spheres. For a valley line the Hessian matrix, at the right scale, is characterized by two large positive eigenvalues ($\lambda_1 >> 0$, $\lambda_2 >> 0$) and one small negative eigenvalue ($\lambda_3 < 0$). At the valley point, the gradient should be orthogonal to the eigenvectors associated to the positive eigenvalues.

Although a formal scale-space analysis is not carried out [11], the automatic scale selection process seeks the optimal reconstruction kernel support, σ_{opt}, such as the following functional of the eigenvalues is maximized

$$\sigma_{opt} = \underset{\sigma}{argmax} \quad \frac{\lambda_1(\sigma) + \lambda_2(\sigma)}{2} - |\lambda_3(\sigma)| \tag{1}$$

where λ_i is the Hessian eigenvalue associated to the eigenvector $\mathbf{e}_i$. This optimal kernel size computed at the selected airway point is used hereon.

2.4 Airway Centering

The first step after a point inside the airway of interest has been chosen is to perform a centering of the point. The center is defined as the centroid, $\mathbf{x}_c$ of the luminal area. The luminal region is defined in the axial plane by means of a thresholding. The threshold is automatically defined as the mean value of the minimum and the maximum intensity values found in two rays in the form of a cross casted from the selected point.

After the centroid is located, a second step refines that location using the properties of the Hessian matrix. The centroid location, $\mathbf{x}_c$, is optimized using a gradient descent in the direction of the projected gradient, $\mathbf{g}_p(\mathbf{x})$, given by

$$\mathbf{g}_p(\mathbf{x}) = (\mathbf{I} - \mathbf{e}_3\mathbf{e}_3^T)\nabla I(\mathbf{x}). \tag{2}$$

The result of the optimization process is the final airway luminal center location, $\mathbf{x}_l$. The Hessian eigenvectors are computed at the optimal scale, σ_{opt}.

When the slice thickness is larger than the in-plane resolution, a full volumetric approach is not reliable, at least in the slice direction. Then, $\mathbf{x}_l$ is only computed in the axial plane by descending in the in-plane 2D gradient direction.

2.5   Airway Local Frame

The Hessian matrix at the optimal scale, σ_{opt}, can be used to estimate the direction of the airway longitudinal axis and define a local reference frame, if not already available through the centerline extraction process. The longitudinal airway axis, $\mathbf{a}_z$ is given by the eigenvector associated to the most negative eigenvalue at the airway center location, $\mathbf{x}_l$, such as $\mathbf{a}_z = \mathbf{e}_3(\mathbf{x}_l)$. A local frame for the airway locations, $\mathbf{x}_l$, is defined as the vectors $\mathbf{a}_x$ and $\mathbf{a}_y$ that form an orthonormal basis with $\mathbf{a}_z$ and at the same time are maximally aligned with the scan axial directions.

Based on the airway local frame and center, a 2D slice containing the whole airway is interpolated for further processing. Airway Inspector allows the user to perform the interpolation in two ways; either in the oblique plane or the native axial plane, if the slice thickness is not suitable for a reliable interpolation in the z-direction. The oblique plane is defined as the plane whose normal is given by the airway longitudinal axis, $\mathbf{a}_z$. The field of view (FOV) of the reformatted airway plane is equal to 25.6 mm, although this value can be modified accordingly to accommodate larger airways.

2.6   Airway wall extraction

The airway wall is defined by casting rays at regular angular increments from the center location, $\mathbf{x}_l$. For each ray, the intensity profile of the 2D reformatted slicer is interpolated at 0.05 mm. Each ray, $A_\rho(r)$, is a profile that comprises a lumen section, a wall section and a parenchyma section. Airway Inspector currently implements four methods for the definition of the inner and outer wall boundaries. In section 3, we will elaborate on these methods.

The result of the wall extraction process is the inner and outer radial locations, $r_i(\rho)$ and $r_o(\rho)$ respectively, at the polar angle ρ. The airway wall boundaries in the local frame defined by the axes $\mathbf{a}_x$ and $\mathbf{a}_y$ are given by point set $x_{i/o} = r_{i/o}(\rho)cos(\rho)$ and $y_{i/o} = r_{i/o}(\rho)sin(\rho)$ in Cartesian coordinates. The current implementation uses 128 rays equally distributed in the interval $[0, 2\pi)$ radians.

2.7   Ellipse Fitting

The airway extraction methods are based on a ideal circular model for the airway. Due to adjacent vessels and other image artifacts, the airway detection can be challenging and airway wall points can be noisy. To yield a more robust measurement, a ellipse model is fitted to the set of points $\{x_{i/o}, y_{i/o}\}$ using the method proposed in [7, 5]. Both the native ray-based measurements as well as the ellipse model parameters are available for the user.

2.8   Airway labeling

Currently the airways are label in four categories: right/left upper/lower lobe, based on the relative location of the airway center $\mathbf{x}_l$ with respect to the CT volume center. The user can redefine this assignment if it is not correct.

2.9   Morphometric statistics

From the polar parametrization of the inner and outer wall, several morphometric quantities are computed for each selected airway. For each quantity, mean, standard deviation, maximum and minimum values are reported. The morphometric quantities that Airway Inspector currently reports are:

- Inner, outer radius and wall thickness.
- Inner and outer perimeter: the perimeter is computed as a piecewise approximation based on the data points
- Inner, outer and wall area: the area is computed as a piecewise approximation of the area for each ray sector.
- Inner and outer ellipse parameters: the parameters of the ellipse fitting process are the length of the major and minor axes and the angle that the major axis with respect to $\mathbf{a}_x$.
- X-ray attenuation: airway inspector provides different metrics for the measured intensity inside the wall. Those metrics include the mean wall intensity inside the wall, peak wall intensity (maximum X-ray attenuation inside the airway wall for each ray) and the inner and outer luminal intensities (X-ray attenuation at the location of the inner and outer wall respectively).

3   Airway wall morphology

Airway Inspector currently implements four different methods for the extraction of the airway wall boundaries and it is poised to be an extensible platform where new airway wall extraction methods can be implemented and compared against is peers. Methods for the estimation of the airway wall can be divided in two categories: parametric [17, 3, 20, 22] and non-parametric [1, 18]. The former methods rely on an estimation of the scanner point spread function (PSF), therefore they can provide measurements beyond the limit imposed by the pixel resolution. The latter do not use any knowledge about the scanner PSF, therefore their measures are more limited in range. Airway Inspector currently only implements non-parametric methods given that they do not imply any knowledge about the acquisition process. The implemented methods can be divided in two categories: traditional methods and new methods based on phase congruency.

3.1   Traditional Methods

FWHM. FWHM defines the wall boundaries at the location where the intensity profile is half the intensity between the peak intensity inside the wall and the intensity of the inner/outer valley areas. Let r_{max} be the location of the maximum of the intensity profile $A_\rho(r)$ and r^i_{min} and r^o_{min} be the locations of the minimum intensity that are observed in the inner and outer walls, respectively. Formally, the FWHM is given by

$$r_i(\rho) = A_\rho^{-1}\left(\frac{A_\rho(r_{max}) - A_{CT}(r^i_{min})}{2}\right); \quad r_o(\rho) = A_\rho^{-1}\left(\frac{A_\rho(r_{max}) - A_{CT}(r^o_{min})}{2}\right),$$

$$(3)$$

where A_ρ^{-1} is the inverse function of A_ρ that maps intensity values into radial locations. As seen from eq. (3), a robust and accurate determination of the wall location based on the FWHM principle relies on a proper estimation of the peak and valley locations (also known as *extrema* points): ρ_{max}, ρ^i_{min} and ρ^o_{min}. These values can be worked out by means of the first order derivative.

$$\frac{dA_\rho}{dr}(r_{max}) = 0; \quad \frac{dA_\rho}{dr}(r^i_{min}) = 0; \quad \text{and} \quad \frac{dA_\rho}{d\rho}(\rho^o_{min}) = 0. \quad (4)$$

Moreover, if we want to distinguish between the peak and the valley points, we can inspect the sign of the second order derivative, such that

$$\frac{d^2A_\rho}{dr}(r_{max}) < 0; \quad \frac{d^2A_\rho}{d\rho}(r^i_{min}) > 0; \quad \text{and} \quad \frac{d^2A_\rho}{dr}(r^o_{min}) > 0. \quad (5)$$

In summary, the first order derivative of the intensity profile is the main operator that allows for the computation of the wall location based on the FWHM principle. FWHM is quite sensitive to the selection of these locations and slight differences in implementation can change the results yielded by FWHM. Airway inspector uses gage to compute the needed derivatives along the ray direction and the Newtwon-Raphson method [24] for the computation of the roots of the first-order derivative.

ZCSEC. Based on the edge detection theory proposed by Marr and Hildreth [12], the location of a boundary is defined as the inflection point between the valley intensity and the peak intensity. This inflection point is located at the zeros of the second order derivative; thus, the airway wall locations can be formally defined as $\frac{d^2A_\rho}{dr^2}(r_i) = 0$ and $\frac{d^2A_\rho}{dr^2}(r_o) = 0$. The zero location of the second order derivatives are computed using Netwon-Raphsons method.

3.2   Phase Congruency Methods

Based on the work by Morrone et al. [14], edge locations are related to those locations where the local Fourier components of the intensity signal are maximally in phase [9]. This *in-phase* behavior has been denoted as phase congruency and it has been used by the computer vision community as a feature descriptor [15, 6]. Airway Inspector defines the airway wall boundaries as the locations of maximal phase congruency. Phase congruency can be computed either by means of multiple CT reconstruction kernels by effectively changing the CT modulation transfer function or by means of a computational approaches using a single reconstruction kernel. We have previously reported these methods for the computation of airway wall location and we have shown their superiority with respect to traditional methods [18].

Single reconstructions. Phase congruency can be estimated by analyzing the input signal with a bank of complex-value quadrature filters, $\mathbf{g}_n(x)$, $n = 0, \ldots, N-1$. The output of a quadrature filter is a complex-value signal whose

magnitude and argument represents the local energy and the local phase of the input signal respectively. The quadrature filter bank used in Airway Inspector is a log-Gabor family given by, in the Fourier domain,

$$G_n(\omega) = e^{-\frac{\log^2(\frac{\omega}{\omega_n})}{2\log^2(\frac{\kappa}{\omega_n})}}, \qquad (6)$$

where ω_n is the central frequency of the filter and κ is related to the filter bandwidth. The filter bank is defined in Airway Inspector by three parameters that can be adjusted by the user: number of filters, N, the bandwidth of the filters, B (in octaves), and a multiplication factor that defines the separation of the filters, $m_{\lambda_{min}}$. The central frequency for each filter in the bank is given by the series $\omega_n = \frac{2\pi}{m_\lambda^n \lambda_{min}}$ where n is the n-th filter and λ_{min} is the spatial wavelength of the highest frequency filter and has a nominal value of 4 samples.

After filtering the 1D intensity profiles with the bank of quadrature filters, Airway Inspector computes phase congruency as

$$\Psi_\theta(r) = \frac{\sqrt{\left(\sum_{n=0}^{N-1} \mathrm{Re}\left\{A(r) * \mathbf{g}_n(r)\right\}\right)^2 + \left(\sum_{n=0}^{N-1} \mathrm{Im}\left\{A(r) * \mathbf{g}_n(r)\right\}\right)^2}}{\sum_{n=0}^{N-1} \|A(r) * \mathbf{g}_n(r)\|} \times \quad (7)$$
$$max\left(cos(\bar{\phi}(x) - \theta), 0\right),$$

where θ depends on the feature type we want to extract by means of phase congruency. Based on the meaning of the local phase, for the detection of the inner wall interface $\theta = \pi/2$ and for a outer wall interface $\theta = \frac{3\pi}{2}$. Then, the inner and outer wall location are estimated by finding the location that maximizes $\Phi_{\pi/2}(x)$ and $\Phi_{3\pi/2}(x)$ respectively. $\bar{\phi}(x)$ is the energy weighted mean local phase and can be approximated by

$$\bar{\phi}(x) = atan2\left(\sum_{n=0}^{N-1} \mathrm{Im}\left\{A(r) * \mathbf{g}_n(r)\right\}, \sum_{n=0}^{N-1} \mathrm{Re}\left\{A(r) * \mathbf{g}_n(r)\right\}\right), \qquad (8)$$

where $atan2(\cdot, \cdot)$ is the four quadrant inverse tangent.

Multiple reconstructions When multiple CT reconstructions are available, maximal phase congruency is expressed as a common crossing point in the intensity profiles of the kernels [18]. If multiple kernels are available, Airway Inspector reformats the airway for each kernel and detects the common crossing point for each 1D ray as the median intersection point for all possible pairwise combinations of kernels. The median operator is used because it is a robust estimator. If we have K kernels, the number of pairwise combinations is $M = \frac{K!}{2(K-2)!}$. Let V be the set of all pairwise combinations and $V(i)$ the $i - th$ pair of the set. Let $Inter_i$ and $Inter_o$ be the intersection operator that computes the intersection point between two kernel profiles for the inner and outer walls, respectively. The inner and outer wall locations are then given by $\rho_{i/o} = \mathrm{median}\left(\bigcup_{i=1}^{M} Inter_{i/o}\{V(i)\}\right)$ respectively.

FWHM   ZCSEC   PC

(a)   (b)   (c)
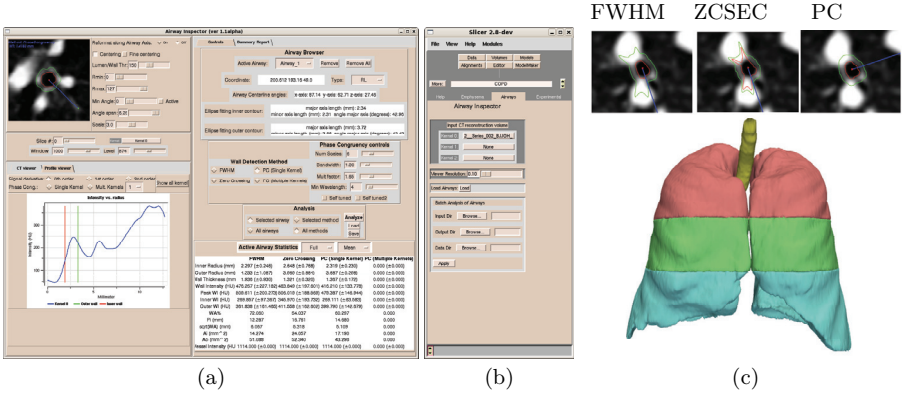
Fig. 2. (a) Airway Inspector GUI interface. (b) COPD module GUI. (c) Top- detected airway with FWHM, ZCSEC and phase congruency. Botton- Lung extraction.

4   Software functionalities

Figure 2a shows a snapshot of the current user interface for Airway Inspector. After the selection of several airway locations, the user can browse different airway locations and can probe the intensity profile by moving the mouse over the airway viewer. The airway wall is also plotted on the airway viewer for the method currently active. The panel below the airway viewer shows a plotting of the intensity profile, intensity derivatives, phase congruency responses ($\Phi_{\pi/2}(x)$ and $\Phi_{3\pi/2}(x)$) and the location of the estimated airway wall at the given ray for the selected method. In the left panel the user can control the analysis by selecting a detection method. The user can analyze either a given airway or all the loaded airways. The analysis can be run for a given wall detection method or all implemented methods. The results are reported in the lower left panel. A general statistical summary is shown as a separate tab on the left panel. Three airway viewer examples for each implemented detection method are shown in Fig. 2c. Phase congruency produces the most consistent results having the best performance in challenges areas like sections of the airway closed to vessels.

Additional features are the possibility of saving the current analysis in a Comma Separated Values (CSV) file format. This file can be later loaded into Airway Inspector and both the airway locations and the corresponding scan are automatically loaded for further analysis. A batch mechanism is in place (see Fig. 2b) that enables a automatic reanalysis of the data. The user can select the directories containing both the Airway Inspector files and the CT scans, and the software automatically loads every case, reanalyzes the data and saves the new analysis files into a directory.

Besides Airway Inspector, the COPD module integrate another package, Emphysema Inspector, for emphysema CT-based quantification based on histogram analysis. As part of this package, the lungs are automatically segmented into

three regions of equal volume, the trachea is extracted and the major vessels inside the parenchyma are also identified as shown in Fig 2c.

5   Conclusions

Airway Inspector is a Free Open Source software package for the morphometric analysis of airway trees. More information about the project, tutorials and source code can be found in www.airwayinspector.org. Although high-quality commercial software packages, like VIDA [21], exist, Airway Inspector is one of the first attempts to offer an open platform for the implementation of new airway morphometry methods that can foster the application of CT-based quantitative research. Airway Inspector is an evolving tool and some of the future features are the following:

- Implementation of a reliable method for automatic extraction of tubular structures in the lung: airway lumen and pulmonary vessels.
- Automatic filter bank design for phase congruency based on the prior knowledge of the structure size that is to measured.
- Automatic labeling of airway generations according to anatomical standards.
- Support for airway generation-based statistics.
- Group-wise registration of airway trees corresponding to a population to enable group statistical analysis.
- Porting to the next 3D Slicer generation (Slicer 3).
- Lobular segmentation.
- Integration with grid computing infrastructures to enable processing of large cohorts.

Further work is also needed to validate the airway extraction and wall detection methods across platforms. Both accuracy and precision are key components to study when designing a validation experiment. These kind of studies are very important for the adoption of computer-based CT morphometry of the lung in clinical research.

References

1. I. Amirav, S. S. Kramer, M. M. Grunstein, and E. A. Hoffman. Assessment of methacholine-induced airway constriction by ultrafast high-resolution computed tomography. *J. Appl. Physiol.*, 75:2239–2250, 1993.
2. Sylvain Bouix, Kaleem Siddiqi, and Allen Tannenbaum. Flux driven automatic centerline extraction. *Medical Image Analysis*, 9(3):209–221, 2005.
3. Y. Cheng, Y. Sato, H. Tanaka, T. Nishii, N. Sugano, H. Nakamura, H. Yoshikawa, S. Wang, and S. Tamura. Accurate thickness measurement of two adjacent sheet structures in ct images. *IEICE Transactions on Information and Systems*, E90-D(1):271–282, 2007.
4. D Eberly, R Gardner, B Morse, and S Pizer. Ridges for image analysis. *Journal of Mathematical Imaging and Vision*, 4:351–371, 1994.
5. A. W. Fitzgibbon, M. Pilu, and R. B. Fischer. Direct least squares fitting of ellipses. In *Proc. of the 13th International Conference on Pattern Recognition*, pages 253–257, September 1996.

6. G. H. Granlund and H. Knutsson. *Signal Processing for Computer Vision*. Kluwer Academic Publishers, 1995. ISBN 0-7923-9530-1.

7. R. Halir and J. Flusser. Numerically stable direct least squares fitting of ellipses. In *Proc. of the 6th International Conference in Central Europe on Computer Graphics, Visualization and Interative Dgital Media (WSCG'98)*, volume 1, pages 125–132, October 1998.

8. Gordon L. Kindlmann. http://teem.sourceforge.net/.

9. Peter D. Kovesi. Image features from phase congruency. *Videre: Journal of Computer Vision Research*, 1(3):1–26, 1999.

10. Thomas M. Lehmann, Claudia Gönner, and Klaus Spitzer. Survery: Interpolation methods in medical image processing. *IEEE Trans. on Medical Imaging*, 18:1049–1063, 1999.

11. T Lindeberg. Edge detection and ridge detection with automatic scale selection. *International Journal of Computer Vision*, 30(2):77–116, November 1998.

12. D. Marr and E. Hildreth. Theory of the edge detection. *Proc. of the Royal Society of London B*, 207:187–217, 1980.

13. D Mitchell and A Netravali. Reconstruction filters in computer graphics. In *Proceedings ACM SIGGRAPH 1988*, pages 221–228, August 1988.

14. M. C. Morrone, D. C. Burr, J. Ross, and R. Owens. Mach bands are phase dependent. *Nature*, 324(1280):250–253, 1986.

15. M. C. Morrone and R. A. Owens. Feature detection from local energy. *Pattern Recognition Letters*, 6:303–313, 1987.

16. Y. Nakano and S. Muro et al. Computed tomographic measurements of airway dimensions and emphysema in smokers. *Am. J. Respir. Crit. Care Med.*, 162:1102–1108, 2000.

17. Joseph M. Reinhardt, Neil D. D'Souza, and Eric A. Hoffman. Accurate measurement of intra-thoracic airways. *IEEE Trans. Medical Imaging*, 16(6):820–827, Dec. 1997.

18. Raúl San José Estépar, George G. Washko, Edwin K. Silverman, John J. Reilly, Ron Kikinis, and Carl-Fredrik Westin. Accurate airway wall estimation using phase congruency. In *MICCAI'06*, volume 4191 of *Lecture Notes in Computer Science*, pages 125–134, October 2006.

19. W. Schroeder, K. Martin, and B. Lorensen. *The visualization toolkit: an object-oriented approach to 3D graphics*. Prentice Hall, 1998.

20. G. J. Streekstra, S. D. Strackee, M. Maas, R. ter Wee, and H. W. Venema. Model-based cartilage thickness measurement in the submillimeter range. *Medical Physics*, 34(9):3562–3570, 2007.

21. Juerg Tschirren, Eric A. Hoffman, Geoffrey McLennan, and Milan Sonka. Intrathoracic airway trees: Segmentation and airway morphology analysis from low-dose ct scans. *IEEE Trans. Medical Imaging*, 24(12):1529–1539, 2005.

22. Oliver Weinheimer, Tobias Achenbach, Carsten Bletz, Christoph Düber, Hans-Ulrich Kauczor, and Claus Peter Heussel. About objective 3-d analysis of airway geometry in computerized tomography. *IEEE Trans. on Medical Imaging*, 27(1):64–74, 2008.

23. Terry S. Yoo, editor. *Insight into Images Principles and Practice for Segmentation, Registration, and Image Analysis*. A K Peters, 2004.

24. Tjalling J. Ypma. Historical development of the newton-raphson method. *SIAM Review*, 37(4):531–551, 1995.

Automated Lymph Node Labeling System

Atilla P. Kiraly[1], Benjamin J. Odry[1], David P. Naidich[2], and Carol Novak[1]

[1] Siemens Corporate Research, Princeton, NJ 08540, USA
[2] New York University Medical Center, New York, NY, 9999, USA

Abstract. Lung cancer staging typically requires locating, measuring, and labeling lymph nodes to determine affected nodes. Until recently, automation and workflow reduction has focused on the first two tasks. According to the classification scheme recommended by the American Joint Committee on Cancer and the Union Internationale Contre le Cancer, pulmonary lymph nodes are divided into four groupings with two to four stations per grouping. We present a system that automatically assigns proper group and station labels to lymph node locations within contrast enhanced chest CT images. The airways and aortic arch are automatically segmented to obtain an anatomic model of the patient. The model provides spatial features, such as distance and angle, used by a support vector machine to automatically provide a label for any given location. The model also provides interactive visual feedback, allowing the user to understand the relationship between the nodes and nearby anatomy for verification and for surgical planning.

1 Introduction

Recently the concept of automatically labeling lymph nodes has been presented to assist in cancer staging [1]. Previous automation approaches mainly focused on lymph node segmentation to assist in evaluation [2,3]. During cancer staging, lymph nodes are evaluated based upon condition and location. The cancer severity not only depends on the condition of the lymph node, but also on its anatomical location. According to the American Joint Committee on Cancer and the Union Internationale Contre le Cancer, pulmonary lymph nodes are divided into four groupings, each with several stations [4]. A scoring or evaluation of a patient involves assessing lymph nodes within each grouping.

We present a demonstration system for automated lymph node labeling and visualization. The user can select any lymph node and receives the label associated with that location. The labeling algorithm is independent of lymph node detection and segmentation methods and can be applied soon after the data is loaded. The topic of assisted labeling was later explored in [5] with a Bayesian approach to define station regions as discussed in [1]. However, not all components used in station definitions were acquired.

2 Method

The method proceeds by first obtaining a physical centerline and surface model of the airways and aorta in the given image. A user input of a lymph node location then produces several physical features relative to the models such as angles and distances that are used as a feature vector on a trained support vector machine (SVM) to produce a label. Further details are described below with complete details in [1].

2.1 Airway and Aortic Arch Modeling

The centerline and surface of the airways are obtained by an adaptive region growing method followed by skeletonization and refinement. The model describes the hierarchy of the airway tree and its physical location. The carina and left and right main bronchus can be determined from the model. The aortic arch is obtained by a tracking process focusing on a 3D response image. The top most region of the arch is then determined. These anatomical features are based on those used to define labels in the staging system.

A user input in a form of a 3D coordinate then produces a feature vector encoding relative distances, angles, and vectors relating the location to specific locations on the airway and aortic arch models. This vector is then passed to the SVM.

2.2 SVM classification

An SVM with a radial basis function is used to determine the label from the feature vector. The SVM was trained and evaluated on a total of 10 images with 86 labeled nodes. The labels in the ground truth were assigned by an experienced radiologist and then verified for by a second reader.

The SVM was first used to test and train on all of the datasets to determine the best features. A total of 8 features both from the airways and aorta were selected to be the most discriminating with the angle in relation the carina deemed the most important feature. The other features included the nearest distance to the airway tree and the distance to the top of the aortic arch.

In order to evaluate the method with the 8 selected features, round-robin testing was performed where all but one image was used to train the classifier with the remaining image used for testing. This procedure was repeated for each image. The results were 100% accuracy for group labeling and 76% for station labeling [1].

Even with a 76% accuracy for stations, staging scores rely more on the group labeling than station labeling to determine severity. Hence, the current method can still serve for automatic group labeling and assist in scoring.

3   System

Figure 1 shows the application. A patient dataset is loaded and the airway and aortic arch models are automatically determined. The user can click on any slice views (transverse, sagittal, coronal) and then have the label automatically generated. Since only a label is generated, the core labeling components of the system are portable and can be easy incorporated into dictation systems to provide further assistance with workflow. In this system, any user specified locations can be saved and documented into a case study. The determined label is automatically added to this case study.

Conversely, since the system has regions associated with station labels, the user can also focus on a particular region of the image given a specific station that they are interested in pursuing.
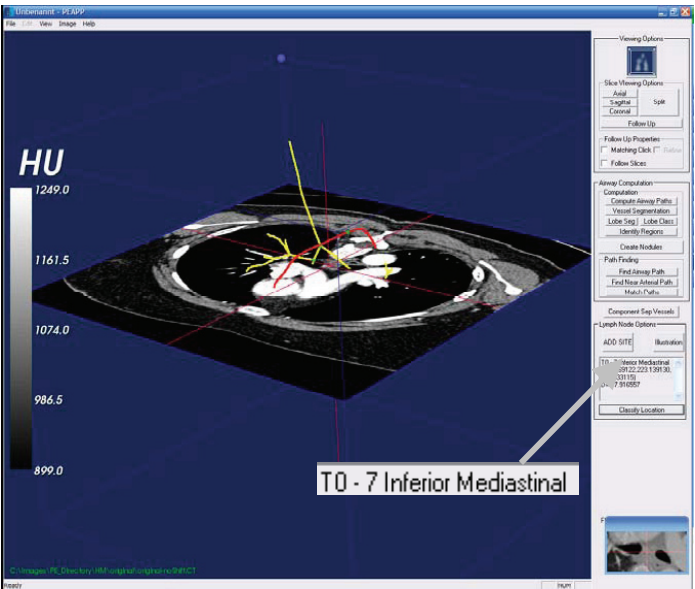
Fig. 1. The demonstration system for automatic lymph node labeling. In this example, the user has selected a location on the axial slice. The "Classify Location" button is pressed to produce the label associated with given lymph node location. This process is then repeated for each lymph node found. The labeled node and its location are then saved into the patient's case study for future reference.

4   Conclusions

We have demonstrated a prototype system for automated lymph node labeling. The interface is simple in that the user simply selects a location to immediately obtain a station label. Also, knowing a specific station, the region bounded by that station can also be obtained. The display of the models obtained allow for verification and better spatial understanding of the station label.

The labeling method is open to different labeling systems and image modalities since it based on physical features. In addition to benefiting the user, the regions defined by the method can be used as precursor inputs to lymph node detection and segmentation methods to help limit search regions.

The evaluation with the SVM provided a 76% accuracy for nodal station labels. Without an accurate model of the brachiocephalic artery, pulmonary artery, or pulmonary ligament, it is difficult to exactly model the labeling system. These are future components that must be incorporated to allow for a label determination on a Bayesian level. Without these components, an SVM allows us to maximize the accuracy with the available models and offer accurate group labels.

The system will be more complete with the incorporation of these additional models. However, the SVM would still be of use in determining useful features to help validate and even possibly improve existing lymph node station schemes. Valuable clues can be garnered to provide a more intuitive boundary description.

References

1. A.P. Kiraly, D.P. Naidich, L. Guendel, L. Zhang, and C.L. Novak, "Novel method and applications for labeling and identifying lymph nodes," SPIE Medical Imaging 2007, vol. 6511, Feb. 2007.
2. J. Dornheim, H. Seim, B. Preim, I. Hertel, and G. Strauss, "Segmentation of Neck Lymph Nodes in CT Datasets with Stable 3D Mass-Spring Models," Medical Image Computing and Computer-Assisted Intervention – MICCAI 2006, Volume 4191/2006, pp. 904-911.
3. G. Unal , G. Slabaugh , A. Ess , A. Yezzi , T. Fang , J. Tyan , M. Requardt , R.Krieg , R.Seethamraju , M. Harisinghani , and R. Weissleder, "Semi-Automatic Lymph Node Segmentation in LN-MRI," Proceedings of the IEEE Int. Conf. Image Processing, 2006.
4. J.P. Ko, E.A. Drucker, J.A. Shepard, C.F. Mountain, C. Dresler, B. Sabloff, and T.C. McLoud, "CT depiction of regional nodal stations for lung cancer staging," AJR Am J Roentgenol. 2000 Mar;174(3):775-82.
5. Lu, K., Merritt, S.A., Higgins, W.E., "Extraction and Visualization of the Central Chest Lymph-Node Stations," SPIE Medical Imaging 2008, Feb. 2008.
6. C. Cortes and V. Vapnik, "Support-Vector Networks," *Machine Learning*, vol. 20, no. 3, 1995, p. 273-297.