



Nicola Giuliani

Fully-Automatic Pulmonary Lobe Segmentation in CT Images

MASTER'S THESIS

to achieve the university degree of
Diplom-Ingenieur

Master's degree programme
Biomedical Engineering

submitted to

Graz University of Technology

Supervisor

Prof. Dr. Rudolf Stollberger
Institute of Medical Engineering

Advisors

Dr. Martin Urschler

Ludwig Boltzmann Institute for Clinical Forensic Imaging, Graz
DI Michael Pienn

Ludwig Boltzmann Institute for Lung Vascular Research, Graz

Graz, Austria, Dec. 2017

Abstract

Pulmonary lobe segmentation is of importance for the localisation and quantification of lung diseases. However, fully-automatic lobe segmentation is still a challenging task, especially in pathological lungs. A new algorithm based on refining an approximate lobe segmentation by minimizing an energy equation using graph cuts is presented. The energy equation combines anatomical information such as the airways, vessels and fissures with prior knowledge on lobar shape. In a first version of the algorithm, the lobe-based labelled airway is used to compute the approximate lobe segmentation. A second, refined version is relying on the lobe-based labelled vasculature. For this purpose an algorithm for lobe-based vessel labelling has been developed. Both versions of the algorithm were evaluated on two different datasets, including the LObe and Lung Analysis 2011 (LOLA11) dataset. The first version of the algorithm achieved a score of 86.6% at the LOLA11 challenge. The second version, including the labelled vasculature, achieved 90.1%, which is currently the highest score in this challenge for a fully automatic method. The benefit of pulmonary lobe segmentation for artery/vein separation is evaluated in further consequence. The clinical use of the algorithm is demonstrated by computing and evaluating quantitative readouts on each lobe.

Kurzfassung

Die Segmentierung von Lungenlappen ist von Bedeutung für die Lokalisierung und Quantifizierung von Lungenkrankheiten. Vollautomatische Lungenlappensegmentierung ist aber bis heute, und insbesondere bei pathologischen Lungen, eine herausfordernde Aufgabe. Ein neuer Algorithmus wird vorgestellt, der durch Minimierung einer Energiegleichung anhand von Graph Cuts eine approximative Lappensegmentierung verfeinert. Die Energiegleichung enthält anatomische Information über die Atemwege, Gefäße und Fissuren kombiniert mit Vorwissen über die Form von Lungenlappen. In einer ersten Version des Algorithmus werden die lappenspezifisch gekennzeichneten Atemwege verwendet um die approximative Lappensegmentierung zu berechnen. Eine zweite, weiterentwickelte Version des Algorithmus beruht auf den lappenspezifisch gekennzeichneten Gefäßen. Für diesen Zweck wurde ein Algorithmus für das lappenspezifische Kennzeichnen von Gefäßen entwickelt. Beide Versionen des Algorithmus wurden auf zwei unterschiedlichen Datensätzen, einschließlich des LObe and Lung Analysis (LOLA11) Datensatzes, evaluiert. Die erste Version des Algorithmus erzielte ein Ergebnis von 86.6% bei der LOLA11 Challenge. Die zweite Version, welche die gekennzeichneten Gefäße inkludiert, erreichte ein Ergebnis von 90.1%, das zurzeit den höchst erzielten Wert für vollautomatische Methoden in dieser Challenge darstellt. Der Nutzen von Lungenlappensegmentierung für Arterien/Venen Separation wird in weitere Folge evaluiert. Ein möglicher klinischer Nutzen des Algorithmus wird durch das Berechnen und Auswerten von quantitativen Maßen auf den einzelnen Lappen gezeigt.

Affidavit

I declare that I have authored this thesis independently, that I have not used other than the declared sources/resources, and that I have explicitly indicated all material which has been quoted either literally or by content from the sources used.

The text document uploaded to TUGRAZonline is identical to the present master's thesis dissertation.

Date

Signature

Acknowledgments

First and foremost I would like to thank Dr. Martin Urschler, for his encouragement and support. His guidance and valuable suggestions made the creation of this thesis possible.

I would like to thank the Ludwig Boltzmann Institute for Lung Vascular Research for the financial support and the possibility to work in a multidisciplinary team on a clinical relevant topic. In particular, I would like to express my gratitude towards DI Michael Pienn for his support during the creation of the thesis, especially for the medical and clinical part.

Finally, I would like to thank my family and friends for all of their love and encouragement. Without them I would not have become who I am today.

Contents

1	Introduction	1
1.1	Motivation	1
1.2	Contributions and Outline	4
2	Related Work	7
2.1	Methods not incorporating auxiliary structures	8
2.2	Methods incorporating auxiliary structures	8
2.3	Discussion	10
3	Method	11
3.1	Airways	12
3.1.1	Segmentation	13
3.1.2	Labelling	13
3.2	Lungs	16
3.3	Vessels	16
3.3.1	Segmentation	16
3.3.2	Labelling	17
3.4	Fissures	19
3.5	Segmentation using Graphs	23
3.5.1	Introduction	23
3.5.2	Minimum-Cut / Maximum-Flow	24
3.5.3	Minimum-Cut / Maximum-Flow Algorithm	26
3.5.4	Solving Energy Functions via Graph Cuts	30
3.5.5	A note on Nomenclature	31
3.5.6	The Alpha-Expansion Algorithm	32
3.6	Lobe Segmentation	34

3.6.1	Derivation of the Energy Equation	35
3.6.2	Refinement of the Segmentation Approach	41
4	Segmentation Evaluation	43
4.1	LBI dataset	44
4.2	LOLA11 dataset	44
4.3	Evaluation of Energy Minimization with Airways as approximation (EMA) on LBI dataset	44
4.4	Evaluation of Energy Minimization with Vessels as approximation (EMV) on LBI dataset	48
4.5	Evaluation of EMA on LOLA11 dataset	52
4.6	Evaluation of EMV on LOLA11 dataset	54
4.7	Discussion	56
5	Lobe based A/V-Separation Evaluation	59
5.1	Lobe wise A/V separation	60
5.2	Using Fissures for A/V separation	60
6	Quantitative Readouts	65
6.1	Definition of Readouts	65
6.1.1	Vessel Segment	65
6.1.2	Sum of Angles Metric	65
6.1.3	Vessel Density	66
6.1.4	Normalized Vessel Volume	66
6.2	Results	66
6.3	Discussion	74
7	Conclusion	75
A	List of Acronyms	77
B	List of Publications	79
B.1	2017	79
B.2	2018	80
	Bibliography	81

List of Figures

1.1	A schematic diagram of the lung.	2
1.2	Anatomical planes in a human.	5
3.1	Segmentation pipeline overview.	12
3.2	Lower airways containing the trachea and bronchi.	13
3.3	Acyclic graph extraction algorithm.	14
3.4	Airway branching points.	15
3.5	Anatomical variations of airway trees.	16
3.6	Vessel labelling with respect to the lobes.	17
3.7	Algorithm for labelling of the vasculature with respect to the lobes.	18
3.8	Complete vs. incomplete fissures.	19
3.9	Heavyside function.	21
3.10	Connected components of the fissure response mask before and after applying the size threshold.	22
3.11	Single-scale vs. multi-scale fissure segmentation result.	22
3.12	Directed weighted and undirected unweighted graph.	24
3.13	Equivalence between the minimum-cut and maximum-flow problem.	25
3.14	Min-Cut / Max-Flow algorithm [4]	26
3.15	Growth stage of the Min-Cut / Max-Flow algorithm in [4]	27
3.16	Augmentation stage of the Min-Cut / Max-Flow algorithm in [4].	28
3.17	Adoption stage of the Min-Cut / Max-Flow algorithm in [4]	29
3.18	“find_parent(p)” function for the adoption stage of the Min-Cut / Max-Flow algorithm in [4]	29
3.19	Example Minimum-Cut / Maximum-Flow Part 1	29
3.20	Example Minimum-Cut / Maximum-Flow Part 2	29
3.21	Example Minimum-Cut / Maximum-Flow Part 3	30

3.22	The alpha-expansion algorithm [5].	32
3.23	Example for minimizing an energy equation using α -expansion.	34
3.24	Example α -expansion graphs.	35
3.25	Approximate lobe segmentation.	36
3.26	Distance map on pulmonary vessels.	38
3.27	Lung lobe segmentation results for case 10 in the LOLA11 dataset.	41
3.28	Segmentation results for EMA and EMV.	42
4.1	Scatter plot for EMA results on the LBI dataset.	45
4.2	Median distances for EMA results on the LBI dataset.	46
4.3	Scatter plots for each fissure of EMA results on the LBI dataset.	47
4.4	EMA segmentation results for case 05 and 06 of the LBI dataset.	48
4.5	EMA segmentation results for case 21 and 22 of the LBI dataset.	48
4.6	Scatter plot for EMV results on the LBI dataset.	49
4.7	Median distances for EMV results on the LBI dataset.	49
4.8	Scatter plots for each fissure of EMV results on the LBI dataset.	50
4.9	EMA vs. EMV segmentation results for case 14 of the LBI dataset.	51
4.10	EMA vs. EMV segmentation results for the left lungs of case 15 and 20 of the LBI dataset.	51
4.11	EMA vs. EMV segmentation results for case 24 of the LBI dataset.	52
4.12	Poor EMV segmentation result for case 19 of the LBI dataset.	52
4.13	EMA segmentation results for cases 01, 13 and 42 of the LOLA11 dataset.	53
4.14	EMV segmentation results for case 29 and 36 of the LOLA11 dataset.	55
4.15	EMV segmentation results for case 01 and 06 of the LOLA11 dataset.	55
4.16	EMA vs. EMV segmentation results for case 07 of the LOLA11 dataset.	56
4.17	EMA vs. EMV segmentation results for case 05 and 21 of the LOLA11 dataset.	56
4.18	Consequences of a wrongly labelled vessel.	58
5.1	A/V labelling evaluation results.	60
5.2	Evaluation results for lobe-based A/V separation.	62
5.3	A/V labelling results on the whole lung vs. on the lobes for case 21.	62
5.4	Adjusted lung mask with removed fissures for A/V labelling.	63
5.5	A/V evaluation results for the removed fissures approach.	63
6.1	Vessel density, SOAM and noramlized vessel volume results.	67
6.2	Correlations between mPAP and SOAM.	68
6.3	ROC curves for SOAM as predictor for PH.	69
6.4	Correlations between mPAP and vessel density.	70
6.5	ROC curves for vessel density as predictor for PH.	71
6.6	Correlations between mPAP and normalized vessel volume.	72
6.7	ROC curves for normalized vessel volume as predictor for PH.	73

List of Tables

2.1	Overview of lung lobe segmentation methods not incorporating auxiliary structures.	9
2.2	Overview of lung lobe segmentation methods incorporating auxiliary structures.	9
3.1	Weights for the edges in the α -expansion algorithm	33
4.1	LOLA11 results for the lobe segmentation algorithm EMA.	53
4.2	LOLA11 results for the lobe segmentation algorithm EMV.	54
4.3	Updated LOLA11 results for lung lobe segmentation including the scores of the algorithms presented in this thesis.	57
5.1	Evaluation results for the A/V separation algorithm computed on the whole lung and computed independently on each lobe.	61
5.2	Evaluation results for the A/V separation algorithm computed on the whole lung and computed on the adapted lung mask.	64

Contents

1.1 Motivation	1
1.2 Contributions and Outline	4

1.1 Motivation

The human lung is divided into five lobes. The lobes are functional regions and anatomically separated by the pulmonary fissures, which are double layers of visceral pleura [1]. A schematic diagram of the lung is shown in Figure 1.1. The lungs are surrounded by a double layer of pulmonary pleura. The outer layer is called parietal pleura while the inner layer is called visceral pleura. A thin layer of pleural fluid is separating the two layers. The pulmonary fissures are formed by the visceral pleura that folds into the lung between the lobes. Usually the right lung contains two pulmonary fissures, the right horizontal and right oblique fissure, while the left lung comprises one pulmonary fissure, the left oblique fissure. From this it follows that the right lung is divided into three lobes, namely the right upper lobe, the right middle lobe and the right lower lobe. Similarly for the left lung, which is divided into the left upper lobe and the left lower lobe. The lack of a third lobe in the left lung can be explained by the location and shape of the heart [67] [26]. However, lung anatomy is known to already vary between individual healthy subjects and can significantly vary in pathological subjects. This includes cases, in which fissures are often incomplete or totally missing [24] [23] and on the contrary, cases, in which additional fissures are found [8]. Raasch et al. [49] reported 40% of incomplete fissures for the upper left major fissures and 46% for the lower left major fissures and even higher numbers for the right lungs: 70% for the upper right oblique fissures, 47% for the lower right oblique fissures and 94% for the right horizontal fissures.

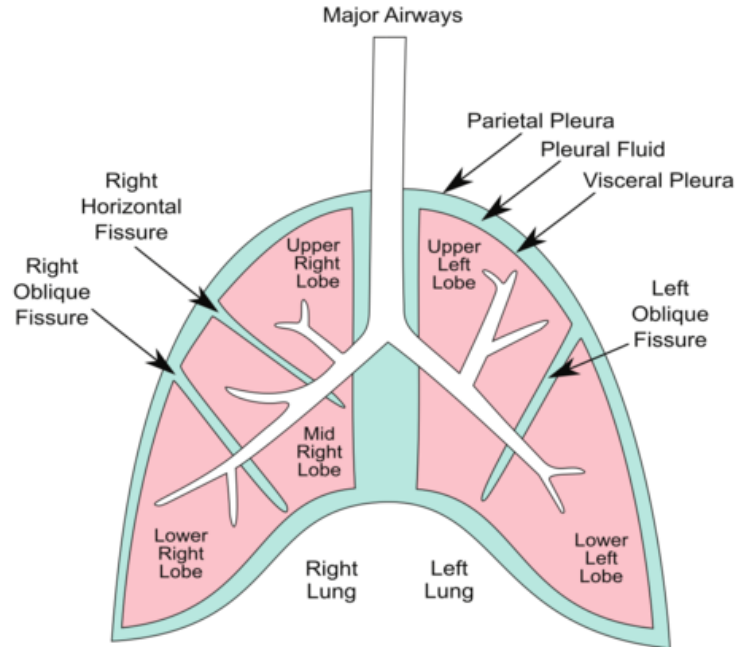


Figure 1.1: A schematic diagram of the lung. The major airways, fissures and lobes are shown. The fissures are formed by visceral pleura, which together with the parietal pleura forms a double layer surrounding the lungs [11].

In addition to being important landmarks for lung lobe segmentation, the pulmonary fissures provide information in localization of pulmonary lesions and assessment of disease processes [8]. Furthermore, fissure integrity and completeness, which is defined as the proportion between the fissures and the complete interlobar boundary, is of increasing interest [46] [64]. One application that greatly benefits from knowing fissure integrity is bronchoscopic lung volume reduction (**BLVR**) [30]. **BLVR** is a term for several procedures for treating hyperinflation in emphysema [42]. This includes a treatment that is based on the placement of endobronchial one-way valves blocking inspiratory flow, which is leading to atelectasis and volume reduction. However, an important factor that has great influence on the effectiveness of this procedure is collateral ventilation (**CV**), which is defined as "the ventilation of alveolar structures through passages or channels that bypass the normal airways" [7]. Atelectasis can thus often not be achieved in cases with **CV**. Missing pulmonary fissures can lead to interlobar **CV**. From this it follows that fissure integrity is an important parameter to assess before therapy [35]. However, the task of evaluating fissure completeness is tedious and time consuming for a radiologist and can be automated with the help of algorithms for fissure and lobe segmentation [45] [64]. Information on the pulmonary boundaries is likewise needed in the application of characterizing perifissural nodules [9], which are intrapulmonary

lymph nodes [27].

Apart from applications, in which identification of the lobar boundary is of indirect use (e.g. fissure integrity), the identification of pulmonary lobes is directly of importance for applications in treatment planning and disease assessment [11]. The identification as well as the segmentation of structures of interest is often a prerequisite for the detection and quantification of abnormalities. Exact quantification is a challenging task for humans, whereas computers generally yield better results [33]. Exact quantification is an important task for measuring the extent of emphysema and volume measurements often provide important diagnostic information [66]. Lung diseases are usually unevenly distributed across the lung. For example several lung diseases like emphysema [32], postprimary tuberculosis [40], silicosis [50] and idiopathic pulmonary fibrosis [58] act at a lobar level. It has been reported that patients suffering from upper-lobe emphysema have a higher chance of short term improvement after lung-volume-reduction surgery than patients suffering from lower-lobe emphysema [29]. A larger multicentric study showed that lung-volume-reduction surgery was beneficial for patients with predominantly upper-lobe emphysema combined with a low maximal workload after rehabilitation and resulted in lower mortality and general greater chances of improvements in symptoms and exercise capacity [13]. However, lung-volume-reduction resulted in higher mortality for patients with predominantly non-upper-lobe emphysema and a high maximal workload after rehabilitation.

In addition to the applications given above, lung lobe segmentation may also be of interest for the non-invasive assessment of Pulmonary Hypertension (PH). PH is defined as increased blood pressure in pulmonary arteries with a Mean Pulmonary Artery Pressure (mPAP) ≥ 25 mmHg [28]. Different types of PH exist and effort has been applied to define clinical classification schemes [56] [17]. Patients suffer from several symptoms including reduced exercise capacity and thoracic pain. Although the survival of patients with PH has improved over recent decades, PH is still a life threatening disease with low survival rates [3]. Invasive right heart catheterization is currently the gold standard for diagnosing PH [17]. New methods for diagnosing PH based on computer-aided analysis of CT images are emerging [25] [44]. The idea behind these methods is based on measuring vessel tortuosity in pulmonary vessels. The hypothesis states that due to the increased pressure in case of PH, the morphology of the pulmonary vessels changes and is showing an increased tortuosity [25]. The algorithm presented in [44] is based on segmenting the pulmonary vessels and separating the arteries from the veins. The tortuosity can subsequently be computed on the arteries to predict whether the patient is suffering from PH or not. To ensure a high prediction certainty, the artery/vein separation has to have high accuracy. However, the algorithm for artery/vein separation still faces problems including cases, in which the vessels of entire lobes are wrongly labelled. The authors suggest that running the algorithm separately for each lobe instead of the whole lung

may solve these issues. Lobe based evaluations of the lung are therefore of great interest and treatment planning can greatly benefit from lung lobe segmentation as in surgical treatments, the location as well as the volume and shape of the lobes can be of importance.

Due to improvements and wide availability of computed tomography (CT) scanners, CT has become the modality of choice for chest diseases [60]. Multiple X-ray measurements from different angles are reconstructed to a single three dimensional image containing multiple slices. In human anatomy three different planes are distinguished: the sagittal (medial) plane, the coronal (frontal) plane and the transverse (horizontal) plane. Figure 1.2 shows examples for different slices of a CT image including the anatomical planes. Varying radiodensities allow the differentiation of anatomical structures in CT images. Radiodensity is defined as the opacity to X-rays i.e. the relative inability of X-rays to pass through the material. This makes visual examination and assessment possible. E.g. the pulmonary fissures appear as bright sheet like structures in thoracic CT images. Using the fissures as point of reference, lung lobe segmentation can be performed by manually tracing the lobar boundaries. However, the described method, which is currently the gold standard, is rarely performed in clinical practice as it is extremely tedious and can take an experienced pulmonary radiologist up to multiple hours for a single scan [11]. As a result of this and a constantly increasing amount of acquired medical image data, the need for fully automatic lung lobe segmentation grows.

1.2 Contributions and Outline

Although several fully-automatic algorithms have been presented (see Chapter 2), lung lobe segmentation still remains an active area of research. In this work a new method for fully-automatic lung lobe segmentation is presented. The approach from [37] is followed, which is based on the computation of anatomical structures including pulmonary fissures, vessels and airways. Unlike in [37], where the anatomical structures are combined into a cost image, on which a watershed segmentation is performed, the anatomical structures in this work are combined into an energy equation, which by minimizing leads to the optimal lobe segmentation. In addition to the anatomical structures, prior knowledge on the shape of lung lobes is included into the energy equation. The main contribution of this work is the development of the energy equation.

The proposed method is presented in detail in Chapter 3. The first step consists of computing an airway segmentation that is based on the work presented in [25]. The airway segmentation is subsequently labelled with respect to the lobes [21]. To reduce the complexity of the following steps and the lung lobe segmentation problem, the lung segmentation derived from [25] is separated into left and right lung using the airways. All following steps are independently computed on the left and the right lung. Vessels trees are segmented using the method described in [44]. The fissures, which are the most important feature, are obtained following and extending the method presented in [37] to a multi-scale




Figure 1.2: Slices of a CT image; top left: coronal slice; top right: sagittal slice; bottom left: transverse slice; bottom right: planes of the human body [51]; red: sagittal plane; yellow: parasagittal plane; blue: coronal plane; green: transverse plane

approach. An approximate lobe segmentation is computed based on a distance map on the labelled airways. The Potts model [68], which minimizes the set of discontinuity is used as prior knowledge on lobar shape. All information is combined into an energy equation, which is minimized using the α -expansion algorithm [5].

The described method is further extended by using a more elaborate approximate lobe segmentation, which is obtained by computing a distance map not only on the labelled airway tree but on the labelled vessel tree. This approach, including the algorithm for lobe-based vessel tree labelling, is seen as a second major contribution.

Evaluation of the methods is shown in Chapter 4. Both methods are evaluated by participating in the LOLA11 (LObe and Lung Analysis 2011) challenge. The LOLA11 challenge is an open challenge part of the MICCAI 2011 conference, containing 55 lung CT images with varying abnormalities. By participating to the challenge, a score repre-

senting the overlap between the automatic segmentation result and a manually created segmentation is obtained. The methods are further evaluated on another dataset containing 25 CT images. In this experiment the accuracy of the lobe boundary was computed by measuring the distance between the automatically obtained boundary and manually placed markers on the fissures.

In addition to evaluating the performance of the fully-automatic lung lobe segmentation algorithm itself, the benefit of incorporating information on pulmonary lobes for the artery/vein separation algorithm presented in [44] is evaluated. Details can be seen in Chapter 5. The clinical applicability is demonstrated in Chapter 6 by computing clinical readouts on pulmonary vessels independently for each lobe. A conclusion of the thesis is given in Chapter 7.

Contents

2.1	Methods not incorporating auxiliary structures	8
2.2	Methods incorporating auxiliary structures	8
2.3	Discussion	10

Pulmonary lobe segmentation has been reported to be one of the most challenging segmentation tasks in the lung [66]. The gold standard for lobe segmentation is still manual segmentation by a radiologist. Manual segmentation is performed by following the pulmonary fissures in regions where they are present and information derived from the airway tree as well as from the pulmonary vasculature is used in regions, in which pulmonary fissures are missing. The same approach is followed by multiple methods for lung lobe segmentation algorithms, which is further described in Chapter 2.2. One challenge that emerges with this approach is to achieve stable segmentations of the auxiliary structures like the airways and the vessels, which in itself is not an easy task. To circumvent this problem another type of methods have been developed, which are not relying on auxiliary structures. However, this need not necessarily to be seen as the best approach as pulmonary fissures are often incomplete and sometimes even entirely missing. These methods are further described in Chapter 2.1. This chapter is intended to give a brief overview of already published works on pulmonary lobe segmentation but is not intended to be a complete review of these works. Most of the methods described in this chapter have already been discussed and compared to each other in [66] and [11]. However, direct comparison is often difficult to perform, as most of the methods have been evaluated on different datasets. The categorization of the methods follows the approach presented in [66]. In Chapter 2.3 a rationale, based on a discussion of the weak and strong points of the available methods, is given for the approach presented in this thesis.

2.1 Methods not incorporating auxiliary structures

Methods not incorporating auxiliary structures often use an atlas based approach. In [69] the authors proposed a method that consists of two main parts. In the first part an atlas is constructed. In the second part fissure contrast is enhanced by a ridgeness operator and an initialization is provided by the atlas. The fissures are further refined using fuzzy reasoning. However, manual interactions are needed in case of incomplete fissures, making this method not a reliable fully automatic approach. In [63] the fissures are extracted by supervised filtering. The lobe segmentation is subsequently obtained by classifying the voxels inside the lung based on their position and their distance to the fissures. Training examples are used to handle cases with incomplete fissures. [47] addressed the problem of incomplete fissures by using implicit radial basis functions to extend the fissures. Several methods have been proposed that followed the approach of fitting splines to the partially detected fissures in order to obtain a full delineation of the lobes. In [52] the authors first apply a particle filter to obtain candidate fissure voxels and remove spurious responses by maximum a posteriori estimation (MAP). The final lobe boundary is obtained by fitting a thin plate spline to the detected fissure points. A similar approach is followed in [22]. First, local surfaces are created by fitting a number of small planes using fissure points. The segmentation is assured to be smooth by using a quadratic B-spline weighting strategy in the next step for combining the local surfaces. Similarly, [55] are first enhancing fissure voxels, removing likely false fissure candidates and finally iteratively fitting a B-spline to the remaining fissure voxels. A different approach is followed by the authors in [53]. They first create lobe boundary surface models using PCA. After initially performing ridge surface feature sampling and removing noise by a simple connected components filtering step, the second step consists of shape model fitting to identify ridge surface particles that truly represent fissures. In Table 2.1 a summery of the above described methods is given. The column “LOLA11” represents the score achieved at the LOLA11 challenge. None of the methods not incorporating auxiliary structures has participated in the challenge.

2.2 Methods incorporating auxiliary structures

Several methods for lung lobe segmentation that are incorporating auxiliary lung structures exist. In Table 2.2 an overview of the methods is given. An anatomy-guided 3D watershed transform approach is followed in [36]. The anatomy is incorporated by computing a distance map on a vessel segmentation. Information on the fissure, where present, is added to the distance map, which is then used for the watershed transform. In [61] the authors use anatomical information from the airways and vessel trees to define a fissure Region Of Interest (ROI). The fissures are further refined by using the contrast information in the ROI. Incomplete fissures are extrapolated to complete the fissure surface. [65] follow a multiatlas approach. Anatomical information, including the fissures and the airway tree, are used to guide the registration of the atlas. The authors in [37] follow

Method	Description	LOLA11
[69]	Atlas, ridgeness, fuzzy reasoning	-
[63]	classifier based on position & distance	-
[47]	implicit radial basis functions	-
[52]	particle filter & MAP & thin plate spline	-
[22]	quadratic B-spline	-
[55]	iterative B-spline	-
[53]	ridge surface sampling & shape model fitting	-

Table 2.1: Overview of lung lobe segmentation methods not incorporating auxiliary structures.

Method	Description	LOLA11
[36]	fissure & distance map on vessels & watershed	-
[61]	airways & vessels & fissure detection in ROI	-
[65] [62]	fissures & airways & multiatlas	0.851
[37]	cost image (airways, vessels, fissures) & watershed	0.881
[48]	adaptive fissure scanning & implicit surface fitting	-
[6]	probabilistic fissure segmentation & groupwise prior	0.884

Table 2.2: Overview of lung lobe segmentation methods incorporating auxiliary structures.

the approach of computing a watershed segmentation on a previously constructed cost image. The cost image is derived by incorporating information on the fissures, vessels and airways. A method based on sagittal adaptive fissure scanning is developed in [48]. This is performed using information derived from the vessels and airways. The potential region is enhanced by a Hessian based line enhancement filter. The fissures are subsequently extracted by performing uniform cost search. Three implicit surface functions are obtained by performing a surface fitting step. A recent work [6] follows the approach of including information from the airways and vessels. The fissures are segmented considering the airways and vessels as priors in a probabilistic setting. In the next step a population prior is produced by combining the fissures using a group-wise registration framework. In the final step, a cost image, which is used for a watershed segmentation is obtained by combining the anatomical information, the fissures, and the group-wise prior.

2.3 Discussion

Even though multiple methods not incorporating auxiliary structures have been presented, none of them has been evaluated by the LOLA11 challenge. This makes it difficult to compare their performances, as each method is evaluated on their own dataset. In contrast, a few methods incorporating auxiliary structures have been evaluated by the LOLA11 challenge. One example is the method presented in [37], which uses information from the airway, vessels and fissures to compute a cost image that is used for a watershed segmentation. As already discussed, the advantage of solely relying on the fissures for lobe segmentation consists of the decrease in dependency of robust segmentation results of the auxiliary structures. However, it seems that this price is worth to be paid, as the frequency of lungs with partially incomplete or even missing fissures is not negligible. Methods based on registering a previously constructed atlas face problems in cases of anomalous anatomy. Since diseases often induce changes in lung anatomy, this poses a major problem for clinical applicability. As one of the requirements for the method developed in this thesis was to be robust against incomplete fissures and varying anatomy, the approach of incorporating auxiliary structures for the segmentation has been followed. By analysing the method in [37], room for improvement in two main points has been found: the fissure detection and watershed segmentation. The single scale fissure detection approach in [37] faced problems in cases of pathological thick fissures, which led to the motivation for extending it to a multiscale approach. The rationale for the replacement of the watershed segmentation by a graph cut based segmentation algorithm is the weakness of the watershed segmentation following shape priors, i.e. that pulmonary lobes have smooth surfaces. This is particularly important in regions absent of fissures, in which a smooth boundary should be interpolated. This shape prior is modelled by the Potts model. To take the fissures, the auxiliary anatomical structures and the shape prior into account, all information is combined into an energy equation by deriving different parameters from the individual structures. Additionally, the airway segmentation is used to generate an approximate lobe segmentation. The minimization of the energy equation is performed using the α -expansion algorithm.

Most methods incorporating vessels that have been presented to date, make use of the vessel segmentation for finding regions of high probability for the lobe boundaries. This is based on the idea that vessels usually do not cross the lobe boundaries. From this it follows that regions with a great distance to vessels have high probabilities for containing the lobe boundary. In this work, the vessel trees are additionally used in an even more sophisticated way. The idea of vessels not crossing the pulmonary lobes is followed one step further. If vessels are generally not crossing the lobar boundaries, then it should hold true that each vessel can be assigned to strictly one lobe. Following this idea, an algorithm for lobe based vessel tree labelling is developed and applied to generate an even more accurate approximate lobe segmentation.

Contents

3.1 Airways	12
3.2 Lungs	16
3.3 Vessels	16
3.4 Fissures	19
3.5 Segmentation using Graphs	23
3.6 Lobe Segmentation	34

The method for pulmonary lobe segmentation presented in this work is based on minimizing an energy equation using the α -expansion algorithm [5]. The aforementioned energy equation has been developed in this work. This approach allows to combine anatomical structures together with prior knowledge about lung lobes into one equation. Figure 3.1 shows an overview of the algorithm. The first level shows the extraction of anatomical structures that are useful for the lobe segmentation. The first step consists of segmenting the airways, which are then labelled corresponding to the lung lobes. Thereafter, the lungs are segmented and separated into left and right lung. Hereto the labelled airway segmentation is used as input for the lung segmentation. After separating the lungs, the vessels are segmented for the left and the right lung. The vessel trees are subsequently used as input for the fissure segmentation. The actual segmentation step is shown in the second row in Figure 3.1. First, an approximate lobe segmentation is computed using the labelled airways. All of the already computed lung structures from the first layer are combined into an energy equation. In the final step the energy equation is minimized, which yields an optimal lobe segmentation.

This section starts by first describing the methods how the anatomical lung structures needed for the segmentation are computed. Subsequently, the minimization of energy equations using α -expansion is explained in detail. Finally, the developed energy equation

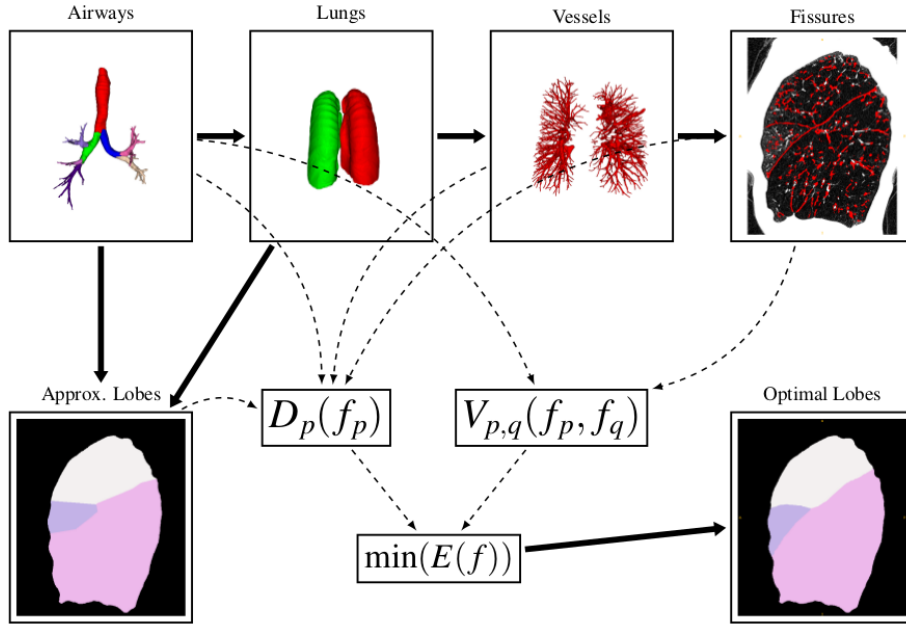


Figure 3.1: An overview of the segmentation pipeline.

for automatic lobe segmentation is presented and described.

3.1 Airways

In humans, the respiratory tract is divided into the upper airways and the lower airways. The upper airways consist of several parts starting from the nose going down to the portion of the larynx above the vocal folds. The lower airways include the trachea, bronchi and bronchioles. As the upper airways are of no interest regarding lung lobe segmentation, the term *airways* will from now on be used to describe the lower airways. Starting from the trachea, the airways branch at each division point into two or more smaller airways. These division points are also called generations. On average the human airway tree has 23 generations. The largest tube in the airways is the trachea. At the carina the trachea branches into two main bronchi for the left and the right lung. The right main bronchus and the left main bronchus branch further into smaller secondary bronchi, which are also called lobar bronchi. Further divisions of the bronchi are called subsegmental bronchi. The bronchi are called bronchioles when they are too small to contain cartilage. The structure of the airway tree follows a known pattern. As each lung lobe should work independently, the main bronchus branches split into the lobar bronchi, which deliver oxygen to the lobes of the corresponding lung. In case of the left lung this means that the left main bronchus subdivides into two lobar bronchi, one for the left upper lobe and one for the left lower

lobe. Likewise for the right lung, in which the right main bronchus subdivides into three lobar bronchi, one for the right upper lobe, one for the right middle lobe and one for the right lower lobe. In Figure 3.2 the airways can be seen.




Figure 3.2: The lower airways containing the trachea and bronchi.¹

3.1.1 Segmentation

The algorithm for airway segmentation [25] is based on an iterative region growing approach. The input CT image is first preprocessed by clamping all voxels with values greater than -400 Hounsfield units. After that, a filter for enhancing the airways is applied to the clamped image. The filter is based on analysing the eigenvectors of the Hessian matrix. Dark tube-like structures are enhanced using this filter. The actual segmentation of the airways is started by first detecting the trachea. This is achieved by scanning for a dark circle on the top-most slice. An iterative 3D region growing algorithm is subsequently started from the detected point inside the trachea. Two thresholds for the region growing are defined and continuously updated during the iterative region growing. Additionally, a mechanism monitoring the number of voxels added in the current iteration is established in order to prevent leakages during the segmentation.

3.1.2 Labelling

As previously described, the branching of the airways follows a known pattern. This pattern can be used for the purpose of labelling the airways with respect to the left and right lung. Using the labelled airway segmentation, it is possible to separate the lung segmentation described in Chapter 3.2 into left and right lung. Further, following the airway anatomy the airway tree can be labelled with respect to the lobes. This is of great

¹Gray, H., & Lewis, W. H. (1918). *Anatomy of the Human Body*. 20 edn. Lea Febiger, Philadelphia.

interest in this work as doing so would already provide the right labels (in terms of lobe segmentation) for a few voxels in the lung. In this work, the labelling of the airway tree was achieved by following the approach proposed in [21].

First, the airway segmentation has to be skeletonized. An undirected acyclic graph G is subsequently extracted from the airway skeleton. The algorithm for the extraction of the undirected acyclic graph consists of four steps and is shown in Figure 3.3. The set of vertices V is initialized with the first voxel of the skeleton corresponding to the trachea. The set S contains all voxels that are part of the skeleton. The edges of the graph G are kept in set E . The parameter ϵ in step 2 is set to 5mm.

Step 1: Find a nearest pair of points $S_i \in V$ and $S_j \in S \setminus V$, where S_i has less than three connected neighbours on G , let $V = V \cup S_j$ and $E = E \cup (i, j)$.
Step 2: Let S_j denote the last inserted point and find its nearest neighbour $S_k \in S \setminus V$. If the distance between S_j and S_k is smaller than a threshold ϵ , then let $V = V \cup S_k$ and $E = E \cup (j, k)$.
Step 3: If a new voxel S_k is added to V , go to Step 2; otherwise, go to Step 4.
Step 4: If $S \setminus V \neq \emptyset$, go to Step 1; otherwise, the graph construction is finished.

Figure 3.3: Pseudo code of the algorithm for the extraction of an undirected acyclic graph described in [21].

Following the algorithm in Figure 3.3, an undirected acyclic graph is obtained that consists of vertices that are either connected to one, two, or three other vertices. A vertex that is solely connected to one other vertex via an edge is called an endpoint. Furthermore, a bifurcation point is defined as a vertex, which is connected to three other vertices. Using the just described undirected acyclic graph, the labelling of the airway skeleton is achieved by first adding a root vertex to the graph. The root vertex is connected to the vertex corresponding to the trachea, which was found in the airway segmentation step described above. Doing so makes it possible to know the orientation of the graph with respect to the anatomy. The main bronchi are found by following the direction from the root vertex to the trachea vertex down to the first bifurcation point. At this point (A in Figure 3.4 on the right) the trachea splits up into the two main bronchi. The labelling of the branches into left and right lung is achieved by detecting the next branching points for both main bronchi. Comparing the x coordinates of the detected branching points (B and M in Figure 3.4 on the right) makes it possible to differentiate between the left and right airway branches. The labelling of the left upper lobe and left lower lobe is done by comparing the z coordinates of the next two branching points (U and L in Figure 3.4 on the right) starting from M. The labelling of the right lung with respect to its lobes is slightly more complicated. The right upper lobe can be found by following the same procedure used for the left lung. Starting from the branching point B, the next two branching points (H and

C in Figure 3.4 on the right) are found. The upper lobe is then found by again comparing the z coordinates. Identifying the right middle and right lower lobe is more complicated due to anatomical variability. The anatomical variability can be seen in Figure 3.5 on the right. Starting from branching point B all end points are detected. Two end points are used for the distinction between right middle and right lower lobe. One end point (E in Figure 3.5 on the left) with the maximal value of $(z - y)$ is selected together with another end point (F in Figure 3.5 on the left) with the minimal value of $(z - y)$. End point E corresponds to the end point of the right middle lobe while end point F corresponds to the end point of the right lower lobe. A branching point (D in Figure 3.5 on the left) is found by performing a depth first search (DFS) starting from E and F. D is the first vertex that is found by the DFS starting from E and F. Subsequently, all vertices between D and E are labelled as right middle lobe while all vertices between D and F are labelled as right lower lobe.




Figure 3.4: Left: Airway segmentation; Right: Skeleton of the airway segmentation on the left. Branching points that are used for the labelling are indicated. [21]

The lobe-based labelling of the airway tree is a good starting point for lung lobe segmentation. To avoid wrong labelled airway trees, topological variations have to be taken into consideration. Figure 3.5 on the right shows a topological variation that would result in a wrong labelled airway tree. To circumvent this problem a self correction step is applied. On the right in Figure 3.5 a supplementary branching point (C) for the right upper lobe can be seen. The self correction mechanism consists in determining and comparing the distances \overline{AB} (distance between branching point A and B) and \overline{BC} (distance between branching point B and C). In the common case (Figure 3.5 in the middle) the ratio $\overline{BC}/\overline{AB}$ is larger than in the exceptional case. A threshold of 0.5 for the ratio $\overline{BC}/\overline{AB}$ is defined for the common case. If the ratio is smaller than the defined threshold the self correction is triggered. As a consequence all voxels in the branch starting at the branching point B and in the branch starting at the branching point C are labelled as right upper lobe. Figure 3.6 on the left shows a result after applying the described airway labelling algorithm.




Figure 3.5: Left: Airway skeleton of a right lung. Branching points that are used for the labelling are indicated. Middle: Example of an airway tree with normal topology; Right: Example of an airway tree with abnormal topology [21]

3.2 Lungs

Segmentation of the lungs can be considered as first step towards lung lobe segmentation as it defines the **ROI**. The problem of lung lobe segmentation can further be simplified by separating the lung into left and right lung. This holds in particular for the left lung, in which case the lobe segmentation reduces to a binary classification problem where each voxel is either assigned the label of the upper lobe or lower lobe.

The methods described in this section are based on the works of [25]. The first step consists of separating bright from dark regions in the input CT image. This binarization is achieved following Otsu's method [43]. An initial segmentation of the lung is obtained by performing a connected components analysis [57] and selecting the largest dark region inside the bright region. The left and right lung are still merged through the trachea. The lungs are subsequently separated by computing the shortest path for each voxel to a reference point. The first branching point at which the trachea splits up into the main bronchi is selected as reference point. Voxels are then assigned the label of the left or right lung depending on their location and hence the direction of their path to the reference point. A morphological closing operation [20] is performed as a postprocessing step to remove holes inside the segmentation.

3.3 Vessels

3.3.1 Segmentation

The vessel segmentation is performed independently for each side of the lung. An algorithm, initially designed for artery and vein separation [44] is used in this work. However, unlike in [44] only the vessel segmentation is used without distinguishing between arteries

and veins.

The algorithm starts with a 4D vessel enhancement step. This is achieved using the Optimally-Oriented Flux (OOF) tubularity filter [38] [2]. To capture vessels of different sizes, the OOF is computed on different scales. This leads to a 4D vessel-enhanced image, in which the first three dimensions correspond to the spatial coordinates and the fourth to the radius of the vessels. The 4D vessel-enhanced image is subsequently used to construct a local maxima graph. The local maxima graph contains vertices corresponding to local maxima points in the enhanced image. The vertices are connected to all other vertices of the graph corresponding to points that are within a defined radius in the image. Using the constructed graph, 4D vessel paths are computed. This yields a directed overcomplete vascular graph which contains meaningful as well as spurious branches. In the last step, meaningful vessel trees are extracted from the local maxima graph. Similar to [59], an optimization procedure based on integer programming is used. Multiple trees are extracted at a time and the root nodes are computed implicitly by minimizing an objective function. The final step consists of resampling the subtrees and computing the actual branching points of the vessels. Figure 3.6 in the middle shows an example of a vessel segmentation obtained from the described algorithm.




Figure 3.6: On the left: labelled airway segmentation of a right lung; in the middle: vessel segmentation of the same lung; on the right: the vessel segmentation is labelled with respect to the lobes; the airway segmentation on the left are used to label the vessel segmentation in the middle; red vessels denote unlabelled vessels.

3.3.2 Labelling

Although [44] and others have already presented algorithms for labelling vessel segmentations according to arteries and veins, not much attention has been paid to label the vessels with respect to the lobes. However, by building on the work presented in [44] and making a couple of observations and assumptions about the lung vasculature, a lobe based vessel labelling has been developed in this work.

The problem with labelling the vessels is that in contrast to the airway tree, the vasculature does not divide into the lobes in a characteristic manner. This makes it impossible to follow the approach presented in Chapter 3.1.2. However, it is still possible to achieve lobe based labelling for the vasculature. The first important observation regarding the lung vasculature is that the pulmonary arteries are situated close to the airways. Using this observation, some vessels can already be labelled by finding vessels in close proximity to the labelled airway segmentation from Chapter 3.1.2. However, as can be seen in Figure 3.6, it is easier to achieve a detailed vessel segmentation than to achieve a detailed airway segmentation. In consequence, the labelling of the vessel segmentation will often be incomplete, due to the airway segmentation not reaching a significant part of the arteries and not reaching veins at all, as they are not closely situated to the airways. Therefore a second important observation, that usually there are no vessels crossing the pulmonary fissures, has to be made. This observation can be reformulated by saying that vessels do usually stay in one lobe. From this it follows that each vessel is only assigned to one lobe i.e. it does not happen that the first part of a vessel is located in the left upper lobe while the second part of the vessel is situated in the left lower lobe. Hence, having only one part of the vessel in close proximity to the airway is enough for labelling the whole vessel. The advantage of using the algorithm for vessel segmentation discussed in Chapter 3.3.1 is that it yields a result in which vessel trees are labelled independently. By using this result and following the two observations described above, it is easy to label a significant part of the pulmonary vasculature. The pseudo code for the algorithm is shown in Figure 3.7. The vessel segmentation for the algorithm described in Figure 3.7 is assumed to contain

1. Start with a labelled airway segmentation (A_S) and a vessel segmentation (V_T)
2. For each voxel a in A_S :
 - 2.1 For each voxel v in V_T in a radius of r_1 to a
 - 2.1.1 Label the whole vessel in V_T corresponding to v with the label of a
3. For each labelled voxel v_l in V_T :
 - 3.1 For each unlabelled voxel v in V_T in a radius of r_2 to v_l
 - 3.1.1 Label the whole vessel in V_T corresponding to v with the label of v_l
4. Return V_T

Figure 3.7: Algorithm for labelling of the vasculature with respect to the lobes.

information to discriminate between vessels. In step 2 of the algorithm, every vessel that is close enough (within a defined radius r_1) to the labelled airway segmentation gets assigned to the same label as their corresponding airway. Depending on the completeness of the airway segmentation, there might still be a considerable amount of unlabelled vessels left. These vessels get labelled in step 3 by looking for already labelled vessels in close proximity (defined by radius r_2). This step could theoretically be repeated multiple times before returning the final labelled vessel tree. On the right in Figure 3.6, a result for a

labelled vessel tree after applying the described algorithm is shown. Red vessels represent vessels that could not be assigned any label during the labelling process.

3.4 Fissures

The fissures are the most important feature when it comes to lung lobe segmentation. In fact, if the fissures were given and accurately labelled, the separation of the lung into lobes would be trivial. For the left lung, the upper lobe consists of all voxels above and the lower lobe of all voxels below the left oblique fissure. For the right lung, the upper lobe is given by all voxels above the right oblique and right horizontal fissure. The lower lobe is given by all voxels below the right oblique fissure. Finally the middle lobe is given by the remaining voxels or all voxels between the right oblique and right horizontal fissure.

Unfortunately, fissures are rarely fully complete [49] and not always easy to detect. However, even incomplete fissures partially define the lobe boundaries. Figure 3.8 on the left shows a sagittal slice of a right lung CT image with almost complete and good visible fissures. On the right of the same figure, another slice of the same CT-Image with incomplete fissures is shown. As apparent in Figure 3.8 on the right, even the incomplete fissures are a good point of reference for the boundary between the lobes.




Figure 3.8: On the left: A sagittal slice of a right lung CT-Scan; Fissures are complete and clearly visible. On the right: A sagittal slice of a right lung CT-Scan; Fissures are incomplete.

The thickness of the fissures can vary depending on the subject as well as within a subject. The difference in size and completeness of the fissures makes it difficult to detect them reliably. However, as fissures define the real boundary it is still worth making efforts segmenting them even if they are incomplete in many cases. One important property for

the detection of the pulmonary fissures is their shape. In CT scans they appear as bright sheet-like structures on dark background. Using the Hessian matrix \mathbf{H} , the local curvature can be analysed.

$$\mathbf{H} = \begin{bmatrix} \frac{\partial^2 f}{\partial x_1^2} & \frac{\partial^2 f}{\partial x_1 \partial x_2} & \cdots & \frac{\partial^2 f}{\partial x_1 \partial x_n} \\ \frac{\partial^2 f}{\partial x_2 \partial x_1} & \frac{\partial^2 f}{\partial x_2^2} & \cdots & \frac{\partial^2 f}{\partial x_2 \partial x_n} \\ \vdots & \vdots & \ddots & \vdots \\ \frac{\partial^2 f}{\partial x_n \partial x_1} & \frac{\partial^2 f}{\partial x_n \partial x_2} & \cdots & \frac{\partial^2 f}{\partial x_n^2} \end{bmatrix}. \quad (3.1)$$

The Hessian matrix \mathbf{H} (3.1) is a square matrix. Its coefficients are second-order partial derivatives of a scalar-valued function (3.2).

$$\mathbf{H}_{i,j} = \frac{\partial^2 f}{\partial x_i \partial x_j}. \quad (3.2)$$

In case of lung CT scans, the Hessian matrix is given by a 3x3 matrix with dimensions x, y and z . The analysis of the local image structure is done by computing the eigenvalues and eigenvectors for the derived Hessian matrix \mathbf{H} . This is done following (3.3), where \mathbf{H} is the given Hessian matrix, \mathbf{I} is the identity matrix, \mathbf{x} is the eigenvector and λ is the corresponding eigenvalue. In case of a 3x3 matrix, three eigenvalues with their corresponding eigenvectors can be found.

$$\mathbf{A}\mathbf{x} = \lambda\mathbf{x} \quad (\mathbf{A} - \lambda\mathbf{I})\mathbf{x} = 0 \quad (3.3)$$

We define the eigenvalues for the three dimensional case of a CT image as $|\lambda_1| \leq |\lambda_2| \leq |\lambda_3|$. In [16] the second order local structure of an image is examined. In case of a plate like structure λ_1 and λ_2 have usually low values while λ_3 has a very high value.

Based on this property, the authors of [37] proposed a fissure similarity measure $S_{Fissure}$ shown in (3.4).

$$S_{Fissure} = F_{Structure} \cdot F_{Sheet} \quad (3.4)$$

Two features $F_{Structure}$ and F_{Sheet} are combined. The first feature ($F_{Structure}$) is used for finding structure in the image. This is done by computing (3.5). The parameter α in (3.5) is set to 50 while the parameter β is set to 35. θ represents a Heavyside function shown in Figure 3.9. The Heavyside function is 0 for all values smaller or equal to 0. By using λ_3 as a parameter for θ , $F_{Structure}$ will be 0 for all voxels with an eigenvalue $\lambda_3 > 0$. This is important to suppress all voxels that are part of dark structures on light backgrounds as fissures are bright structures on dark backgrounds.

$$F_{Structure} = \theta(-\lambda_3) e^{\frac{-(\lambda_3 - \alpha)^6}{\beta^6}} \quad (3.5)$$

The second feature (F_{Sheet}) is used to detect sheet-like structures and can be seen in (3.6). The parameter γ is set to 25. F_{Sheet} gets smaller with increasing λ_2 . This leads to the suppression of vessels as fissures have usually smaller values for λ_2 .

$$F_{Sheet} = e^{\frac{-\lambda_2^6}{\gamma^6}} \quad (3.6)$$




Figure 3.9: Heavyside function.

Using the Hessian matrix \mathbf{H} and the fissure similarity measure $S_{Fissure}$ a mask of potential fissure voxels M_C is computed. This is done by first computing \mathbf{H} for every voxel in the image and subsequently computing $S_{Fissure}$. All voxels that satisfy the condition $S_{Fissure} > 0.1$ are kept in the mask. Figure 3.10 on the left shows an example for a mask M_C . It can easily be seen that the image still contains a lot of spurious results. To overcome the problem of spurious fissure voxels, M_C is filtered by a 3D-vector-based connected component analysis with a 6-neighbourhood. In case of a sheet, the eigenvector corresponding to λ_3 points perpendicular to the direction of the structure. Since a fissure usually does not bend much, the eigenvectors corresponding to λ_3 of two neighbouring voxels that are part of fissures should point into similar directions. Using this property a similarity measure for adjacent voxels can be formulated. The similarity of two vectors can be determined by computing the inner product of the normalized vectors. Two neighbouring voxels in the mask M_C are considered as connected, if the aforementioned inner product is larger or equal to 0.98. This leads to an image similar to the one shown in Figure 3.10 on the left, in which all voxels belonging to a connected component share the same colour. To reduce the number of falsely detected fissure voxels, the connected components are filtered using a size threshold. All components smaller than 100ml are rejected. Figure 3.10 on the right shows the resulting image after applying the size threshold to the image shown on the left.

For the computation of the Hessian matrix \mathbf{H} the differentiation is defined as a convolution with derivatives of Gaussians [16]. The mask M_C of potential fissure voxels shown in Figure 3.10 was obtained with a parameter $\sigma = 1.0mm$. This works fine in most cases when fissures have a normal thickness. However, in cases of pathological lungs, in which fissures can be a lot thicker than in healthy lungs, a parameter of $\sigma = 1.0mm$ can often be too small to capture these fissures. To circumvent this problem the fissures can be



Figure 3.10: Left: The mask M_C of enhanced fissure voxels. Voxels sharing the same colour belong to the same component after a vector based connected components analysis. Right: The remaining connected components after applying a size threshold to the image on the left.




Figure 3.11: Fissure segmentation results for a right lung with pathological thick fissures. The voxels detected as fissure from the algorithm are shown in red. Left: Result for a single scale approach. A parameter of $\sigma = 1.0mm$ was used. The algorithm is not even capturing parts of the fissure because of their pathological thickness. Right: Result for a multiscale approach including values for σ in the range of $[0.5mm - 2.0mm]$. The fissures are almost completely detected. However, the result also contains a lot of false positives.

computed on multiple scales. Empirical analysis showed, that \mathbf{H} computed with different values for σ in the range of $[0.5 - 2.0]$ can detect different fissure sizes. Combining the results computed on different scales can capture a much larger percentage of the whole fissure than computing \mathbf{H} on a single scale. The drawback of increasing σ to large values up to $2mm$ is that this can lead to a lot of vessels being falsely detected as fissures. To circumvent this problem the vessel segmentation described in Chapter 3.3 is dilated and subtracted from the multiscale fissure segmentation. In Figure 3.11 on the right, an

example for a multiscale fissure segmentation result is shown. On the left in Figure 3.11 the result using a parameter $\sigma = 1.0$ is shown. In the multiscale result, the pulmonary fissures are almost fully segmented despite of their pathological thickness. However, together with the pathological thick fissures there are still a lot of falsely detected structures in the segmentation result even though the vessel segmentation got subtracted. Although this problem could probably be solved by further filtering the resulting fissure image, not much effort is invested in this work as the proposed segmentation algorithm in Chapter 3.6 showed to be robust against false positives.

3.5 Segmentation using Graphs

Following the methods presented in the previous chapters, a segmentation of airways, lungs, vessels and fissures could be achieved. As the real goal of this work is pulmonary lobe segmentation, the question arises how to use the aforementioned auxiliary lung structures for partitioning the lung into its lobes (i.e. assigning each voxel one of five labels corresponding to the lobes). In this context it is common to follow a graph-based segmentation approach.

3.5.1 Introduction

A graph $G = (V, E)$ is an ordered pair comprising a set of vertices (nodes) V and a set of edges E . Edges are associated with two vertices and can thus be seen as a pair of vertices. This pair can either be ordered, in case of *directed* graphs, or unordered, in case of *undirected* graphs. Edges can be either weighted or unweighted. A graph G is called a weighted graph if weights are assigned to its edges E .

Figure 3.12 shows an example for a directed and weighted graph on the left. It consists of a set of vertices $V = \{\alpha, \beta, \gamma, \delta\}$ and a set of edges $E = \{(\alpha, \beta), (\beta, \gamma), (\gamma, \delta)\}$. The edges go from α to β , from β to γ and from γ to δ as indicated by the arrows in the graph while the weights assigned to the edges are 3, 1 and 4 respectively. The graph on the right side in Figure 3.12 is an example for an undirected and unweighted graph. Its set of vertices is given by $V = \{a, b, c, d, e, f\}$ and the edges are $E = \{\{a, b\}, \{b, c\}, \{c, d\}, \{d, e\}, \{e, f\}\}$. In this case the edges have no weights as the graph is unweighted. Due to the fact that the graph is undirected, the edges are written as a set of vertices $\{x, y\}$ instead of (x, y) as in the case of the directed graph. It is important to note here that (x, y) means specifically that the edge goes from x to y following that $(x, y) \neq (y, x)$.

Given a directed graph $G = (V, E)$ with positive edge weights and a source (s) vertex and a sink (t) vertex, we can define an $s - t$ -cut $C = S, T$ on G . The cut C partitions our vertices in two disjoint sets S and T , such that $s \in S$ and $t \in T$. Since the result of the cut C is a binary partition (two disjoint sets), it can be interpreted as a binary labelling,




Figure 3.12: Example graphs for a directed, weighted graph (left) and an undirected, unweighted graph (right)

such that all vertices in set S are assigned one label and all vertices in set T are assigned another label.

The cost of the cut C is the sum of costs (or weights) of all edges that need to be cut (that go from S to T) to obtain two disjoint sets.

$$c(S, T) = \sum_{\substack{e_{\{i,j\}} \in E, \\ |\{i,j\} \cap S| = 1}} c_e$$

A network $N = (G, c)$ is a directed graph $G = (V, E)$ with a non-negative *capacity-function* $c : E \rightarrow \mathbb{R}_{\infty}$. If a source node s and a sink node t are present, then $N = (G, c, s, t)$ is called a **flow network** [18].

A flow network is basically a directed and weighted graph with two distinguishable nodes s and t that has the additional information of a flow going from s to t . This means that each edge has a capacity and each edge receives a flow. One important restriction to point out is that for each node the amount of incoming flow must equal the amount of outgoing flow. This restriction has not to be satisfied for the source and the sink node, as the source node has only outgoing flow and the sink node only incoming flow. An example for a flow network is given in Figure 3.13. The numbers on the edges denote the flow (first value) and the capacity of each edge (second value) separated by “/”.

One concept, which will often appear in the context of maximum-flow algorithms (Chapter 3.5.2) is the **residual graph** G_f of the graph G with the flow f on G . G_f has the same node set V as G . The edge set is identical too, except for the weights that are assigned to them. Each edge $e(a, b)$ is assigned the weight $w_e - f(e)$. The corresponding reverse edge $e'(b, a)$ is assigned the weight $f(e)$.

3.5.2 Minimum-Cut / Maximum-Flow

An $s - t$ -cut on G is called minimum-cut if the cost of the cut is the minimum among all costs resulting from all possible $s - t$ -cuts on G . According to the

max-flow min-cut-theorem [14] [12], the value of the maximum flow in a graph is the same as the value of the minimum cut on the corresponding flow network. The maximum-flow problem is defined as maximizing the flow $|f|$ from s to t in a flow network.

A natural way of understanding this theorem is by looking at the graph as a system of pipes. Each pipe represents an edge in the graph and its corresponding weight can be seen as the diameter of the pipe. The maximum flow is then represented by the maximum amount of liquid that can go from one end (s -vertex) to the other end (t -vertex) of the pipe system. Further, we can also say that the limiting element for the maximum flow in the graph is given by the diameters of the smallest pipes in a layer of the system. There can not be more flow in the system than the amount of liquid that can pass through the limiting element. So finding this limiting element will give us the maximum flow. The limiting element is given by the layer of pipes with the smallest diameters, which represent edges with small weights. Thus, the maximum flow is the sum of the weights of the edges part of the limiting element, which has previously been defined as the minimum cut on the graph G .




Figure 3.13: A minimum-cut on a graph. On the left the dashed red line shows the minimum-cut. On the right the red edges denote the limiting edges in the flow graph. In the right graph the first value of the edge weights denotes the flow in the edge while the second value denotes the capacity.

Figure 3.13 shows an example for a minimum-cut on a graph. The dashed red line on the left graph represents the minimum-cut. All edges that get crossed by the dashed red line are part of the cut, so that $C = \{(s, \alpha), (\beta, \delta), (\beta, \gamma)\}$. The cost of the cut is given by the weights of the edges that get cut and is therefore $c(S, T) = 5 + 3 + 1 = 9$. All other possible ways of cutting the graph so that the source node s is in S and the sink node t is in T result in higher costs, which makes this cut the minimum-cut on G .

On the right side in Figure 3.13 the same graph as on the left is shown. In this case the graph is displayed as a *flow network*. The capacities of the edges are displayed as the second value on the edges and the flow on the edges are represented by the first value. The edge (s, β) for example has a flow of 4 and a capacity of 5. The maximum amount of incoming flow for the sink node t is given by the capacities of the two incoming edges (γ, t) and (δ, t) . They both have a capacity of 6, which means that the maximum amount of flow that can go in t is given by $6 + 6 = 12$. However, the maximum amount of flow

in the network is also limited by the flow that flows out of the source node s , which is in this case given by the capacities of the edges (s, α) and (s, β) and is therefore $5 + 5 = 10$. As the maximum outgoing flow from s is 10, the maximum incoming flow into t is also limited to 10. Node β has two outgoing edges $((\beta, \gamma), (\beta, \delta))$ and its maximum amount of outgoing flow is given by 4. Consequently, these two edges are limiting the amount of incoming flow to β , which is now also 4. Node δ has a maximum outgoing flow of 6 and is already receiving 3 incoming flow from β , resulting in a maximum amount of 3 more flow from α . Now α can forward 2 remaining flow to γ . Finally, t has a total amount of incoming flow of 9 that is determined by the outgoing flow of γ and the outgoing flow of δ . The red edges mark the edges of the limiting factor, which means that there can not be more flow from the source node s to the sink node t that can pass through these edges. Hence, the maximum flow is given by the capacity of these edges and equals to $|f| = 5 + 3 + 1 = 9$. This corresponds exactly to the cost of the minimum-cut on the left side in Figure 3.13.

3.5.3 Minimum-Cut / Maximum-Flow Algorithm

Several different algorithms for solving the minimum cut / maximum-flow problem on graphs exist [19] [10] [15]. One of them was presented by Boykov and Kolmogorov in [4]. This algorithm has been compared to several other graph cut algorithms and shown to outperform them with respect to speed [4].

```

initialize:  $S = \{s\}, T = \{t\}, A = \{s, t\}, O = \emptyset$ 
while true:
    grow  $S$  or  $T$  to find an augmenting path  $P$  from  $s$  to  $t$ 
    if  $P = \emptyset$  terminate
    augment on  $P$ 
    adopt orphans
end while

```

Figure 3.14: Min-Cut / Max-Flow algorithm [4]

An overview of the main steps is given in Figure 3.14. The graph that is used for this algorithm is of the structure described in Chapter 3.5.2, which means that it has a source node s and a sink node t . It is based on finding augmenting paths. Two separate search trees (S and T), as well as a set of active nodes (A) and a set of orphan nodes (O) are kept. Three different stages form the core of the algorithm. Before the iteration over the three main stages begins, an initialization is performed: The search tree S is initialized with the source node s and the search tree T is initialized with the sink node t . Additionally, the two nodes (s and t) are added to the active nodes set A . The orphans set O stays empty. As shown in Figure 3.14, the three stages (*growth*, *augmentation*,

adoption) get repeated until an empty path P is returned from the *growth* stage. Every iteration starts with the *growth* stage, in which a path P from the source s to the sink t is found. Then the previously found path P is augmented in the *augmentation* stage. The third stage is the *adoption* stage, in which the trees are restored. Regarding notation, $tree_cap(p \rightarrow q)$ denotes the residual capacity of either edge (p, q) if p is in tree S ($TREE(p) = S$) or edge (q, p) if p is in tree T ($TREE(p) = T$).

During the *growth* stage (details can be seen in Figure 3.15) the search trees are grown. Therefore, the active nodes in set A are processed. Every neighbouring node that is reachable from the current active node ($tree_cap(p \rightarrow q) > 0$) gets examined to determine whether it can be added to the same tree as the current active node. This is the case if the neighbour node is currently not part of a tree. When being added, the neighbour node is set to active and therefore added to the set A . Additionally, the active node is set as parent of the neighbour node. A new path P is found if the neighbouring node is in a different search tree as the active node. As soon as a new path P is found, the *growth* stage is being terminated and the path P is returned. The *growth* stage exits with an empty path $P = \emptyset$ only if there are no active nodes left in the set A . In further consequence this leads to the termination of the algorithm.

```

while  $A \neq \emptyset$ 
  pick an active node  $p \in A$ 
  for every neighbour  $q$  such that  $tree\_cap(p \rightarrow q) > 0$ 
    if  $TREE(q) = \emptyset$  then add  $q$  to search tree as an active node:
       $TREE(q) := TREE(p), PARENT(q) := p, A := A \cup \{q\}$ 
    if  $TREE(q) \neq \emptyset$  and  $TREE(q) \neq TREE(p)$  return  $P = PATH_{s \rightarrow t}$ 
  end for
  remove  $p$  from  $A$ 
end while
return  $P = \emptyset$ 

```

Figure 3.15: Growth stage of the Min-Cut / Max-Flow algorithm in [4]

The non-empty path P that was found during the *growth* stage is now being augmented. The pseudo code for the augmentation stage is shown in Figure 3.16. First, the bottleneck capacity Δ in P is found. The residual graph G_f is updated by pushing flow Δ through P . Doing so will cause at least one edge (p, q) in P to become saturated, which will result in the corresponding edge (p, q) getting an edge weight of zero ($w_{e_{(p,q)}} = 0$). In further consequence node q will no longer be reachable from node p and is thus becoming an orphan. It is important to distinguish three different cases. In the first case both nodes p and q are in tree S and node p is added to the set of orphans O as described above. However, if both nodes are in tree T , node p is added to the

set of orphans O . That is because the path $P_{s \rightarrow t}$ is always going in the direction from the source node s to the sink node t , which are the roots of the search trees S and T , respectively. So in the case of having both nodes in tree T , node p is the one becoming an orphan as it is the child of node q with respect to the search tree T . The third case is that node p and q are not in the same search tree. In that case neither node p nor node q is becoming an orphan as there is no parent / child relationship between those two.

```

find the bottleneck capacity  $\Delta$  on  $P$ 
update the residual graph by pushing flow  $\Delta$  through  $P$ 
for each edge  $(p, q)$  in  $P$  that becomes saturated
    if  $TREE(p) = TREE(q) = S$  then set  $PARENT(q) := \emptyset$  and  $O := O \cup \{q\}$ 
    if  $TREE(p) = TREE(q) = T$  then set  $PARENT(p) := \emptyset$  and  $O := O \cup \{p\}$ 
end for

```

Figure 3.16: Augmentation stage of the Min-Cut / Max-Flow algorithm in [4].

The last stage is called *augmentation* (Figure 3.17). Its goal is to restore the search trees and adopt the orphans. Each orphan gets processed and removed from the set of orphans O . The `find_parent` function is shown in Figure 3.18. A new valid parent q for a node p is found if there is an edge from q to the node p that is greater than zero, both nodes are in the same tree and the origin of the node q is either source or sink. This is important, because due to the *adoption* stage, some nodes in the search trees S and T are successors of orphan nodes. If a valid parent is found, $PARENT(p)$ is set to the found node q . In the case that no valid parent is found for node p , `find_parent` returns *FALSE* and all neighbours q of p that are in the same tree as p get processed. The neighbour node q is added to the active set A if there is an edge from q to p with capacity greater zero. If the parent of the neighbour node q is the orphan node p , node q becomes an orphan too. After all neighbours that had the same tree as p got processed, p gets removed from its tree and is removed from the active set A .

An example on how the algorithm works is shown in Figure 3.19, Figure 3.20 and Figure 3.21. Green and blue coloured vertices are in search tree S and T respectively, while black coloured vertices are currently in no search tree. Active vertices are visualized by a red border and rectangular-shaped vertices denote orphans. Dotted edges in the graphs represent edges that are part of the current path found in the growth stage of the algorithm. Red dotted edges denote edges that got saturated during the augmentation stage.

The graph is initialized in Figure 3.19 on the left with vertex s in search tree S and vertex t in search tree T . Both vertices s and t are set as active. The graph in the middle in Figure 3.19 shows the state after the first growth stage is done and the found path

```

while  $O \neq \emptyset$ 
  pick and orphan node  $p \in O$  and remove it from  $O$ 
  if  $find\_parent(p) = FALSE$ 
    for all neighbours  $q$  of  $p$  such that  $TREE(q) = TREE(p)$ :
      if  $tree\_cap(q \rightarrow p) > 0$  add  $q$  to the active set  $A$ 
      if  $PARENT(q) = p$  add  $q$  to  $O$  and set  $PARENT(q) := \emptyset$ 
    end for
     $TREE(p) := \emptyset, A := A - \{p\}$ 
  end while
  
```

Figure 3.17: Adoption stage of the Min-Cut / Max-Flow algorithm in [4]

```

for every neighbour  $q$  such that  $tree\_cap(q \rightarrow p) > 0$ 
  if  $TREE(q) = TREE(p)$  and  $ORIGIN(q) \in \{s, t\}$ :
     $PARENT(p) = q$ 
  return  $TRUE$ 
end for
return  $FALSE$ 
  
```

Figure 3.18: “find_parent(p)” function for the adoption stage of the Min-Cut / Max-Flow algorithm in [4]




Figure 3.19: Example Minimum-Cut / Maximum-Flow Part 1




Figure 3.20: Example Minimum-Cut / Maximum-Flow Part 2




Figure 3.21: Example Minimum-Cut / Maximum-Flow Part 3

was augmented. The edge (α, γ) has the smallest capacity of all edges in the path and is hence saturated by pushing the maximum flow through the path. On the right side in Figure 3.19 the state of the graph after the second growth and augmentation stage is shown. The edge (s, α) got saturated and hence α becomes an orphan. As visualized in Figure 3.20, no new parent is found during the adoption stage for vertex α leading to its removal from search tree S . After several cycles, vertex α is finally added to the search tree T on the left in Figure 3.21. The result is visualized on the right in Figure 3.21. The two search trees (S and T) are separated by three edges $\{(s, \alpha), (\beta, \delta), (\beta, \gamma)\}$, which have a combined flow of $5 + 3 + 1 = 9$.

3.5.4 Solving Energy Functions via Graph Cuts

As discussed earlier, a cut C on a graph G can be interpreted as a binary labelling and can thus be used for two label image segmentation. However, graph cuts can also be used to (approximately) solve more sophisticated problems like energy equations of the form

$$E(f) = \sum_{p,q \in N} V_{p,q}(f_p, f_q) + \sum_{p \in P} D_p(f_p) \quad (3.7)$$

The energy presented in (3.7) consists of a unary ($V_{p,q}(f_p, f_q)$) and a binary ($D_p(f_p)$) term. The unary term can also be called data penalty term as its purpose is to calculate a cost (penalty) for assigning a specific label to a pixel. For example the squared difference to the initial labelling can be used as the unary term, which will cause the system to try to stay close to the initial labelling.

$$D_p(f_p) = (f_p - I_p)^2$$

The binary term ($V_{p,q}(f_p, f_q)$) is used to penalize interactions between pixels and is hence often called data interaction term. One example for the data interaction term is the Potts model

$$V_{a,b} = k \cdot P(a, b)$$

with

$$\begin{aligned} P(a, b) &= 1 && \text{if } a \neq b, \\ &= 0 && \text{else} \end{aligned} \tag{3.8}$$

The Potts model penalizes neighbouring pixels with different labels and is thus leading to solutions with constant regions and smooth boundaries between those regions. Further, the Potts model is a regularizer, which in addition to keeping the boundaries smooth, penalizes the total interface length.

Finding a global minimum for the energy presented in (3.7), with the just introduced models for the data term and interaction term, can be shown to be NP-hard [5]. Nevertheless, algorithms exist to find local minimums that are guaranteed to be in a known factor to the global minimum. One such algorithm is the α – *expansion* algorithm [5].

3.5.5 A note on Nomenclature

At this point the question may arise why in the context of finding an optimal image segmentation, the term "energy minimization" is used. To answer this question it is helpful to consider the special case of the two-state Potts model [68], which is generally referred to as the Ising model [31]. It was developed by physicist Ernst Ising to explain empirically observed facts about ferromagnetic materials. In the two-dimensional case, the Ising model can be viewed as a grid of points representing spins (up or down). Ising then assigned an energy to each possible configuration of the system. The energy was of the form of (3.9), where the first term represents the energy caused by the interaction of the spins and the second term represents the effect of an external magnetic field of intensity H . The constants J and m represent properties of the material.

$$U(\omega) = -J \sum_{i,j} \sigma_i(\omega) \sigma_j(\omega) - mH \sum_i \sigma_i(\omega) \tag{3.9}$$

It is now possible to assign to each configuration probabilities that are proportional to

$$e^{-\frac{1}{kT}U(\omega)}$$

where k is a constant and T is the temperature. This leads to a probability measure of

$$P(\omega) = \frac{e^{-\frac{1}{kT}U(\omega)}}{Z}$$

where Z is the partition function. The configuration of highest probability $P(\omega)$ is then obtained at the minimum value of $U(\omega)$, i.e. the energy has to be minimized to find the optimal solution. A generalization of this representation is called a *Markov random field* [34]. The underlying structure of a Markov random field is an undirected graph, in which the nodes represent random variables and the edges represent dependencies. The

random variables can be any type of random variable that is satisfying the Markov property e.g. the spins of a ferromagnetic material in the context of physics or pixels of an image in the context of computer vision. Minimizing the energy gives the optimal solution, i.e. solution with the highest probability. The energy equation needs to be adapted according to the desired output. In a machine learning context the desired energy equation can be obtained by learning from data [39]. In this case a cost function or loss function is defined to associate low energies to desired configurations. The optimal energy equation is learned by minimizing the cost function. Minimizing the energy equation on the other hand is considered as *inference*, i.e. we obtain the optimal solution for the given input data X and the specific energy equation. Although in this work the energy equation is not obtained by learning from data and hence using the term *cost function* or *loss function* for the energy function should not lead to confusion between learning and inference, we still stick to the common nomenclature (energy equation) that is used in these fields.

3.5.6 The Alpha-Expansion Algorithm

1. Start with an arbitrary labelling f
2. Set success := 0
3. For each label $\alpha \in \mathcal{L}$
 - 3.1 Find $\hat{f} = \arg \min E(\hat{f})$ among \hat{f} within one α -expansion move of f
 - 3.2 If $E(\hat{f}) < E(f)$, set $f = \hat{f}$ and success := 1
4. If success = 1 goto 2
5. Return f

Figure 3.22: The alpha-expansion algorithm [5].

As the name suggests, the α -expansion algorithm is based on the computation of expansion moves. In an α -expansion, pixels that were previously not labelled α can be assigned the label α . However, it is not possible to assign an α -labelled pixel another label. The α -expansion algorithm from Boykov et al. is shown in Figure 3.22. The core of the algorithm is step 3, in which the energy is minimized for a label α with respect to one α -expansion move. This corresponds to the computation of a minimum-cut on a specifically designed graph. The construction of the graph will be shown below.

Create a *Source*(α) node and a *Sink*($\bar{\alpha}$) node. For every pixel (voxel) in the image $p \in \mathcal{P}$ add a vertex to the graph. Add two edges to every vertex (except *source* and *sink*): One edge to the *source* node and one edge to the *sink* node. We will call all edges that go to either *source* or *sink* *t-links*.

In the next step edges are created between all vertices that represent image pixels. If two neighbouring pixels $p, q \in \mathcal{N}$ have the same label $f_p = f_q$ they are directly connected via an *n-link*.

edge	weight	for
$t_p^{\bar{\alpha}}$	∞	$p \in \mathcal{P}_\alpha$
$t_p^{\bar{\alpha}}$	$D_p(f_p)$	$p \notin \mathcal{P}_\alpha$
t_p^α	$D_p(\alpha)$	$p \in \mathcal{P}$
$e_{\{p,a\}}$	$V(f_p, \alpha)$	$\{p, q\} \in \mathcal{N}, f_p \neq f_q$
$e_{\{a,q\}}$	$V(\alpha, f_q)$	
$t_a^{\bar{\alpha}}$	$V(f_p, f_q)$	
$e_{\{p,q\}}$	$V(f_p, \alpha)$	$\{p, q\} \in \mathcal{N}, f_p = f_q$

Table 3.1: Weights for the edges in the α -expansion algorithm.

If $f_p \neq f_q$ an auxiliary node $a_{\{p,q\}}$ is created and a set of three edges are added to the graph. One edge between p and $a_{\{p,q\}}$, one edge between $a_{\{p,q\}}$ and q and one edge between the auxiliary node $a_{\{p,q\}}$ and the *sink* node. The set of all edges can be written as

$$\varepsilon_\alpha = \left\{ \bigcup_{p \in \mathcal{P}} \{t_p^\alpha, t_p^{\bar{\alpha}}\}, \bigcup_{\substack{\{p,q\} \in \mathcal{N} \\ f_p \neq f_q}} \varepsilon_{\{p,q\}}, \bigcup_{\substack{\{p,q\} \in \mathcal{N} \\ f_p = f_q}} e_{\{p,q\}} \right\}$$

with

$$\varepsilon_{\{p,q\}} = \{e_{\{p,a\}}, e_{\{a,q\}}, t_a^{\bar{\alpha}}\}$$

And the set of all vertices can be written as

$$\mathcal{V}_\alpha = \{\alpha, \bar{\alpha}, \mathcal{P}, \bigcup_{\substack{\{p,q\} \in \mathcal{N} \\ f_p \neq f_q}} a_{\{p,q\}}\}$$

The weights assigned to the edges are shown in Table 3.1.

Figure 3.23 shows a 3x3 pixel binary image. The blue label is given the value 0 while the red label is given the value 1. It is now possible to minimize an energy function that is given by (3.10).

$$E(f) = \sum_{p,q \in \mathcal{N}} V_{p,q}(f_p, f_q) + \sum_{p \in \mathcal{P}} D_p(f_p), \quad \text{with} \quad (3.10)$$

$$\begin{aligned}
 V_{p,q}(f_p, f_q) &= 0 \quad \text{if } f_p = f_q \\
 &= 1 \quad \text{if } f_p \neq f_q \\
 D_p(f_p) &= 1.5 * |f_p - I_p|
 \end{aligned}$$

Changing the colour of a pixel from red to blue or vice versa leads to an increase of the global cost of 1.5, while maintaining the current colour (label) does not add any cost. For each neighbour that does not share the same colour as the current pixel the cost is increased by 1. Note that a 4-connected neighbourhood is used for this example. This means that the pixel (x, y) is connected to every pixel $(x \pm 1, y)$ and $(x, y \pm 1)$

Figure 3.24 on the left shows the α -expansion move for the red label. The graph is constructed following the rules described above. The grid represents the original 3x3 binary pixel image. The colour of the nodes in the middle of each pixel denotes the value of the pixel in the image. Edges with edge weight of zero are removed in the visualization for sake of readability. In order to change the label of a pixel its edge to the source node has to be cut. It is easy to see that the computation of a minimum-cut on graphs constructed following these rules corresponds to α -moves as nodes having the same label as the source node get attached to the sink by edges with infinite cost, which makes it impossible to be cut. The cut edges are visualized by red dashed lines. Figure 3.24 on the right shows the second cut on the graph constructed for the blue label using the α -expansion algorithm. This cut leads to no change in the labels of the pixels as the optimal result has already been found in the first cut. The final image can be seen on the right in Figure 3.23.




Figure 3.23: Left: Example 3x3 pixel image; Right: The result of minimizing the energy shown in equation 3.10 on the 3x3 pixel image.

In case of two label image segmentation as in this particularly easy example, the α -expansion algorithm can stop after computing a single minimum-cut for each label. By separating the lung into left and right lung and subsequently computing the lobe segmentation inside the lungs, the lobe segmentation problem for the left lung reduces to a two label image segmentation, which can thus be solved by a single cycle of the α -expansion algorithm.

3.6 Lobe Segmentation

As described and demonstrated before, the α -expansion algorithm is capable and suitable for minimizing complex energy equations of the form shown in (3.7). However,




Figure 3.24: Red dashed edges are part of the minimum cuts. Left: Example for an α -expansion for the red label. Right: An α -expansion for the blue label. The result of the α -expansion on the left is used as an input. Auxiliary nodes are removed for simplicity.

using the energy in (3.10) like in the example, would not make a lot of sense when it comes to lung lobe segmentation. This can be seen by analysing the unary term $D_p(f_p)$ of the equation. The term is solely checking if the value that a pixel gets assigned to differs from the value it has in the original image and increases the energy accordingly. This makes sense in the example but not for segmenting the lung into its lobes because the original values in the CT image contain no direct information for the lobes. When it comes to lung lobe segmentation, the goal is to separate the lung into five distinct lobes, which is achieved by assigning each voxel in the lung one of five labels, where each label corresponds to one lung lobe. The labels for the lobes can be chosen arbitrarily and have no direct meaning concerning the original pixel values in the CT image. Consequently, using the original pixel values directly in the unary term is not a promising approach. Whereas using lung structures that do contain information about the lobes would make a lot more sense. The challenge is now to combine all of the available information into an energy equation that leads to a meaningful lung lobe segmentation.

3.6.1 Derivation of the Energy Equation

In order to find the sought energy equation, several considerations about the lung anatomy and structures have to be made. First, we already know the right labels regarding the lobes for a small amount of voxels, namely the voxels of the airway tree. This was achieved by segmenting the airway tree and labelling it according to the lobes. It follows naturally that voxels in proximity of a specifically labelled airway branch have a high probability of sharing the same label as the branch. Thus, it appears that using a distance map on the different labelled airway branches and assigning voxels to the label of its nearest airway

branch may give a good approximation to the lung lobes. Computing this approximation is consequently also leading to a reduction in different voxel values from several hundreds or thousands to only five values (corresponding to lung lobes). Figure 3.25 on the left shows a sagittal slice of the CT image of a right lung. In the middle in Figure 3.25 the result of the above described airway distance based labelling is shown. When displayed as overlay, as shown on the right in Figure 3.25, it is clearly visible that this approach already leads to a good approximation to the lobes. From now on this approximation will be called the approximate lobe labelling.



Figure 3.25: On the left: A sagittal slice of a right lung CT-Scan; In the middle: voxels are assigned labels corresponding to their nearest airway branch; On the right: airway branch labelling is shown as overlay on the original CT image.

It would now perfectly make sense to use the approximate lobe labelling as an input for the α -expansion algorithm and to increase the energy by a specific amount (e.g. 1) for each voxel that gets assigned a label that differs from its original label. Note that in contrast to the example discussed in Chapter 3.5.6, no difference between the original and new label is calculated and used as a value for the increase in energy. This means that there is no difference in changing the label of a voxel from right middle lobe to right upper lobe or right lower lobe. For now the energy equation for the lung lobe segmentation would look as follows:

$$E(f) = \sum_{p,q \in N} V_{p,q}(f_p, f_q) + \sum_{p \in P} D_p(f_p), \quad \text{with} \quad (3.11)$$

$$V_{p,q}(f_p, f_q) = \text{undefined} \quad D_p(f_p) = \begin{cases} 0 & \text{if } f_p = I_p \\ 1 & \text{if } f_p \neq I_p \end{cases}$$

All the energy equation does is to penalize every discrepancy between the original lobe labelling and any other labelling. Therefore no label is changing and the system is static. The next logical step consists in including a term which motivates the system to adapt its labelling. In Chapter 3.5.4 we already introduced the Potts model (3.8), which penalizes neighbouring pixels with different labels. By minimizing the set of discontinuity, the Potts model is leading to solutions consisting of a minimum amount of constant regions with

the minimum sum of all smooth boundaries between those regions. Including the Potts model into the energy equation makes perfect sense as the lung lobes are self-contained structures with smooth surfaces. The unary and binary term of the energy equation are now given by

$$V_{p,q}(f_p, f_q) = k_p \cdot P(a, b) \quad D_p(f_p) = \begin{cases} 0 & \text{if } f_p = I_p \\ 1 & \text{if } f_p \neq I_p \end{cases} \quad (3.12)$$

Minimizing this energy equation would now lead to a solution with smooth boundaries while staying close to the approximate lobe labelling. The solution yielded by the minimization may be an improvement compared to the approximate lobe labelling, however it is still not very sophisticated as apart from the initialization it does not consider any lung structure. In Chapter 1 we already mentioned that the lung lobes are defined by the pulmonary fissures. The approximate lobe labelling is a good approximation to the lung lobes but it is not perfectly accurate when it comes to drawing the boundary on the fissures. Thus we could say that voxels are more likely to change its labels in proximity of the pulmonary fissures than in other regions. This can simply be modelled by setting the unary term to zero for voxels that are part of the fissures. A new factor k_f for the fissures is introduced leading to following terms:

$$V_{p,q}(f_p, f_q) = k_p \cdot P(a, b), \quad \text{with} \quad D_p(f_p) = \begin{cases} 0 & \text{if } f_p = I_p \\ k_f & \text{if } f_p \neq I_p \end{cases}, \quad \text{with} \quad (3.13)$$

$$P(a, b) = \begin{cases} 0 & \text{if } a = b \\ 1 & \text{else} \end{cases} \quad k_f = \begin{cases} 0 & \text{if } p \in \text{fissures} \\ 1 & \text{else} \end{cases}$$

Including the fissures significantly improves the result for CT images where fissures can be found by the proposed algorithm. However, as already discussed fissures are not always complete depending on the subject and severity of the diseases. Even if the fissures are present, it is often a hard task to fully detect them, which necessitates information from other structures in order to guarantee more stable results. A suitable structure to additionally use is the vessel tree. The vessels are usually separated into subtrees with respect to the lobes. Normally, the fissures are complete and vessels are not crossing them. By looking at an exemplary vessel tree as shown in Figure 3.26 on the left, one can see the just described properties. From the observations made about the vessel tree, we can deduce that the fissures are likely found in regions absent of vessels. This property can be used by computing a distance map on the vessel tree and using the resulting values in the image as probabilities for fissure occurrence. This means that for each voxel the probability of being part of one of the fissures is getting higher with greater distance to

the vessels. When it comes to minimizing the energy equation one must keep in mind that the α -expansion algorithm is computing minimum cuts on specifically constructed graphs. For the lung lobes this connotes that we want the graph to be cut at the lobe boundaries, which means that we want the edges in these regions to have low weights. For that reason the distance map on the vasculature is inverted, which is leading to low values where the probability is high for fissures and high values otherwise. Figure 3.26 in the middle shows the inverted distance map computed on the vessels. On the right in Figure 3.26 the true boundary is visualized as overlay in red on the inverted distance map. It can easily be seen that the distance map on the vasculature is a good additional indicator for locating the fissures.

By including a new factor k_v for the vessels into the unary term, the energy equation changes to (3.14). The vessels factor k_v takes values in the range between 0 and 1 and is incorporated by a multiplication. Doing so leads to very small, down to zero values for voxels that are either part of fissures or have high fissure probability due to the distance map on the vessels. In contrast, if no fissures are found because of low contrast or incomplete fissures in a region where the boundary actually should be drawn, there is still a chance that D_p will be low in this region because of k_v . Only if no fissures are found and vessels are close to the boundary, D_p will be high and thus falsely not contribute to cut in this region.




Figure 3.26: On the left: The vasculature of a right lung. The lobe boundaries can be guessed by looking at regions absent of vessels. In the middle: An inverted distance map computed on the vasculature. Dark regions denote regions with high fissure probability. On the right: The real boundary is displayed as an overlay.

$$\begin{aligned}
V_{p,q}(f_p, f_q) &= k_p \cdot P(a, b), \quad \text{with} & D_p(f_p) &= 0 \quad \text{if } f_p = I_p \\
& & &= k_f \cdot k_v \quad \text{if } f_p \neq I_p, \quad \text{with} \\
P(a, b) &= 0 \quad \text{if } a = b & k_f &= 0 \quad \text{if } p \in \text{fissures} \\
&= 1 \quad \text{else} & &= 1 \quad \text{else} \\
& & k_v &\in [0.0, 1.0]
\end{aligned} \tag{3.14}$$

Hitherto, minimizing the energy equation given by the unary and binary terms that were presented in (3.14), leads to good results. However, there is still room for improvements. For the calculation of the approximate lobe labelling, distance maps to the labelled airway branches had been used. The intuition for that was that the labelling of the airway tree is considered trustworthy and that voxels in proximity of a branch have a high probability of being part of the same lobe as the branch. At the same time voxels that are far away from any airway branch are not necessarily assigned the right label. They may be a bit closer to a specific branch than to another but still be significantly far away. This uncertainty due to distance can again be modelled by computing a distance map on the airway tree. With increasing distance the uncertainty increases, which should lead to a more liberal system in terms of label change in this area. To foster label changing, D_p must be small, thus the distance map on the airway tree has to be inverted. This property will be incorporated by the factor k_{a_d} .

In addition to using the overall distance to the airway tree, a second property concerning the airways can be used. When it comes to the approximate lobe labelling, there might still be false labelled voxels in close proximity to the airway tree near the first few branching points after the carina. Voxels in these regions are close to the airway tree but may still be wrong labelled because they may be few voxels closer to the wrong branch. Consequently there is an additional insecurity based on the relative distance to the nearest branch and the subsequent branch. Computing this relative distance can be done by first computing the distances to all branches (2 for the left lung and 3 for the right lung) and sorting them in ascending order. The relative distance would then be obtained by dividing the first distance by the second. A new factor k_{a_r} is introduced to model the relative distance.

These two properties can simply be combined into one factor k_a by a multiplication.

This leads to the following terms:

$$\begin{aligned}
V_{p,q}(f_p, f_q) &= k_p \cdot P(a, b), \quad \text{with} & D_p(f_p) &= 0 \quad \text{if } f_p = I_p \\
& & &= k_f \cdot k_v \cdot k_a \quad \text{if } f_p \neq I_p, \quad \text{with} \\
P(a, b) &= 0 \quad \text{if } a = b & k_f &= 0 \quad \text{if } p \in \text{fissures} \\
&= 1 \quad \text{else} & &= 1 \quad \text{else} \\
& & k_v &\in [0.0, 1.0] \\
& & k_a &= k_{a_d} \cdot k_{a_r}
\end{aligned} \tag{3.15}$$

The unary term D_p in (3.15) ensures that a lot of voxels change their labels in regions where the optimal lobe boundary is assumed. By using the Potts model in the binary term $V_{p,q}$, a smooth boundary near the supposed boundary is found. This does not lead to the best result since on top of having regions of minimal cross-sectional areas, the boundary should follow the fissures. In order to achieve this, the fissures have to be included into the binary term $V_{p,q}$. Placing the boundary along the fissures should not increase the energy. Therefore, a factor k_{f_b} is added to the binary term as follows:

$$\begin{aligned}
E(f) &= \sum_{p,q \in N} V_{p,q}(f_p, f_q) + \sum_{p \in P} D_p(f_p), \quad \text{with} \\
V_{p,q}(f_p, f_q) &= k_{f_b} \cdot k_p \cdot P(a, b), \quad \text{with} & D_p(f_p) &= 0 \quad \text{if } f_p = I_p \\
& & &= k_f \cdot k_v \cdot k_a \quad \text{if } f_p \neq I_p, \quad \text{with} \\
P(a, b) &= 0 \quad \text{if } a = b & k_f &= 0 \quad \text{if } p \in \text{fissures} \\
&= 1 \quad \text{else} & &= 1 \quad \text{else} \\
k_{f_b} &= 0.1 \quad \text{if } p, q \in \text{fissures} & k_v &\in [0.0, 1.0] \\
&= 1 \quad \text{else} & k_a &= k_{a_d} \cdot k_{a_r}
\end{aligned} \tag{3.16}$$

The factor k_{f_b} is either 1 if neither voxel p nor voxel q are part of the fissures and 0.1 else. By setting k_{f_b} to 0.1 we favour the system to change labels on the fissures since at the same time k_f in the unary term will be zero. Consequently, it will be cheaper in terms of cost to cut the t -link to the α vertex, which will assign the voxel the label α . On the contrary, setting both k_f and k_{f_b} to 0, could lead to isolated vertices in the graph. The isolated vertices would not change their labels and regions with alternating labels would occur in the segmentation. However, lung lobes are self-contained structures and fragmented regions are not a good representation for them. To circumvent this problem the Potts model is combined with the factor k_{f_b} in the binary term. Doing so will ensure that the boundary will be drawn on the fissures where fissures are found. Whereas, in

regions, in which no fissures are found, a smooth boundary will be interpolated in between the fissures. Figure 3.27 shows an example for a lung lobe segmentation result.




Figure 3.27: Lung lobe segmentation results for case 10 in the LOLA11 dataset (Chapter 4.2).

3.6.2 Refinement of the Segmentation Approach

As can be seen in Figure 3.27, minimizing the energy equation derived in Chapter 3.6.1 yields good results. However, since the approximate lobe segmentation is derived from the labelled airway tree, the quality of the airway segmentation is crucial for the whole segmentation process. Good results in the airway segmentation and labelling yield good approximate lobe segmentations and are therefore leading to good lobe segmentation results. However, if airway segmentations are weak, i.e. not reaching enough depth in the CT image, the quality of the approximate lobe segmentation decreases. In consequence, the labelled vessel tree derived in Chapter 3.3.2 can be used to address the problem of weak airway segmentations. The approximate lobe segmentation is computed on a combined image of labelled airways and vessels. This is generally leading to much better approximations, which are further leading to better lobe segmentation results. Following this approach technically does not change anything to the energy equation derived in Chapter 3.6.1. However, the parameter k_a does now not only represent the labelled airways but also the labelled vessels. Figure 3.28 on the left shows an example for a poor lobe segmentation result. As can be seen in the image, the boundary between right upper and right lower lobe does not follow the fissure. This can be explained by a weak approximate lobe segmentation, which was solely derived from the airways. A better approximation can be obtained by using the labelled vasculature. This is consequently leading to a better lobe segmentation result, which is shown on the right in Figure 3.28.

The refined version of the algorithm will be called from now on **EMV**, while the algorithm described in Chapter 3.6.1 will be called **EMA**.




Figure 3.28: Lung lobe segmentation results for the left lung of LOLA11 case 1. Left: the result obtained from the algorithm presented in Chapter 3.6.1; Right: result obtained from the algorithm presented in Chapter 3.6.2.

Segmentation Evaluation

Contents

4.1 LBI dataset	44
4.2 LOLA11 dataset	44
4.3 Evaluation of EMA on LBI dataset	44
4.4 Evaluation of EMV on LBI dataset	48
4.5 Evaluation of EMA on LOLA11 dataset	52
4.6 Evaluation of EMV on LOLA11 dataset	54
4.7 Discussion	56

The evaluation of the algorithm developed in this thesis is described in the following chapter. To provide an accurate and meaningful evaluation, the selected dataset is of great importance. Therefore, the evaluation was not only performed on an in-house dataset (Chapter 4.1) but also on a public available one (Chapter 4.2). In Chapter 4.3 the algorithm (EMA) presented in Chapter 3.6.1 is evaluated on the in-house dataset described in Chapter 4.1. The refined version of the algorithm, which is incorporating the labelled vessels (EMV; see Chapter 3.6.2) is evaluated on the same dataset in Chapter 4.4. Evaluations for both versions of the algorithm (EMA and EMV), are performed on the public dataset presented in Chapter 4.2 in the experiments shown in Chapter 4.5 and Chapter 4.6 respectively. In Chapter 4.7 the results of the performed experiments are discussed.

In all four experiments, the multi-scale fissure segmentation step is computed on a total of 16 different values for σ , starting from 0.5mm and increasing by 0.1mm in each step up to 2.0mm. The subtracted vessel segmentation was dilated with a voxel radius of 1. For the two EMV experiments, the radius r_1 for the vessel labelling algorithm was set to 3mm, while the radius r_2 was set to 1mm.

4.1 LBI dataset

The dataset described in this chapter is an in-house dataset consisting of 25 CT images. In addition to the CT data, clinical information such as the pulmonary pressure of the subjects is available. A manual reference standard for the evaluation of the algorithm was created based on this dataset. Two individuals that were not involved in the development of the algorithm manually placed approximately 500 markers on the fissures on each CT image. Distinct markers were used to allow an individual evaluation for the left oblique, right oblique and right horizontal fissure. No markers were placed in regions, where the fissures were not visible. Following this procedure led to one CT image not containing any markers at all, and two without markers in the right horizontal fissure, as the fissures were not visible.

4.2 LOLA11 dataset

This dataset is provided by the LObe and Lung Analysis 2011 (LOLA11) challenge [41]. It is comprised of 55 chest CT scans that are coming from a variety of sources representing clinically common scanners and protocols. The CT scans include a broad range of different abnormalities. The dataset includes a reference standard for lung and lobe segmentations, which were manually created. In contrast to the 55 CT scans, the reference standard is not made publicly available.

Evaluation is performed by computing overlaps between the segmentation results yielded from the algorithm and the manually created reference standard. The overlap between the automatic and reference segmentation is defined as the volume of their intersection divided by the volume of their union. Voxels that are part of a defined slack border of 2mm within the manually drawn border, are not taken into account for evaluation. Following this evaluation procedure, a score for each lobe is obtained by computing the mean overlap between the automatic and reference segmentation. The overall score for the lobe segmentation is then given by the mean of the means.

4.3 Evaluation of EMA on LBI dataset

In this experiment, the performance of the algorithm (EMA) presented in Chapter 3.6.1 is evaluated. That is to say that solely the airways are used for the computation of the approximate lobe segmentation. The evaluation was performed on the dataset described in Chapter 4.1. For each manually placed marker, the distance to the nearest automatically detected lobar boundary was measured. The distances were measured as positive or negative depending on their relative location to the automatically detected lobe boundary. For the left lung this means that markers that are part of the lower lobe in the automatic segmentation result are considered as positive while markers that are part of the upper lobe are considered

as negative. In case of the right lung, negative values were either used for markers for the horizontal fissure that were not part of the middle lobe in the automatic segmentation result, or markers for the oblique fissure that were not part of the lower lobe.

Following the described evaluation protocol the median distance was 0.00mm (interquartile range: -0.64-0.00mm). Median distances for the left oblique fissures were 0.00mm (0.00-0.78mm). The right oblique fissures median distance was 0.00mm (-0.64-0.00mm) and the median distance for the right horizontal fissure was 0.00mm (-1.04-0.00mm).

Using the same evaluation protocol with absolute distances instead of signed distances, the median distance was 1.04mm (0.88-1.10mm). Median distances for the left oblique fissures were 0.93mm (0.87-1.05mm). The right oblique fissures median distance was 0.94mm (0.87-1.06mm) and the median distance for the right horizontal fissure was 1.06mm (1.03-1.10mm).

Figure 4.1 shows the combined results for all fissures and each case. No results are available for dataset LBI_Pilot11 due to a missing manual reference as described in Chapter 4.1. In Figure 4.2 on the left the median distances for all cases are visualised. On the right in Figure 4.2 the median absolute distances are shown. On the top, in the middle and the bottom in Figure 4.3 the results for the left oblique, right oblique and right horizontal fissures of each dataset are shown, respectively.




Figure 4.1: Scatter plot including the combined results of all fissures. The results were obtained by computing EMA.

Figure 4.4 shows the segmentation results for case 05 and case 06. As can be seen in the images, the detected boundaries follow precisely the fissures. The right oblique fissure in the right most image in Figure 4.4 is completely missing. Even though the right horizontal fissure is not visible in this scan, the algorithm finds a plausible boundary between the




Figure 4.2: Median distances for all fissures. The results were obtained by computing [EMA](#).

right upper lobe and the right middle lobe. Figure 4.5 shows the segmentation results for case 21 and case 22. Similar to the results in Figure 4.4, the boundaries are precisely following the fissures. In case of missing fissures, a meaningful boundary is found. As can easily be seen in Figure 4.2 on the right, lobar results in only 5 cases show median distances larger than 10mm. This includes 3 cases for the left oblique fissure and 1 for the right oblique and right horizontal fissure. The three cases for the left oblique fissure are shown on the leftmost image in Figure 4.9, on the leftmost image in Figure 4.10 and in the third image from the left in Figure 4.10 respectively. The cases for the right oblique and right horizontal fissure are shown on the left in Figure 4.11 and in the third image from the left in Figure 4.9 respectively.




Figure 4.3: Top: left oblique fissures; middle: right oblique fissures; bottom: right horizontal fissures; The results were obtained by computing EMA.




Figure 4.4: Segmentation results computed by EMA for case 05 (images on the left) and case 06 (images on the right) from the in-house dataset (Chapter 4.2).




Figure 4.5: Segmentation results computed by EMA for case 21 (images on the left) and case 22 (images on the right) from the in-house dataset (Chapter 4.2).

4.4 Evaluation of EMV on LBI dataset

This experiment is basically the same as the one presented in Chapter 4.3. The only difference is that the labelled vessel tree was used to compute the approximate lobe segmentation as described in Chapter 3.6.2 (EMV).

Following the described evaluation protocol the median distance was 0.00mm (interquartile range: 0.00-0.00mm). Median distances for the left oblique fissures were 0.00mm (0.00-0.78mm). The right oblique fissures median distance was 0.00mm (0.00-0.00mm) and the median distance for the right horizontal fissure was -0.84mm (-1.05-0.00mm).

Using the same evaluation protocol with absolute distances instead of signed distances, the median distance was 0.927mm (0.86-1.06mm). Median distances for the left oblique fissures were 0.91mm (0.86-1.03mm). The right oblique fissures median distance was 0.92mm (0.86-1.06mm) and the median distance for the right horizontal fissure was 1.06mm (1.03-1.10mm).

Figure 4.6 shows the scatter plot for all cases. Figure 4.7 shows the median distances on the left and the median absolute distances on the right. Figure 4.8 shows the individual results for each fissure and case.




Figure 4.6: Scatter plot including the combined results of all fissures. The results were obtained by computing EMV.




Figure 4.7: Median distances for all fissures. The results were obtained by computing EMV.

As can be seen in Figure 4.7 on the right, all cases that performed poorly (i.e. cases with over 10mm absolute distance) in the previous experiment, achieved significantly better results by including the labelled vessels. In the first and third image from the left in Figure 4.9, the poor segmentation results of case 14 from the previous experiment are shown, while the much stronger results from this experiment are shown in the second and fourth image. Similarly the first and third images from the left in Figure 4.10 show the weak segmentation results for the left lungs of cases 15 and 20 from the previous experiment. The improved results for the same lungs, which were obtained in this experiment, are shown in the second and fourth image. The segmentation result for the right lung of case 24 is shown in Figure 4.11. In contrast to the result obtained in the previous experiment, the boundary between the right upper and right lower lobe is closely following the right oblique fissure. However, the result did not improve for the boundary between the right middle and right lower lobe.

Although not all scans have a median absolute distance down to zero, the results are




Figure 4.8: Top: left oblique fissures; middle: right oblique fissures; bottom: right horizontal fissures; The results were obtained by computing [EMV](#).




Figure 4.9: Segmentation results for case 14 (left and right lung) of the in-house dataset. The results computed by EMA are shown in the first and third images from the left and the results computed by EMV are shown in the second and fourth image from the left.




Figure 4.10: The segmentation results for the left lung of case 15 of the in-house dataset are shown in the first (EMA) and second (EMV) images. The third image shows the result for the left lung of case 20 computed by EMA while the result computed by EMV is shown on the fourth image from the left.

very promising, as only one case has a median absolute distance above 10mm for the left oblique fissure. This case is shown in Figure 4.12 on the left. As can be seen in the image, the lower part of the left oblique fissure is pathologically thick. The pathologically thick part of the fissure got detected by the vessel segmentation algorithm as shown in the second image from the left in Figure 4.12. Unfortunately, the airway segmentation also produced a very weak result, as shown in the third image from the left in Figure 4.12. The combination of the poor airway segmentation and fissures falsely detected as vessels, led to the labelled vessel tree shown in Figure 4.12 on the right. Starting from a weak and partially wrong labelled vessel tree, it is impossible that the energy minimization yields a good result. However, as described above, there have to be several poor performing steps in the process to come to a comparable result.

All things considered, the algorithm showed to produce reliable and stable results and almost all weak performing cases from the previous experiment performed much better.




Figure 4.11: Segmentation results for case 24 (right lung) of the in-house dataset. The result computed by [EMA](#) is shown on the left and the result computed by [EMV](#) is shown on the right.




Figure 4.12: From left to right: 1) Poor segmentation result for the left lung of case 19. The lower part of the left oblique fissure is pathologically thick. 2) Result for the vessel tree reconstruction. The pathologically thick part of the fissure gets falsely detected by the vessel segmentation. 3) Labelled airway tree. 4) labelled vessel tree; The red labelled vessels denote vessels that could not get labelled.

4.5 Evaluation of [EMA](#) on LOLA11 dataset

This experiment was conducted on the publicly available dataset provided by the organizers of the LOLA11 challenge described in Chapter 4.2. The algorithm presented in Chapter 3.6.1 ([EMA](#)) is executed on each scan in the dataset. As the provided dataset contains cases with bad image quality, the airway segmentation described in Chapter 3.1.1 sometimes faced problems to create airway segmentations that include the

	mean	SD	min	Q1	median	Q3	max
left upper lobe	0.929	0.118	0.272	0.922	0.971	0.991	0.997
left lower lobe	0.884	0.231	0.000	0.919	0.965	0.983	0.993
right upper lobe	0.873	0.169	0.000	0.853	0.938	0.977	0.996
right middle lobe	0.714	0.322	0.000	0.575	0.863	0.941	0.994
right lower lobe	0.928	0.104	0.341	0.909	0.973	0.983	0.995
score	0.866						

Table 4.1: LOLA11 results for the lobe segmentation algorithm [EMA](#) described in Chapter 3.6.1.

lobar bronchi. To circumvent this problem, a second airway segmentation is obtained by using the `PartialLungLabelMapImageFilter` from the Chest Imaging Platform [54]. In cases it still did not reach the lobar bronchi, the airway segmentation got manually refined. Table 4.1 shows the evaluation results obtained from the LOLA11 organizers. It includes the results for each lobe independently and an overall score (see Chapter 4.2) including all lobes. In this experiment a score of 0.866 was achieved, which comes close to the current state of the art with a score of 0.884.




Figure 4.13: LOLA11 lobe segmentation results computed with [EMA](#). From left to right: Case 01 left, case 13 right, case 42 left (includes focal regions of emphysema), case 42 right (the arrow indicates the invisible fissure).

Figure 4.13 shows four examples for good segmentation results computed by [EMA](#). In the first image from the left, a sagittal slice of the left lung of LOLA11 case 01 is shown. As can be seen in the image, the computed boundary is precisely following the fissure. In the second image from the left, a result of similar precision of the right lung of case 13 is shown. The third image from the left shows a slice of the left lung of case 42, which is accurately segmented although the lung includes focal regions of emphysema. The segmentation result for the right lung of case 42 is shown on the fourth image from the left. The segmentation result follows closely the fissures in regions they are present. Although the upper part of the oblique fissure is not visible (indicated by the red arrow),

	mean	SD	min	Q1	median	Q3	max
left upper lobe	0.940	0.127	0.178	0.955	0.98	0.992	0.997
left lower lobe	0.896	0.231	0.000	0.942	0.973	0.986	0.995
right upper lobe	0.898	0.158	0.000	0.885	0.944	0.981	0.997
right middle lobe	0.819	0.222	0.000	0.779	0.897	0.951	0.994
right lower lobe	0.953	0.076	0.489	0.940	0.975	0.987	0.996
score	0.901						

Table 4.2: Lung lobe segmentation results for the 55 scans in LOLA11 computed with [EMV](#).

the algorithm produces a meaningful boundary between the right upper and right lower lobe.

Even though the algorithm yields convincing results in some challenging cases like shown in [Figure 4.13](#), it still faces problems in others. The first and third image from the left in [Figure 4.15](#) show results for the right lung of case 11 and left lung of case 06 respectively. In the first image, the boundary between the right upper lobe and right middle lobe does not follow the right horizontal fissure. In the third image from the left, the algorithm is not even closely finding the right boundary. The case represents a pathologically deformed lung with focal regions of emphysema and a drastically narrowed left upper lobe. Another poor segmentation result is obtained for case 07, which is shown in the first (left lung) and third (right lung) image of left in [Figure 4.16](#). Due to bad image quality and a lot of noise, the algorithm produces unsatisfying results.

4.6 Evaluation of [EMV](#) on LOLA11 dataset

In this experiment, the protocol of [Chapter 4.5](#) is repeated with the refined version of the algorithm ([EMV](#)) described in [Chapter 3.6.2](#). The hypothesis that better results can be achieved by using the labelled vessels segmentation for computing a better approximate lobe segmentation, has shown to be true. In fact the current highest score of 0.884 has been exceeded with a score of 0.901 as can be seen in [Table 4.2](#) and [Table 4.3](#).

[Figure 4.14](#) shows the segmentation results for case 29 (first and second image from the left) and case 36 (third and fourth image from the left) computed with [EMV](#). The detected boundaries are precisely following the fissures in all four images. A meaningful boundary is even found in the second image from the left between the upper and middle lobe although the horizontal fissure is completely missing.

More challenging cases that performed poorly in the previous experiment (computation with [EMA](#)) are shown in [Figure 4.15](#). In the second image from the left, the right lung of LOLA11 case 11 is shown. In contrast to the first image from the left, which was obtained with [EMA](#), the boundary between the right upper and middle lobe is precisely following the fissure. The improvement of the left lung of case 06, which is




Figure 4.14: LOLA11 segmentation results computed with [EMV](#). From left to right: Case 29 left, 29 right, 36 left, 36 right.




Figure 4.15: LOLA11 segmentation results. From left to right: Case 01 right ([EMA](#)), case 01 right ([EMV](#)), case 06 left ([EMA](#)), case 06 left ([EMV](#)).

shown in the fourth image from the left is impressive. While the algorithm ([EMA](#)) used in the previous experiment has produced a boundary not even close to the fissure, the algorithm incorporating the labelled vessel trees ([EMV](#)) yields an almost perfect result. This segmentation result is all the more impressive as the lung is significantly deformed.

Another example for a segmentation result that significantly improved by using [EMV](#) instead of [EMA](#), is shown in [Figure 4.16](#). The improved results for the left and right lung are shown in the second and fourth image from the left respectively. As can be seen, using the labelled vasculature adds enough information to the algorithm for finding a very precise solution even in case of noisy scans.

[Figure 4.17](#) shows two exceptional cases. The first and second image from the left show the left and right lung of case 05 respectively. The algorithm ([EMV](#)) has no problems finding the right segmentation, even though the lung is anatomically deformed due to a very strong scoliosis. The third image from the left shows the left lung of case 21. As opposed to the manually created reference standard, in which no boundary between the




Figure 4.16: LOLA11 segmentation results. From left to right: Case 07 left (EMA), case 07 left (EMV), case 07 right (EMA), case 07 right (EMV).




Figure 4.17: LOLA11 segmentation results computed with EMV. From left to right: Case 05 left, case 05 right, case 21 left, case 21 right.

upper and lower lobe was drawn for this case [37], the algorithm (EMV) finds a plausible delineation. The fourth image from the left shows the segmentation result for the right lung of case 21. No lobar border was drawn for the right horizontal fissure in the manually created reference standard for this case [37]. As apparent in the image, only one fissure is visible in this image and the algorithm perfectly separates the lung into two lobes with the fissure as boundary.

4.7 Discussion

The proposed algorithm for fully automatic lung lobe segmentation that was developed in this thesis has shown good results in the experiments presented in this chapter. The algorithm not only got evaluated on an in-house dataset of 25 CT images but also on a publicly available dataset of 55 CT images. In Chapter 4.3 the algorithm presented in Chapter 3.6.1 was evaluated on the in-house dataset presented in Chapter 4.1 and achieved good results. The same algorithm was evaluated on the publicly available dataset described in Chapter 4.2 and achieved a score of 0.866. With this score the

Method	Description	LOLA11
[65]	fissures & airways & multiatlas	0.851
[37]	cost image (airways, vessels, fissures) & watershed	0.881
[6]	probabilistic fissure segmentation & groupwise prior	0.884
	approximation (labelled airways) & EM (airways, fissures, vessels)	0.866
	approximation (labelled vessels) & EM (airways, fissures, vessels)	0.901

Table 4.3: Updated LOLA11 results for lung lobe segmentation including the scores of the algorithms presented in this thesis.

algorithm performs in the range of the other state-of-the-art methods that were evaluated by the LOLA11 challenge. The interactive approaches still show distinctly better results [41], which come however, with the drawback of not being fully-automatic and needing manual intervention during the segmentation process. By refining the algorithm and using the lobe-based vasculature for computing a better approximate lobe segmentation as described in Chapter 3.6.2, an even higher score of 0.901 (see Table 4.3) was achieved at the LOLA11 challenge dataset. With this score the previously best result of 0.884 got exceeded, which makes the EMV algorithm the best reported fully-automatic method. The positive impact of using the labelled vasculature for finding a better approximation on the resulting lobe segmentation was also shown by the experiment presented in Chapter 4.4.

As no individual result for each case in the LOLA11 dataset is provided, it is difficult to give a detailed analysis of the algorithm. However, visual inspection has shown that the approach of labelling the vasculature and using it to compute an approximation has significantly improved a lot of cases, in which the airway segmentation was too weak to yield a good approximation. This adheres to the expectations, as it is normally easier to achieve a detailed vessel segmentation than an airway segmentation. On the other hand it is crucial that the labelling of the vessels is not erroneous. Visual inspection has shown that following the algorithm presented in Chapter 3.3.2 led to good results. However, there are still cases, in which errors occurred. Figure 4.18 in the first image from the left shows a labelled vessel tree, which contains a wrong labelled vessel in the middle lobe. By looking at the tree reconstruction result shown in Figure 4.18 in the second image from the right it is apparent that the error comes from the tree reconstruction step, in which two vessels from different lobes are reconstructed into one tree. The third and fourth image from the left in Figure 4.18 show the approximate lobe segmentation computed on the labelled vessel tree. This kind of error was observed a few times, which makes it worth the effort to correct it in the future to achieve even better results.




Figure 4.18: From left to right: 1) Labelled vessel trees containing a wrong labelled vessel in the middle lobe. 2) Tree reconstruction result. The wrong labelled vessel is shown in red. 3) Approximate lobe segmentation computed on the labelled vessel trees. 4) A sagittal slice of the approximate lobe segmentation.

By comparing Table 4.1 and Table 4.2 it is apparent that improvements were mostly achieved in the first quantile (Q1). This means that the approach of including the labelled vessel tree primarily led to better results in cases that previously performed poorly. The increase of the score for the right middle lobe is also remarkable. The middle lobe is often the most difficult lobe to segment, as it is normally the smallest lobe in terms of volume and fissures are often missing at its boundaries. Adding the labelled vessel tree is providing a strong foundation and can therefore greatly help for segmenting the middle lobe, especially in cases where it is small and fissures are missing.

In conclusion the results show that the approach of minimizing the developed energy equation yields very good results. By first computing an approximate lobe segmentation, the minimization of the energy equation becomes a refinement step. The advantage of this lies in the fact that the anatomical structures used in the energy equation do not need to be absolutely uncorrupted. This provides the possibility to use a multi-scale fissure segmentation approach as described in Chapter 3.4, which is able to detect big parts of pathological thick fissures in conjunction with falsely detected structures. The approach of using the labelled vasculature for finding a better approximation has significantly helped in increasing the performance of the algorithm, which has finally led to an algorithm with the currently highest score reported on the LOLA11 challenge dataset.

Lobe based A/V-Separation Evaluation

Contents

5.1 Lobe wise A/V separation	60
5.2 Using Fissures for A/V separation	60

In this chapter the question whether incorporating information on the pulmonary lobes can help to improve the A/V labelling is evaluated. On that point, the performance of the A/V separation algorithm [44] is first evaluated by computing it on the whole lung. The dataset described in Chapter 4.1 is used to evaluate the performance. A manual segmentation and A/V labelling has been created for each scan in the dataset. The overlap between the automatically generated A/V labelling and the manually defined reference is computed using following formula:

$$\text{agreement} = \frac{|A_{ref} \cap A_{test}| + |V_{ref} \cap V_{test}|}{|(A_{ref} \cup V_{ref}) \cap (A_{test} \cup V_{test})|} \quad (5.1)$$

Figure 5.1 shows three A/V labelling evaluation examples. Arteries are visualised in blue, veins in red and wrong labelled vessels in green. The left lungs are displayed in brown while the right lungs are visualised in grey. The first image from left in Figure 5.1 shows the results for case 10. As can be seen in the image, the right middle lobe is completely mislabelled. The second image and third image show case 17 and case 25 respectively. In both cases the left lower lobe is completely mislabelled.

In order to evaluate the benefit of incorporating information on lung lobes into the A/V separation algorithm, two different experiments are performed. In Chapter 5.1 the A/V separation is carried out in each lobe separately. In Chapter 5.2 information of the identified fissures is included into the A/V separation.




Figure 5.1: A/V labelling evaluation results. Arteries are shown in red while veins are shown in blue. Wrong labelled vessels are coloured in green. From left to right: Case 10, case 17 and case 25.

5.1 Lobe wise A/V separation

In this experiment the A/V separation is recomputed separately for each lung lobe on the same dataset. The lobe segmentation results obtained from the experiment in Chapter 4.4 are used to this end. Table 5.1 shows the results for this experiment. The table contains the A/V separation results computed on the whole lung (Lung A/V) and the A/V separation results computed independently for each lobe (Lobe A/V). The arrows are coloured if the absolute difference between the two values corresponding to the two computations is greater than 0.03. No results for the A/V separations are available for case 3, as the lobe segmentation algorithm only detected two instead of three lobes. As can be seen in Table 5.1, computing the A/V separation independently for each lobe does not increase the accuracy but instead slightly decreases it. Only two cases that previously performed poorly achieved significantly better results. Case 21 and 24 increased by approximately 17% and 10% respectively. On the other hand, the performance of a lot of other cases decreased. This observation is also visualised in Figure 5.2. The A/V labelling evaluation result computed on the whole lungs are shown on the left in Figure 5.3. The result obtained by computing the A/V separation independently on each lobe is visualised on the right in Figure 5.3.

5.2 Using Fissures for A/V separation

In this experiment the information on lung lobes is incorporated in a different way. Instead of computing the A/V separation independently for each lobe, the lung mask that is used for the computation of the A/V separation is adjusted by removing the fissures from it. As the multi-scale fissure segmentation algorithm described in Chapter 3.4 produces results

Case	Lung A/V	Lobe A/V	Difference	Trend
1	0.967	0.9639	-0.0031	↓
2	0.9546	0.9592	0.0046	↑
3	0.6556	-	-	-
4	0.8438	0.7806	-0.0633	↓
5	0.9849	0.9615	-0.0234	↓
6	0.9862	0.9858	-0.0004	↓
7	0.9943	0.9885	-0.0058	↓
8	0.9887	0.977	-0.0117	↓
9	0.9215	0.932	0.0105	↑
10	0.8015	0.8078	0.0063	↑
11	0.5652	0.4652	-0.1	↓
12	0.963	0.9069	-0.0561	↓
13	0.9834	0.9816	-0.0017	↓
14	0.8977	0.9139	0.0162	↑
15	0.9852	0.9841	-0.0012	↓
16	0.6908	0.7153	0.0245	↑
17	0.7102	0.6846	-0.0255	↓
18	0.7958	0.6987	-0.0971	↓
19	0.9669	0.8917	-0.0751	↓
20	0.9198	0.9269	0.0071	↑
21	0.7609	0.9336	0.1727	↑
22	0.9836	0.9352	-0.0484	↓
23	0.9692	0.9699	0.0007	↑
24	0.332	0.4291	0.0971	↑
25	0.7903	0.7178	-0.0724	↓
Overall mean	0.8565	0.8204	-0.0361	↓

Table 5.1: Evaluation results for the A/V separation algorithm computed on the whole lung (Lung A/V) and computed independently on each lobe (Lobe A/V). The fourth column indicates whether the performance increases (↑) or decreases(↓). The arrows are coloured if the absolute performance difference is greater than 0.03.

that include a lot of false positives, the true fissures need first to be detected. This is done by converting the lung lobe segmentation result into pulmonary lobe boundaries. The multi-scale fissure segmentation result is subsequently filtered by only keeping voxels that are part of the previously computed boundary. The process of the construction of the lung masks that are needed in this experiment, is visualised in Figure 5.4. The idea behind this approach is that computing the A/V separation completely independently on each lobe may remove valuable information. By only removing the fissures from the mask and hence introducing partial boundaries into the lung volume, the loss of information is less extreme. Vessels from different lobes can still influence each other in terms of A/V labelling in regions where no fissures are present.

The results obtained in this experiment are shown in Table 5.2. Although better




Figure 5.2: Left: Bland Altman diagram. X-axis: average % overlap of the lobe wise A/V separation and the A/V separation computed on the whole lung. Y-axis: difference in the overlap percentage between the A/V separation results computed on the whole lung and the lobes. Right: The change in overlap % for each case is visualised.




Figure 5.3: A/V labelling evaluation results for case 21. Arteries are shown in red while veins are shown in blue. Wrong labelled vessels are coloured in green. Left: A/V separation computed separately on the whole left and whole right lung. The left upper lobe is almost completely mislabelled. Right: A/V separation result computed independently on each lung lobe.

results were achieved using the approach of removing the fissures from the lung mask compared to computing the A/V separation independently for each lobe, the results are still disappointing. As can be seen in Table 5.2, the overall mean only increased by 0.26 percentage points, which is insubstantial. By comparing Table 5.1 with Table 5.2 it is apparent that in this experiment far less cases experienced a performance loss compared to the previous experiment. However, case 21, which improved by 21 percentage points for the lobe wise computation, did not improve at all in this experiment. In Figure 5.5




Figure 5.4: From left to right: 1) Lobe segmentation result as overlay on the original CT scan. 2) Lobe segmentation converted to boundary and displayed as overlay on original CT scan. 3) Fissure segmentation result including a lot of false positives as overlay on the original CT scan. 4) Final lung mask, in which true fissures are removed.

it is clearly visible that in contrast to the previous experiment the performance difference for each individual case is smaller. This comes up to one's expectations, as the approach of only removing the fissures from the lung mask is introducing less change compared to computing the A/V separation independently on each lobe.

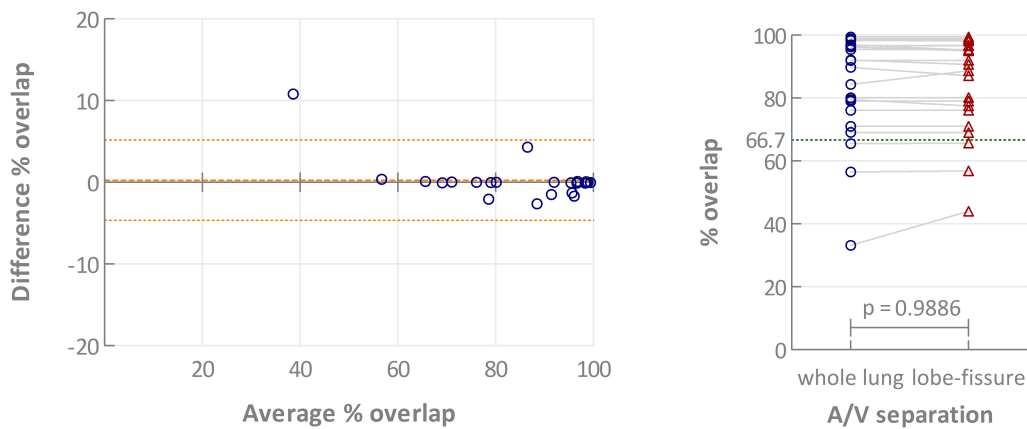


Figure 5.5: Left: Bland Altman diagram. X-axis: average % overlap of the A/V separation results computed on the adapted lung mask (removed fissures) and the A/V separation computed on the whole lung. Y-axis: difference in the overlap percentage between the A/V separation results computed on the whole lung and computed on the adapted lung mask. Right: The change in overlap % for each case is visualised.

Case	Lung A/V	Removed Fissures A/V	Difference	Trend
1	0.967	0.9662	-0.0007	↓
2	0.9546	0.9542	-0.0004	↓
3	0.6556	0.6568	0.0012	↑
4	0.8438	0.8871	0.0433	↑
5	0.9849	0.9858	0.0008	↑
6	0.9862	0.9862	0	↓
7	0.9943	0.9943	0	↓
8	0.9887	0.9889	0.0002	↑
9	0.9215	0.9067	-0.0149	↓
10	0.8015	0.8017	0.0002	↑
11	0.5652	0.5691	0.0039	↑
12	0.963	0.9502	-0.0129	↓
13	0.9834	0.9825	-0.0009	↓
14	0.8977	0.8716	-0.0261	↓
15	0.9852	0.9853	0.0001	↑
16	0.6908	0.6905	-0.0004	↓
17	0.7102	0.7107	0.0005	↑
18	0.7958	0.7755	-0.0204	↓
19	0.9669	0.968	0.0011	↑
20	0.9198	0.9201	0.0003	↑
21	0.7609	0.761	0.0002	↑
22	0.9836	0.9836	0	↓
23	0.9692	0.9525	-0.0167	↓
24	0.332	0.44	0.108	↑
25	0.7903	0.7902	0	↓
	0.8565	0.8591	0.0026	↑

Table 5.2: Evaluation results for the A/V separation algorithm computed on the whole lung (Lung A/V) and computed on the adapted lung mask (Removed Fissures A/V). The fourth column indicates whether the performance increases (↑) or decreases(↓). The arrows are coloured if the absolute performance difference is greater than 0.03.

Quantitative Readouts

Contents

6.1 Definition of Readouts	65
6.2 Results	66
6.3 Discussion	74

To demonstrate the clinical applicability of pulmonary lobe segmentation, several quantitative readouts are computed on the pulmonary vessels and evaluated for each lobe. The dataset described in Chapter 4.1 is used for that purpose. Moreover, the computed quantitative readouts are correlated to the pulmonary pressure of the subjects to identify possible indicators for PH. In Chapter 6.1 the quantitative readouts are defined. The results are shown in Chapter 6.2 and subsequently discussed in Chapter 6.3.

6.1 Definition of Readouts

Several quantitative readouts including the number of vessel segments, Sum of Angles Metric (SOAM), vessel density and normalized vessel volume are computed.

6.1.1 Vessel Segment

A vessel may consist of several vessel segments. A vessel segment is a part of a vessel that is located between either two branching or end points. From this it follows that two new segments are formed at each bifurcation.

6.1.2 Sum of Angles Metric

The SOAM is a metric for quantifying the tortuosity. It is computed by first defining a series of n points \mathbf{u}_1 to \mathbf{u}_n at an interval of 4mm on a vessel segment. This vessel path further contains $n - 1$ normalized direction vectors \mathbf{v}_1 to \mathbf{v}_{n-1} . The SOAM is then

computed by summing up the angles of all consecutive directional vectors \mathbf{v}_i as defined in (6.1).

$$\text{SOAM}(p) = \frac{\sum_{i=1}^{n-2} \cos^{-1}(\mathbf{v}_i \cdot \mathbf{v}_{i+1})}{\sum_{i=1}^{n-1} |\mathbf{u}_i - \mathbf{u}_{i+1}|} \quad (6.1)$$

with

$$\mathbf{v}_i = \frac{\mathbf{u}_{i+1} - \mathbf{u}_i}{|\mathbf{u}_{i+1} - \mathbf{u}_i|} \quad (6.2)$$

6.1.3 Vessel Density

The vessel density is defined as

$$V_d = \frac{N}{V_L} \quad (6.3)$$

where N is the number of vessel segments and V_L is the lung or lobe volume.

6.1.4 Normalized Vessel Volume

The normalized vessel volume is defined as

$$V_N = \frac{V_V}{V_L} \quad (6.4)$$

where V_V is the volume of all vessels in the lung or lobe with volume V_L .

6.2 Results

Figure 6.1 shows the results for vessel density, SOAM and normalized vessel volume on the top left, top right and bottom respectively. The readouts are shown on the y-axes while the individual results for all five lung lobes (LU, LL, RU, RM, RL) are shown on the x-axes.

Correlations between PH and SOAM, PH and vessel density, PH and normalized vessel volume are shown in Figure 6.2, Figure 6.4 and Figure 6.6 respectively. The figures contain results computed on the whole lung and on each lobe separately.

The Receiver-Operating-Characteristic (ROC) curves for SOAM, vessel density and normalized vessel volume as predictors for PH are shown in Figure 6.3, Figure 6.5 and Figure 6.7 respectively.

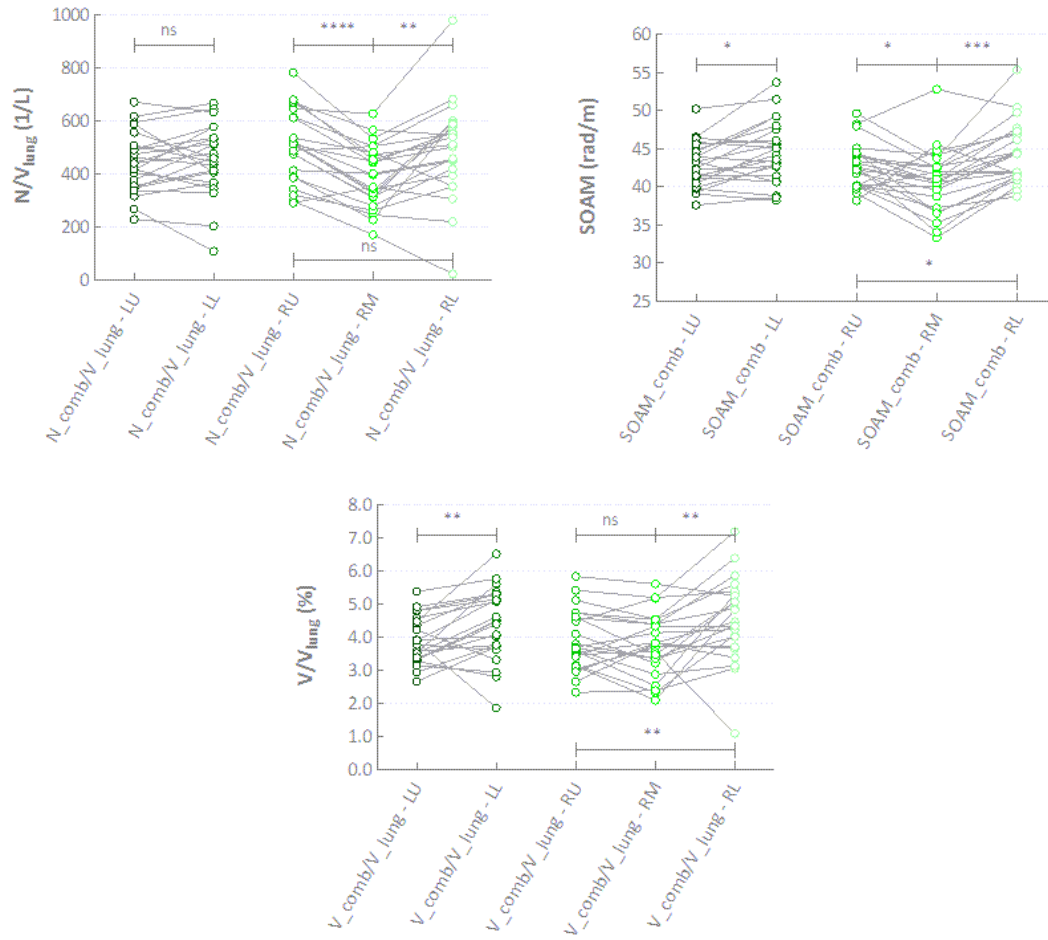


Figure 6.1: Top left: Vessel density for all five lung lobes. Top right: SOAM for all five lung lobes. Bottom: Normalized vessel volume for all five lung lobes. LU: left upper lobe; LL: left lower lobe; RU: right upper lobe; RM: right middle lobe; RL: right lower lobe.

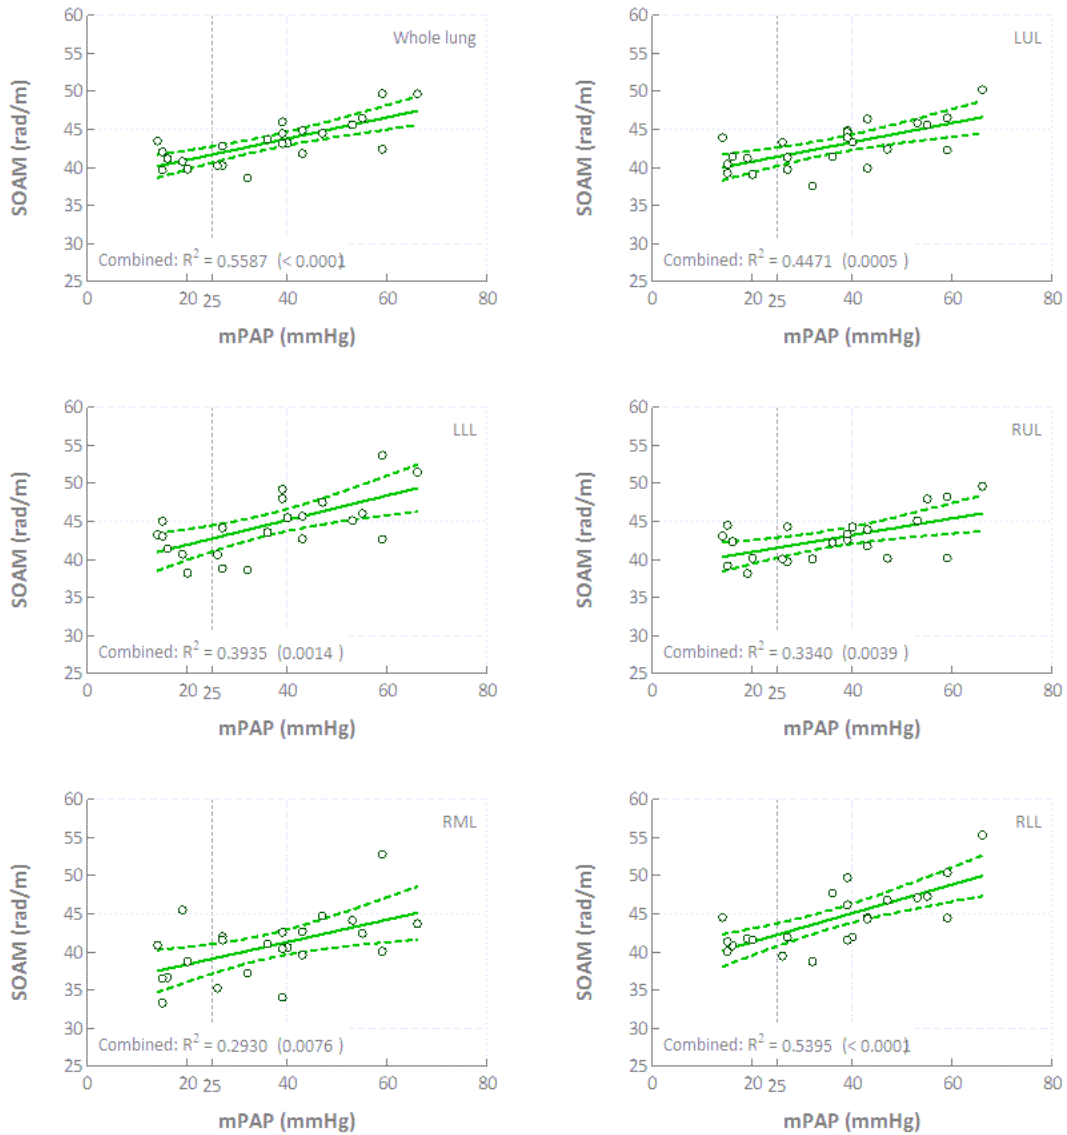


Figure 6.2: From left to right and top to bottom: Correlation between mPAP and SOAM for the whole lung, left upper lobe, left lower lobe, right upper lobe, right middle lobe and right lower lobe.

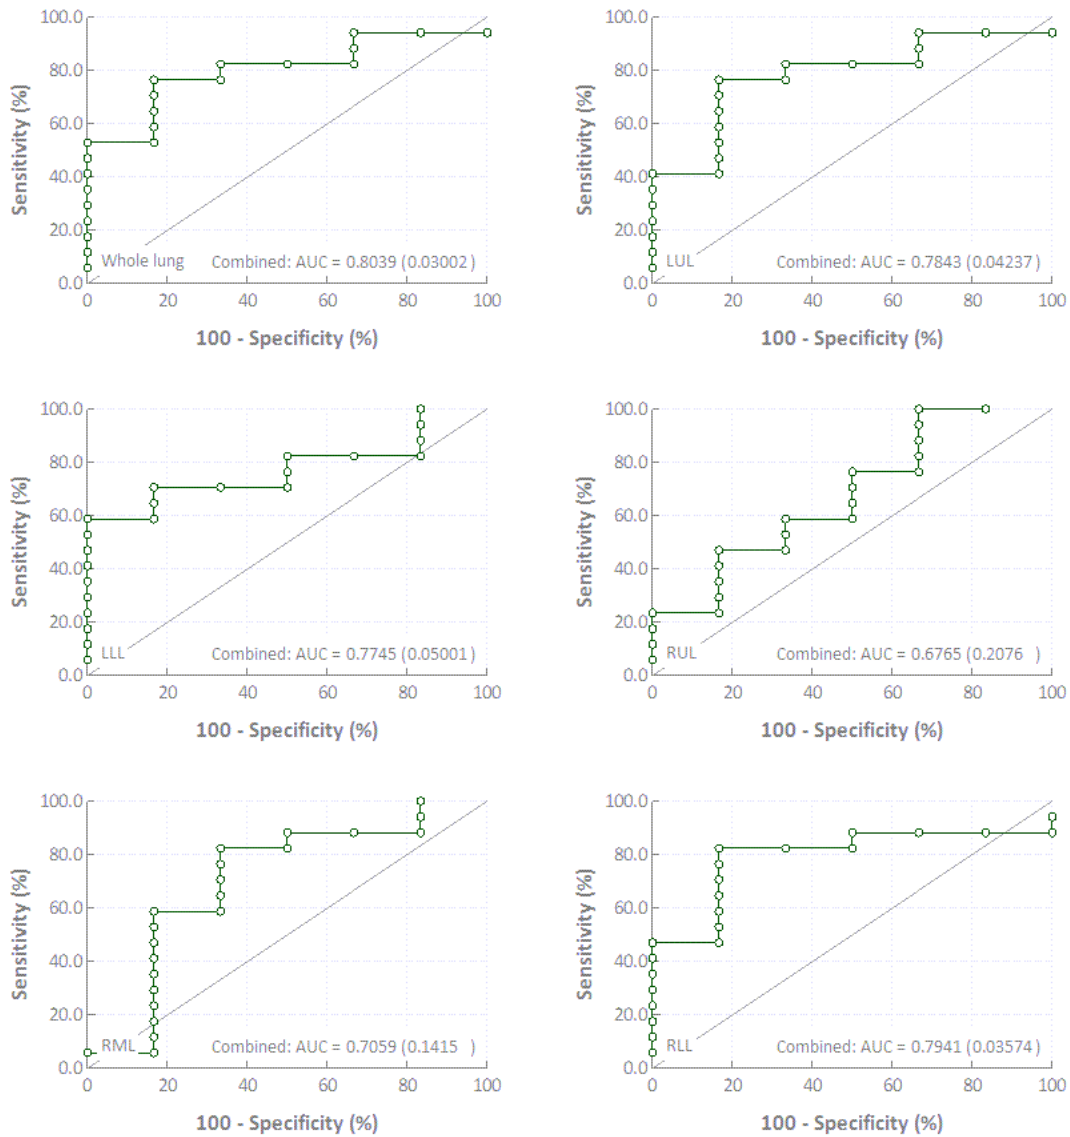


Figure 6.3: From left to right and top to bottom: ROC curve for SOAM as predictor for PH computed on the whole lung, left upper lobe, left lower lobe, right upper lobe, right middle lobe and right lower lobe.

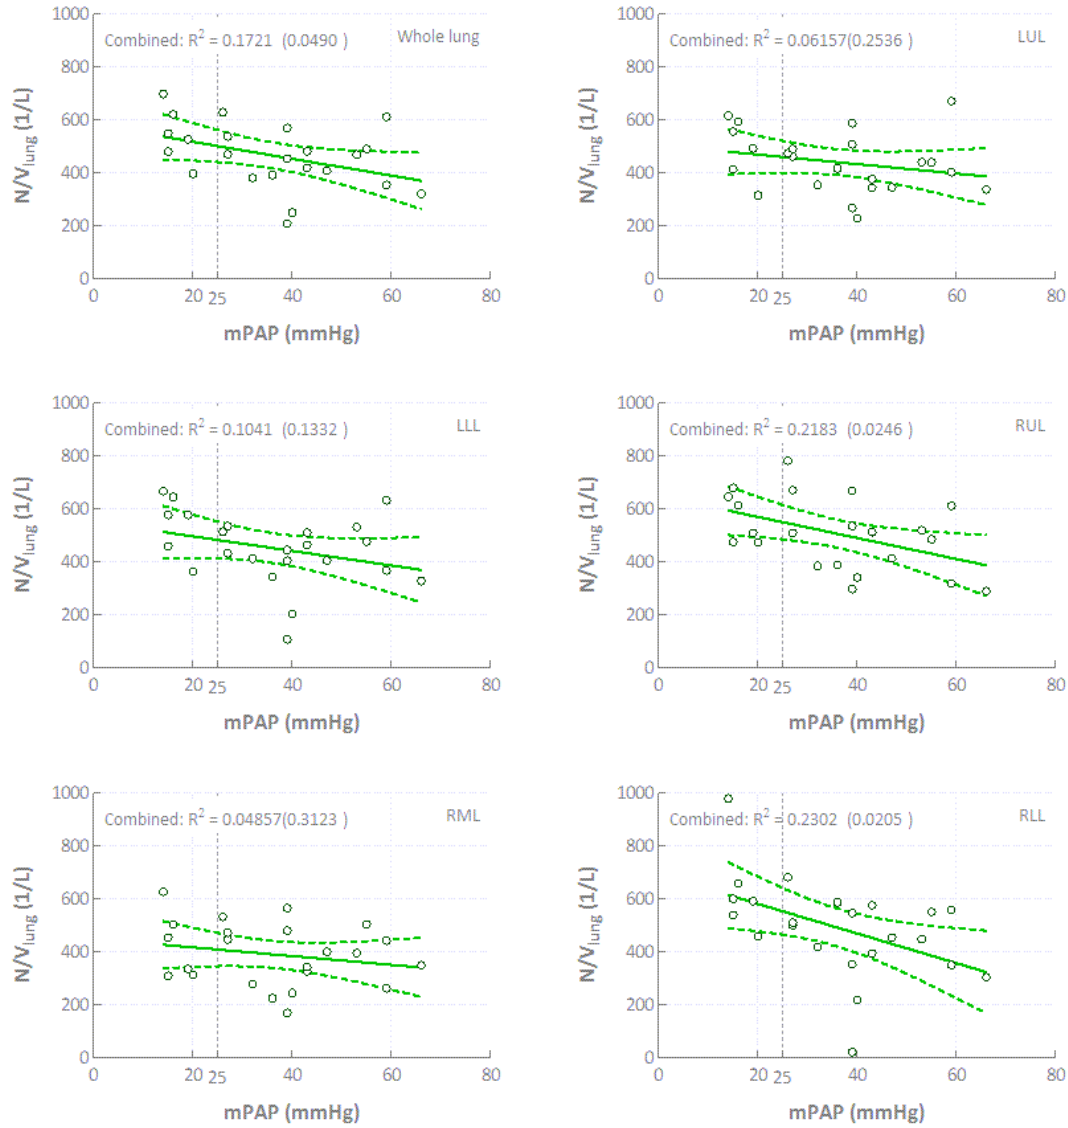


Figure 6.4: From left to right and top to bottom: Correlation between mPAP and vessel density for the whole lung, left upper lobe, left lower lobe, right upper lobe, right middle lobe and right lower lobe.

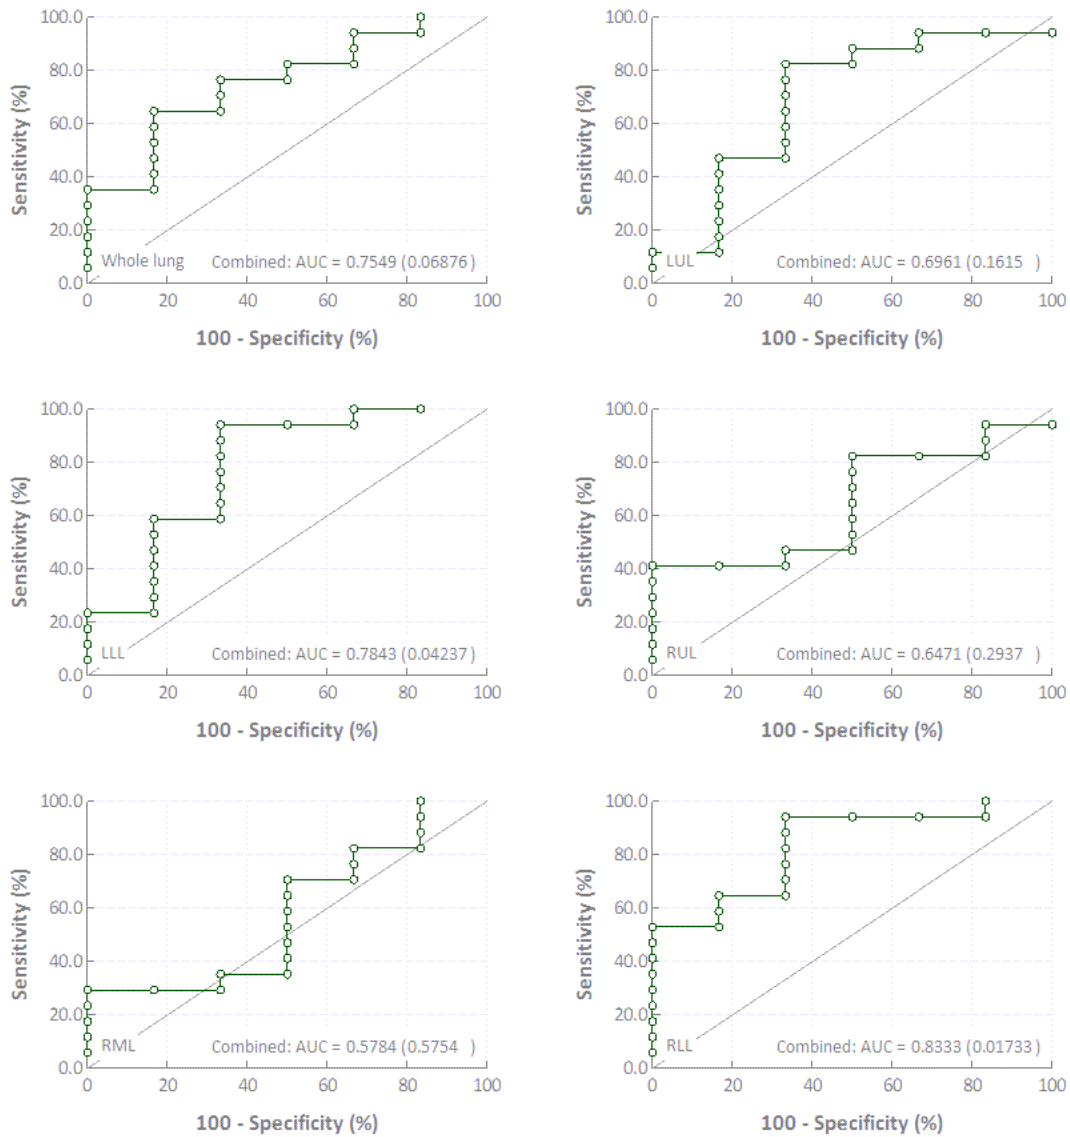


Figure 6.5: From left to right and top to bottom: ROC curve for vessel density as predictor for PH computed on the whole lung, left upper lobe, left lower lobe, right upper lobe, right middle lobe and right lower lobe.

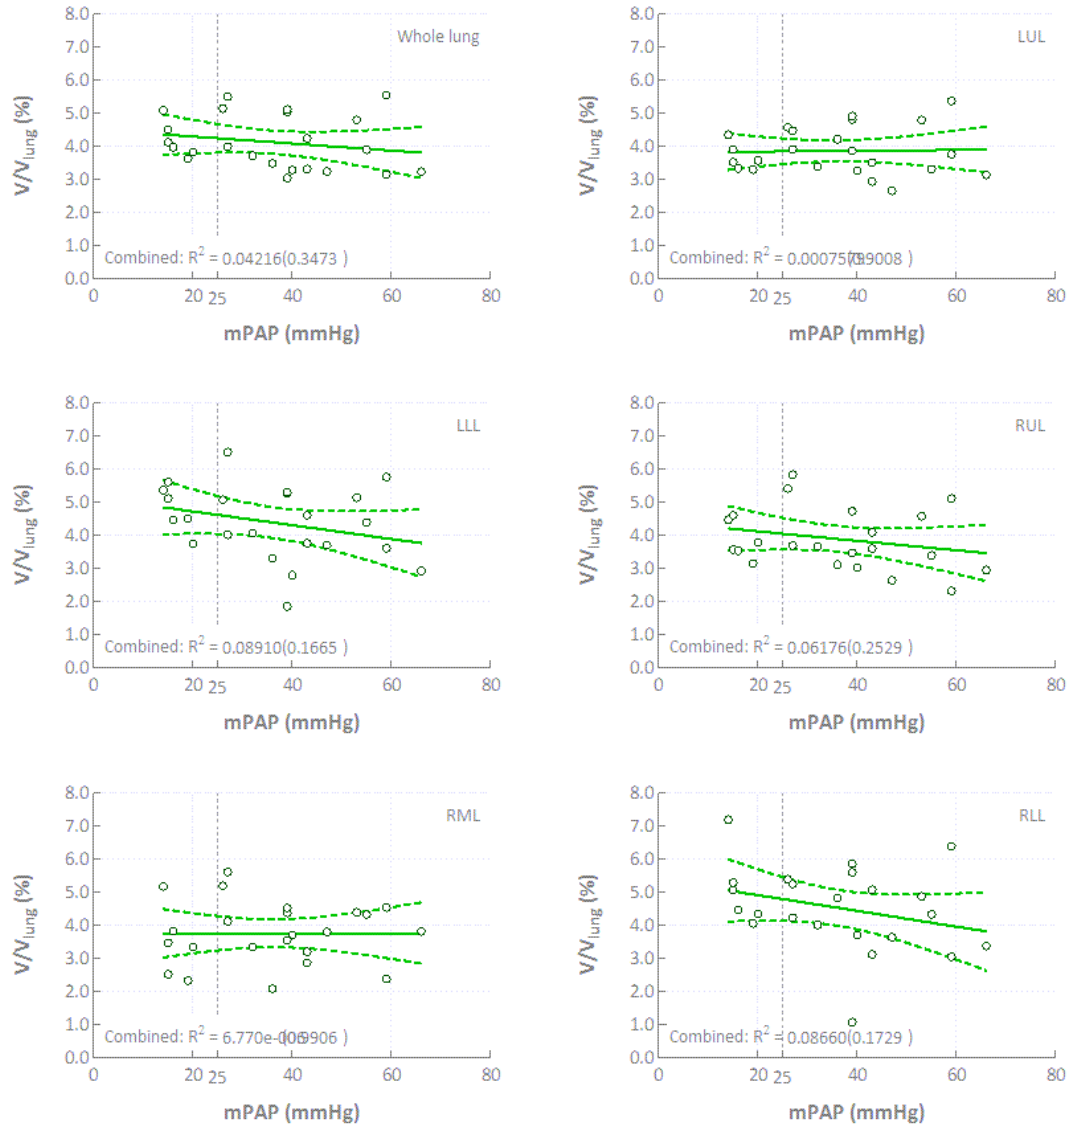


Figure 6.6: From left to right and top to bottom: Correlation between mPAP and normalized vessel volume for the whole lung, left upper lobe, left lower lobe, right upper lobe, right middle lobe and right lower lobe.

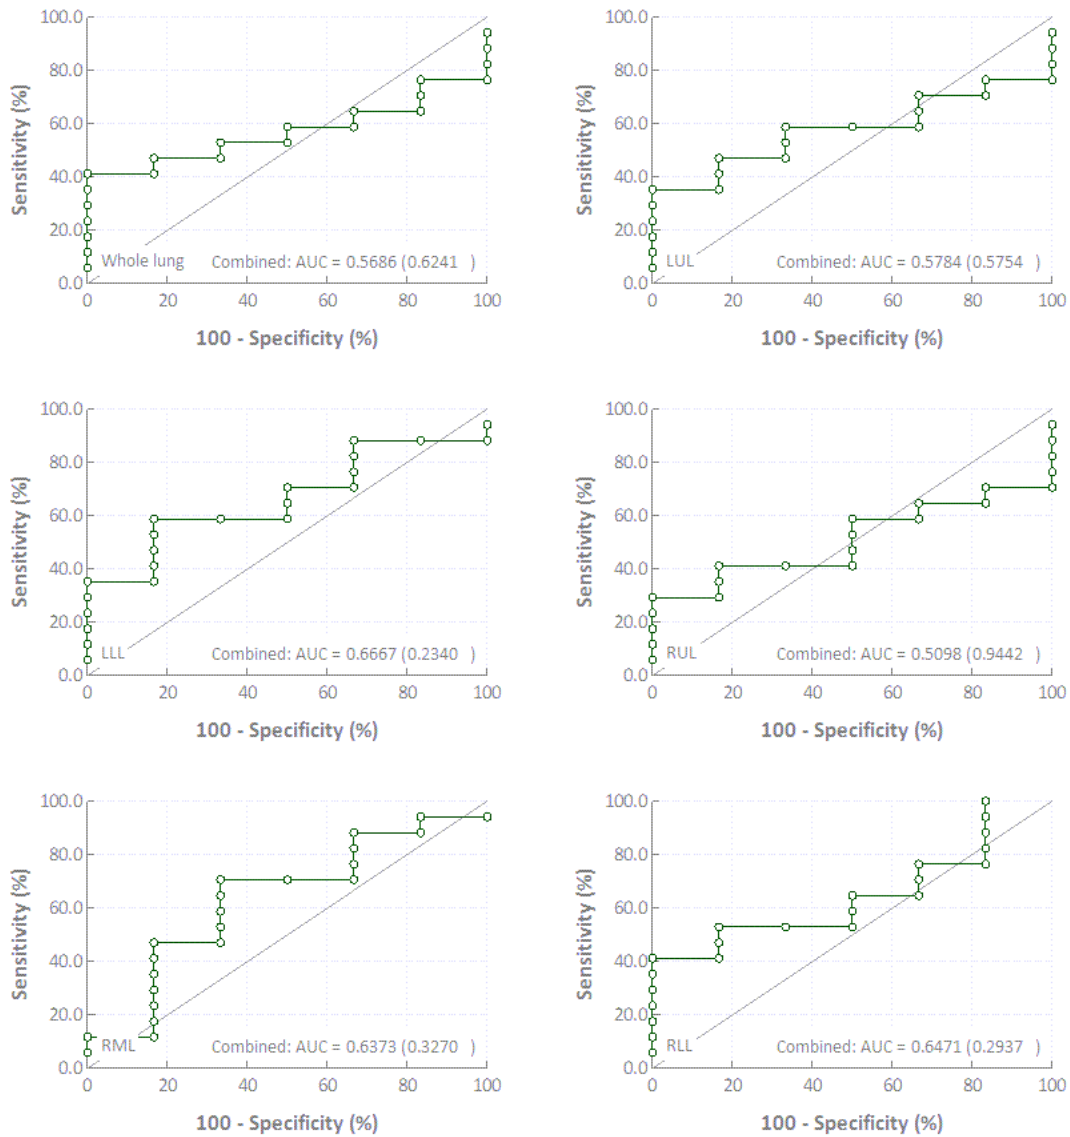


Figure 6.7: From left to right and top to bottom: ROC curve for normalized vessel volume as predictor for PH computed on the whole lung, left upper lobe, left lower lobe, right upper lobe, right middle lobe and right lower lobe.

6.3 Discussion

Several readouts, including the vessel density, **SOAM** and normalized vessel volume were computed on the vasculature of the whole lung and of each lobe independently. Figure 6.1 shows the results for the vessel density, **SOAM** and normalized vessel volume computed on each lobe. As can be seen in the figure, the vessel density is significantly lower in the right middle lobe compared to the right upper and right lower lobe. However, it is difficult to find a meaningful explanation for this. The **SOAM** is increased in the lower lobes compared to the upper lobes. An explanation for this could be the location of the lobes and with this the constant higher pressure that is exerted on the lower lobes. However, following this explanation, the right middle lobe should have a higher **SOAM** compared to the right upper lobe, which is not the case in the presented dataset. When it comes to normalized vessel volume, the data shows a similar pattern. The normalized vessel volume in the lower lobes is increased compared to the normalized vessel volume in the upper lobes. Since there is no significant difference in vessel density between the upper and lower lobes, the increased normalized vessel density in the lower lobes could be explained by thicker vessels. Thicker vessels, on the other hand, can again be explained as a response to higher pressure.

Correlations between the readouts and pulmonary pressure of the subjects were computed for the whole lung and for the lung lobes. It is apparent from Figure 6.2 that there is a correlation between the **SOAM** and pulmonary pressure. The **ROC** curves presented in Figure 6.3 show that the **SOAM** is a potential predictor for pulmonary hypertension. The evaluation of the **SOAM** on the whole lung seems to be best for predicting pulmonary hypertension i.e. the information on hypertension is not concentrated in a single lobe.

Normalized vessel volume shows no significant correlation to pulmonary pressure as shown in Figure 6.6 and is also not useful for predicting pulmonary hypertension as depicted in Figure 6.7.

Vessel density seems to be slightly correlated to pulmonary pressure as can be seen in Figure 6.4. In the used dataset, the lower lobes appear being potentially useful for predicting pulmonary hypertension as can be seen in Figure 6.5.

The low number of subjects in the dataset is seen as a limitation of the results. It is therefore difficult to give a precise explanation of the generated readouts, as noise may be interpreted as patterns. Further analyses with larger datasets have to be performed in future studies to be able to draw more reliable conclusions. However, a step into a promising direction has been taken with this analysis and the clinical usefulness of fully automatic pulmonary lobe segmentation has successfully been demonstrated.

Conclusion

In this thesis a fully-automatic algorithm for pulmonary lobe segmentation has been developed. To ensure clinical applicability, two variations of the algorithm have been evaluated on two different datasets that contain a lot of varying cases and pathologies. A score, representing the mean overlap between the automatic and reference segmentations for all 55 cases, of 86.6% and 90.1% was reached, the latter of which is currently the highest score that has been achieved at the LOLA11 challenge.

Using the improved lobe segmentation algorithm, the question whether information on lung lobes can improve an existing A/V separation algorithm [44] was evaluated. In a first attempt the A/V separation was computed independently on each lobe. The second approach was to compute the A/V separation on the whole lung with adapted lung masks, in which the pulmonary fissures were removed. However, the performance of the A/V separation was not increased in either approach and hence the hypothesis that pulmonary lobe segmentation may improve the performance of A/V separation could not be confirmed.

The clinical applicability of the algorithms was shown by computing clinical readouts on the vessel segmentation and evaluating the results for each lobe. An increased SOAM has shown to be a potential predictor for pulmonary hypertension. The results have further suggested that an increased vessel density in the lower lung lobes may be an indicator for pulmonary hypertension.

The approach of first computing an approximate lobe segmentation and then minimizing an energy equation that contains information on auxiliary lung structures, has shown promising results. Analysing the results has shown that poor lobe segmentations were almost always attributable to poor segmentation results of auxiliary structures. Labelling the pulmonary vasculature with respect to the lobes and using it for lobe segmentation has shown to be a promising new approach, as suggested by the improved LOLA11 result of 90.1%.



List of Acronyms

- BLVR** bronchoscopic lung volume reduction
- CT** computed tomography
- CV** collateral ventilation
- DFS** depth first search
- EMA** Energy Minimization with Airways as approximation
- EMV** Energy Minimization with Vessels as approximation
- mPAP** Mean Pulmonary Artery Pressure
- OOF** Optimally-Oriented Flux
- PH** Pulmonary Hypertension
- ROC** Receiver-Operating-Characteristic
- ROI** Region Of Interest
- SOAM** Sum of Angles Metric



List of Publications

B.1 2017

Fully-automatic Lung Fissure Detection from Thoracic Computed Tomography

Nicola Giuliani, Martin Urschler, Andrea Olschewski, Horst Olschewski and Michael Pienn
In: *Wiener klinische Wochenschrift*, 129(19):743-780
Oktober 2017, Innsbruck, AT
(Accepted for poster presentation)

Abstract: BACKGROUND: Fully-automatic lung fissure detection from computed tomography (CT) images can aid the planning of valve implantation in COPD patients. METHODS: Lung lobes in thoracic CT images were segmented without user intervention. Lobar bronchi were labelled to initialize the lobe segmentation. Fissures were detected based on an eigenvalue analysis of the Hessian-matrix. In regions where no fissure was detected, the distance to lung vessels indicates regions with high probabilities for lobe boundaries. Minimization of an energy equation containing all above information yielded optimal lobe boundaries. For the evaluation of the algorithm two non-biased individuals placed 500 markers along the fissures. Distances between markers and the closest lobe boundary in the segmentation were determined as measure of segmentation accuracy. The distance in the cranial and caudal direction was measured as positive and negative, respectively. RESULTS: Out of 25 CT datasets, 2 were excluded because of undefined fissures in the image and failing bronchus labelling, respectively. In 2 further datasets, the right horizontal fissure was not visible. The median distance was 0.00mm (IQR: -0.86-0.00mm). The right and left oblique fissures' median distances were 0.00mm (0.00-0.80mm) and 0.00mm (-0.86-0.00mm), respectively. Median distance for the right horizontal fissure was 0.00mm (-1.05-0.00mm). CONCLUSION: Our fully-automatic algorithm provides a stable and reliable identification of lung fissures.

B.2 2018

Pulmonary Lobe Segmentation in CT Images using Alpha-Expansion

Nicola Giuliani, Christian Payer, Michael Pienn, Horst Olschewski and Martin Urschler

In: *Proceedings of the 13th International Joint Conference on Computer Vision, Imaging and Computer Graphics Theory and Applications*

January 2018, Madeira, PT

(Accepted for poster presentation)

Abstract: Fully-automatic lung lobe segmentation in pathological lungs is still a challenging task. A new approach for automatic lung lobe segmentation is presented based on airways, vessels, fissures and prior knowledge on lobar shape. The anatomical information and prior knowledge are combined into an energy equation, which is minimized via graph cuts to yield an optimal segmentation. The algorithm is quantitatively validated on an in-house dataset of 25 scans and on the LObe and Lung Analysis 2011 (LOLA11) dataset, which contains a range of different challenging lungs (total of 55) with respect to lobe segmentation. Both experiments achieved solid results including a median absolute distance from manually set fissure markers of 1.04mm (interquartile range: 0.88-1.09mm) on the in-house dataset and a score of 0.866 on the LOLA11 dataset. We conclude that our proposed method is robust even in case of pathologies.

Bibliography

- [1] Aziz, A., Ashizawa, K., Nagaoki, K., and Hayashi, K. (2004). High resolution CT anatomy of the pulmonary fissures. *Journal of thoracic imaging*, 19(3):186–191. (page 1)
- [2] Benmansour, F., Türetken, E., and Fua, P. (2013). Tubular geodesics using oriented flux: An ITK implementation. Technical report. (page 17)
- [3] Benza, R. L., Miller, D. P., Barst, R. J., Badesch, D. B., Frost, A. E., and McGoon, M. D. (2012). An evaluation of long-term survival from time of diagnosis in pulmonary arterial hypertension from the REVEAL Registry. *CHEST Journal*, 142(2):448–456. (page 3)
- [4] Boykov, Y. and Kolmogorov, V. (2004). An experimental comparison of min-cut/max-flow algorithms for energy minimization in vision. *IEEE transactions on pattern analysis and machine intelligence*, 26(9):1124–1137. (page xiii, 26, 27, 28, 29)
- [5] Boykov, Y., Veksler, O., and Zabih, R. (2001). Fast approximate energy minimization via graph cuts. *IEEE Transactions on pattern analysis and machine intelligence*, 23(11):1222–1239. (page xiv, 5, 11, 31, 32)
- [6] Bragman, F. J., McClelland, J. R., Jacob, J., Hurst, J. R., and Hawkes, D. J. (2017). Pulmonary lobe segmentation with probabilistic segmentation of the fissures and a groupwise fissure prior. *IEEE transactions on medical imaging*, 36(8):1650–1663. (page 9, 57)
- [7] Cetti, E. J., Moore, A. J., and Geddes, D. M. (2006). Collateral ventilation. *Thorax*, 61(5):371–373. (page 2)
- [8] Cronin, P., Gross, B. H., Kelly, A. M., Patel, S., Kazerooni, E. A., and Carlos, R. C. (2010). Normal and accessory fissures of the lung: evaluation with contiguous volumetric thin-section multidetector CT. *European journal of radiology*, 75(2):e1–e8. (page 1, 2)
- [9] de Hoop, B., van Ginneken, B., Gietema, H., and Prokop, M. (2012). Pulmonary perifissural nodules on CT scans: rapid growth is not a predictor of malignancy. *Radiology*, 265(2):611–616. (page 2)
- [10] Dinic, E. (1970). Algorithm for Solution of a Problem of Maximum Flow in Networks with Power Estimation. In *Soviet Math. Dokl.*, volume 11, pages 1277–1280. (page 26)
- [11] Doel, T., Gavaghan, D. J., and Grau, V. (2015). Review of automatic pulmonary lobe segmentation methods from CT. *Computerized Medical Imaging and Graphics*, 40:13–29. (page 2, 3, 4, 7)
- [12] Elias, P., Feinstein, A., and Shannon, C. (1956). A note on the maximum flow through a network. *IRE Transactions on Information Theory*, 2(4):117–119. (page 25)

- [13] Fishman, A., Martinez, F., Naunheim, K., Piantadosi, S., Wise, R., Ries, A., Weimann, G., and Wood, D. E. (2003). A randomized trial comparing lung-volume-reduction surgery with medical therapy for severe emphysema. *New England Journal of Medicine*, 348(21):2059–2073. (page 3)
- [14] Ford, L. R. and Fulkerson, D. R. (1956). Maximal flow through a network. *Canadian journal of Mathematics*, 8(3):399–404. (page 25)
- [15] Ford Jr, L. R. and Fulkerson, D. R. (2015). *Flows in networks*. Princeton university press. (page 26)
- [16] Frangi, A. F., Niessen, W. J., Vincken, K. L., and Viergever, M. A. (1998). Multiscale vessel enhancement filtering. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 130–137. Springer. (page 20, 21)
- [17] Galiè, N., Humbert, M., Vachiery, J.-L., Gibbs, S., Lang, I., Torbicki, A., Simonneau, G., Peacock, A., Vonk Noordegraaf, A., Beghetti, M., et al. (2015). 2015 ESC/ERS Guidelines for the diagnosis and treatment of pulmonary hypertension: the Joint Task Force for the Diagnosis and Treatment of Pulmonary Hypertension of the European Society of Cardiology (ESC) and the European Respiratory Society (ERS): endorsed by: Association for European Paediatric and Congenital Cardiology (AEPC), International Society for Heart and Lung Transplantation (ISHLT). *European heart journal*, 37(1):67–119. (page 3)
- [18] Goldberg, A. V., Tardos, É., and Tarjan, R. (1989). Network flow algorithm. Technical report, Cornell University Operations Research and Industrial Engineering. (page 24)
- [19] Goldberg, A. V. and Tarjan, R. E. (1988). A new approach to the maximum-flow problem. *Journal of the ACM (JACM)*, 35(4):921–940. (page 26)
- [20] Gonzalez, R. C., Woods, R. E., et al. (1992). Digital image processing. (page 16)
- [21] Gu, S., Wang, Z., Siegfried, J. M., Wilson, D., Bigbee, W. L., and Pu, J. (2012a). Automated lobe-based airway labeling. *Journal of Biomedical Imaging*, 2012:1. (page 4, 14, 15, 16)
- [22] Gu, S., Zheng, Q., Siegfried, J., and Pu, J. (2012b). Robust pulmonary lobe segmentation against incomplete fissures. In *Proc. SPIE*, volume 8315, page 831535. (page 8, 9)
- [23] Gülsün, M., Arıyürek, O. M., Cömert, R. B., and Karabulut, N. (2006). Variability of the pulmonary oblique fissures presented by high-resolution computed tomography. *Surgical and Radiologic Anatomy*, 28(3):293–299. (page 1)

- [24] Hayashi, K., Aziz, A., Ashizawa, K., Hayashi, H., Nagaoki, K., and Otsuji, H. (2001). Radiographic and CT appearances of the major fissures. *Radiographics*, 21(4):861–874. (page 1)
- [25] Helmberger, M., Pienn, M., Urschler, M., Kullnig, P., Stollberger, R., Kovacs, G., Olschewski, A., Olschewski, H., and Bálint, Z. (2014). Quantification of tortuosity and fractal dimension of the lung vessels in pulmonary hypertension patients. *PLoS one*, 9(1):e87515. (page 3, 4, 13, 16)
- [26] Hlastala, M. P. and Berger, A. J. (2001). *Physiology of respiration*. Oxford University Press. (page 1)
- [27] Hochhegger, B., Hochhegger, D. Q. d. R., Irion, K., Sartori, A. P., Gazzoni, F. F., and Marchiori, E. (2013). Intrapulmonary lymph node: a common and underrecognized tomography finding. *Jornal Brasileiro de Pneumologia*, 39(6):757–758. (page 3)
- [28] Hoepfer, M. M., Bogaard, H. J., Condliffe, R., Frantz, R., Khanna, D., Kurzyna, M., Langleben, D., Manes, A., Satoh, T., Torres, F., et al. (2013). Definitions and diagnosis of pulmonary hypertension. *Journal of the American College of Cardiology*, 62(25 Supplement):D42–D50. (page 3)
- [29] Ingenito, E. P., Loring, S. H., Moy, M. L., Mentzer, S. J., Swanson, S. J., Hunsaker, A., McKEE, C. C., and Reilly, J. J. (2001). Comparison of physiological and radiological screening for lung volume reduction surgery. *American journal of respiratory and critical care medicine*, 163(5):1068–1073. (page 3)
- [30] Ingenito, E. P., Wood, D. E., and Utz, J. P. (2008). Bronchoscopic lung volume reduction in severe emphysema. *Proceedings of the American Thoracic Society*, 5(4):454–460. (page 2)
- [31] Ising, E. (1925). Beitrag zur theorie des ferromagnetismus. *Zeitschrift für Physik A Hadrons and Nuclei*, 31(1):253–258. (page 31)
- [32] Jeffery, P. K. (1998). Structural and inflammatory changes in COPD: a comparison with asthma. *Thorax*, 53(2):129. (page 3)
- [33] Kim, S. S., Seo, J. B., Lee, H. Y., Nevrekar, D. V., Forssen, A. V., Crapo, J. D., Schroeder, J. D., and Lynch, D. A. (2013). Chronic obstructive pulmonary disease: lobe-based visual assessment of volumetric CT by using standard images-comparison with quantitative CT and pulmonary function test in the COPD Gene study. *Radiology*, 266(2):626–635. (page 3)
- [34] Kindermann, R., Snell, J., and Society, A. M. (1980). *Markov Random Fields and Their Applications*. AMS books online. American Mathematical Society. (page 31)

- [35] Koster, T. D. and Slebos, D.-J. (2016). The fissure: interlobar collateral ventilation and implications for endoscopic therapy in emphysema. *International journal of chronic obstructive pulmonary disease*, 11:765. (page 2)
- [36] Kuhnigk, J.-M., Hahn, H., Hindennach, M., Dicken, V., Krass, S., and Peitgen, H.-O. (2003). Lung lobe segmentation by anatomy-guided 3 D watershed transform. In *Proceedings of SPIE*, volume 5032, pages 1482–1490. (page 8, 9)
- [37] Lassen, B., van Rikxoort, E. M., Schmidt, M., Kerkstra, S., van Ginneken, B., and Kuhnigk, J.-M. (2013). Automatic segmentation of the pulmonary lobes from chest CT scans based on fissures, vessels, and bronchi. *IEEE transactions on medical imaging*, 32(2):210–222. (page 4, 8, 9, 10, 20, 56, 57)
- [38] Law, M. W. and Chung, A. C. (2008). Three dimensional curvilinear structure detection using optimally oriented flux. In *European conference on computer vision*, pages 368–382. Springer. (page 17)
- [39] LeCun, Y., Chopra, S., Hadsell, R., Ranzato, M., and Huang, F. (2006). A tutorial on energy-based learning. *Predicting structured data*, 1:0. (page 32)
- [40] Leung, A. N. (1999). Pulmonary tuberculosis: the essentials. *Radiology*, 210(2):307–322. (page 3)
- [41] LOLA11 (2017). [Online; accessed 30-October-2017]. (page 44, 57)
- [42] Maxfield, R. A. (2004). New and emerging minimally invasive techniques for lung volume reduction. *CHEST Journal*, 125(2):777–783. (page 2)
- [43] Otsu, N. (1979). A threshold selection method from gray-level histograms. *IEEE transactions on systems, man, and cybernetics*, 9(1):62–66. (page 16)
- [44] Payer, C., Pienn, M., Bálint, Z., Shekhovtsov, A., Talakic, E., Nagy, E., Olschewski, A., Olschewski, H., and Urschler, M. (2016). Automated integer programming based separation of arteries and veins from thoracic CT images. *Medical image analysis*, 34:109–122. (page 3, 4, 6, 16, 17, 59, 75)
- [45] Pu, J., Fuhrman, C., Durick, J., Leader, J. K., Klym, A., Scirba, F. C., and Gur, D. (2010). Computerized assessment of pulmonary fissure integrity using high resolution CT. *Medical physics*, 37(9):4661–4672. (page 2)
- [46] Pu, J., Wang, Z., Gu, S., Fuhrman, C., Leader, J. K., Meng, X., Tedrow, J., and Scirba, F. C. (2014). Pulmonary fissure integrity and collateral ventilation in COPD patients. *PloS one*, 9(5):e96631. (page 2)
- [47] Pu, J., Zheng, B., Leader, J. K., Fuhrman, C., Knollmann, F., Klym, A., and Gur, D. (2009). Pulmonary lobe segmentation in CT examinations using implicit surface fitting. *IEEE transactions on medical imaging*, 28(12):1986–1996. (page 8, 9)

- [48] Qi, S., van Triest, H. J., Yue, Y., Xu, M., and Kang, Y. (2014). Automatic pulmonary fissure detection and lobe segmentation in CT chest images. *Biomedical engineering online*, 13(1):59. (page 9)
- [49] Raasch, B., Carsky, E., Lane, E., O’callaghan, J., and Heitzman, E. (1982). Radiographic anatomy of the interlobar fissures: a study of 100 specimens. *American Journal of Roentgenology*, 138(6):1043–1049. (page 1, 19)
- [50] Rees, D. and Murray, J. (2007). Silica, silicosis and tuberculosis. *The International Journal of Tuberculosis and Lung Disease*, 11(5):474. (page 3)
- [51] Richfield, D. (2014). Medical gallery of David Richfield. *Wikiversity Journal of Medicine*, 1(2). (page 5)
- [52] Ross, J. C., Estépar, R. S. J., Kindlmann, G., Díaz, A., Westin, C.-F., Silverman, E. K., and Washko, G. R. (2010). Automatic lung lobe segmentation using particles, thin plate splines, and maximum a posteriori estimation. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 163–171. Springer. (page 8, 9)
- [53] Ross, J. C., Kindlmann, G. L., Okajima, Y., Hatabu, H., Díaz, A. A., Silverman, E. K., Washko, G. R., Dy, J., and Estépar, R. S. J. (2013). Pulmonary lobe segmentation based on ridge surface sampling and shape model fitting. *Medical physics*, 40(12). (page 8, 9)
- [54] San Jose Estepar, R., Ross, J. C., Harmouche, R., Onieva, J., Diaz, A. A., and Washko, G. R. (2015). Chest imaging platform: an open-source library and workstation for quantitative chest imaging. In *C66. Lung imaging II: New probes and emerging technologies*, pages A4975–A4975. Am Thoracic Soc. (page 53)
- [55] Shamonin, D. P., Staring, M., Bakker, M. E., Xiao, C., Stolk, J., Reiber, J. H., and Stoel, B. C. (2012). Automatic lung lobe segmentation of COPD patients using iterative B-spline fitting. In *SPIE Medical Imaging*, pages 83140W–83140W. International Society for Optics and Photonics. (page 8, 9)
- [56] Simonneau, G., Robbins, I. M., Beghetti, M., Channick, R. N., Delcroix, M., Denton, C. P., Elliott, C. G., Gaine, S. P., Gladwin, M. T., Jing, Z.-C., et al. (2009). Updated clinical classification of pulmonary hypertension. *Journal of the American College of Cardiology*, 54(1):S43–S54. (page 3)
- [57] Sonka, M., Hlavac, V., and Boyle, R. (1998). *Image Processing: Analysis and Machine Vision*. CL-Engineering, 2 edition. (page 16)
- [58] Travis, W. D., King, T. E., Bateman, E. D., Lynch, D. A., Capron, F., Center, D., Colby, T. V., Cordier, J. F., DuBois, R. M., Galvin, J., et al. (2002). American

- thoracic society/european respiratory society international multidisciplinary consensus classification of the idiopathic interstitial pneumonias. *American journal of respiratory and critical care medicine*, 165(2):277–304. (page 3)
- [59] Turetken, E., Benmansour, F., Andres, B., Pfister, H., and Fua, P. (2013). Reconstructing loopy curvilinear structures using integer programming. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1822–1829. (page 17)
- [60] Turner, M. O., Mayo, J. R., Müller, N. L., Schulzer, M., and FitzGerald, J. M. (2006). The value of thoracic computed tomography scans in clinical diagnosis: a prospective study. *Canadian respiratory journal*, 13(6):311–316. (page 4)
- [61] Ukil, S. and Reinhardt, J. M. (2009). Anatomy-guided lung lobe segmentation in X-ray CT images. *IEEE transactions on medical imaging*, 28(2):202–214. (page 8, 9)
- [62] Van Rikxoort, E. and Van Ginneken, B. (2011). Automatic segmentation of the lungs and lobes from thoracic CT scans. In *Proc. 4th Int. Workshop Pulmonary Image Anal*, pages 261–268. (page 9)
- [63] Van Rikxoort, E. M., de Hoop, B., van de Vorst, S., Prokop, M., and van Ginneken, B. (2009). Automatic segmentation of pulmonary segments from volumetric chest CT scans. *IEEE transactions on medical imaging*, 28(4):621–630. (page 8, 9)
- [64] van Rikxoort, E. M., Goldin, J. G., Galperin-Aizenberg, M., Abtin, F., Kim, H. J., Lu, P., van Ginneken, B., Shaw, G., and Brown, M. S. (2012). A method for the automatic quantification of the completeness of pulmonary fissures: evaluation in a database of subjects with severe emphysema. *European radiology*, 22(2):302–309. (page 2)
- [65] Van Rikxoort, E. M., Prokop, M., de Hoop, B., Viergever, M. A., Pluim, J. P., and van Ginneken, B. (2010). Automatic segmentation of pulmonary lobes robust against incomplete fissures. *IEEE transactions on medical imaging*, 29(6):1286–1296. (page 8, 9, 57)
- [66] Van Rikxoort, E. M. and Van Ginneken, B. (2013). Automated segmentation of pulmonary structures in thoracic computed tomography scans: a review. *Physics in medicine and biology*, 58(17):R187. (page 3, 7)
- [67] West, J. B. (2012). *Respiratory physiology: the essentials*. Lippincott Williams & Wilkins. (page 1)
- [68] Wu, F.-Y. (1982). The potts model. *Reviews of modern physics*, 54(1):235. (page 5, 31)

-
- [69] Zhang, L., Hoffman, E. A., and Reinhardt, J. M. (2006). Atlas-driven lung lobe segmentation in volumetric X-ray CT images. *IEEE transactions on medical imaging*, 25(1):1–16. (page [8](#), [9](#))