

Three-dimensional graph cuts for human airway analysis

Master thesis

Jens Petersen

Supervisor: Mads Nielsen
Institute of Computer Science
University of Copenhagen
August 31, 2010

ABSTRACT

Airway abnormalities in the form of morphological changes are associated with Chronic Obstructive Pulmonary Disease (COPD). Computerized Tomography (CT) is becoming an increasingly popular imaging tool to investigate these pathologies. This thesis focuses on airway wall segmentation algorithms applicable for such investigations.

We present a globally optimal three-dimensional graph cut algorithm for the purpose. Its novelty include the construction of a graph based on columns following flow lines calculated from the convolution of an initial segmentation of the airways with different kernels. We evaluate the method within a framework of 649 manually annotated two-dimensional cross-sectional images from 15 different subjects and show that the obtained results are superior to a previously developed two-dimensional method and a similar three-dimensional method using straight columns pointing in the initial segmentation surface normal direction.

We quantify the method's use as a COPD diagnostic tool in a series of large scale tests, involving 2,512 CT scans from a lung cancer screening study, comparing the results with the two-dimensional method. The majority of the measurements conducted with the proposed method are shown to be statistically more reproducible, correlate more with lung function and have statistically better diagnostic ability quantified as area under the receiver operating characteristic for different disease stages.

We investigate the association between subject health as quantified by answers to the St. George's Respiratory Questionnaire and airway measurements obtained in 690 subjects and show moderate but significant correlations with especially disease symptoms and what impact it has on their psychological and social functioning.

The results presented in this thesis, demonstrate the use of CT and the proposed method for COPD investigation and diagnosis.

RESUMÉ

Luftvejsuregelmæssigheder i form af morfologiske ændringer er forbundet med Kronisk Obstruktiv Lungesygdom (KOL). ComputerTomografi (CT) er ved at blive et populært redskab til at undersøge disse patologier. Dette speciale fokuserer på metoder, der kan anvendes til at segmentere luftvejsvæggene til brug ved sådanne undersøgelser.

Vi præsenterer en global optimal metode, som via tredimensionelle grafsnit kan løse problemet. Det nyskabende ved metoden er blandt andet konstruktionen af en graf, hvis kolonner følger integralkurver udregnet fra foldningen af en initial segmentation af luftvejene og forskellige kerner. Metoden evalueres ved hjælp af 649 manuelt annoterede todimensionelle tværsnit fra 15 forskellige personer. De opnåede resultater er bedre end for en tidligere udviklet todimensionel metode og en lignende tredimensionel metode, der bruger rette kolonner pegende i normalretningen af den initiale segmentationsoverflade.

Vi kvantificerer hvor brugbar metoden er som KOL diagnosticeringsværktøj i en test af stor skala, der involverer mere end 2.512 CT skanninger fra en screeningundersøgelse af lungekræft og sammenligner resultaterne med den todimensionelle metode. Størstedelen af de foretagne målinger er statistisk mere reproducerbare, viser større sammenhæng med lungefunktion og har statistisk bedre diagnostiske egenskaber, kvantificeret med arealet under ROC for forskellige sygdomsstadier.

Vi undersøger associationen mellem helbred, kvantificeret med St. George's Respiratory Questionnaire, og luftvejsmål i 690 personer, og kan vise moderate men signifikante sammenhænge mellem specielt sygdomssymptomer og hvordan sygdommen påvirker deres psykologiske og sociale funktion.

Resultaterne præsenteret i dette speciale, viser at CT og den udviklede metode kan bruges til at undersøge og diagnosticere KOL.

Contents

1	Introduction	6
1.1	Motivation	6
1.2	Problem domain	7
1.3	Document outline	7
1.4	Requirements on the reader	8
1.5	Acknowledgements	8
1.6	Used software	8
2	Previous Methods	9
2.1	Airway segmentation	9
2.2	Airway wall segmentation	9
2.2.1	Optimal net	11
2.2.2	Smoothness constraints	15
2.2.3	Coupled surfaces	16
2.2.4	Edge penalties	17
2.2.5	Segmentation of non-flat surfaces	18
2.2.6	Medial axes columns	18
2.2.7	Electric lines of force inspired columns	20
2.2.8	Cost functions	20
2.3	Airway abnormality measures	22
3	Method	23
3.1	Airway segmentation	24
3.2	Tree extraction	24
3.3	Graph construction	25
3.3.1	Surface mesh construction	30
3.3.2	Medial axes columns	31
3.3.3	Flow line columns	31
3.3.4	Numerical integration of columns	33
3.4	Cost functions	35
3.5	Training and evaluation	35
3.5.1	Manually segmented cross-sections	36
3.5.2	Algorithm	36
3.5.3	Practical issues	37
3.6	Leak detection	37
3.7	Measurements	37
3.8	Statistical investigations	38

4	Material	39
4.1	DLCST	39
4.2	CBQ	40
4.3	Incomplete data	40
5	Results	41
5.1	Airway segmentation and tree extraction	41
5.2	Parameters	41
5.2.1	Graph resolution	41
5.2.2	Flow line regularisation	42
5.2.3	Normal error tolerance and medial axes neighbours	42
5.2.4	Derivative weighting, smoothness and separation penalties	43
5.2.5	Leak detection	44
5.3	Training	44
5.4	Performance	45
5.5	Comparison with manual segmentations	45
5.6	Generation based comparison	47
5.6.1	Differences in leak detection methods	47
5.6.2	Reproducibility	47
5.6.3	Correlation with lung function	49
5.7	Per scan measures	50
5.7.1	Reproducibility and correlation with lung function	50
5.7.2	COPD Diagnostic ability	51
5.7.3	Bifurcation regions	52
5.7.4	St. George’s Respiratory Questionnaire	52
6	Discussion	53
6.1	Medial axes and normal directions from Delaunay balls	53
6.2	Two versus three-dimensional methods	54
6.2.1	Measure differences	54
7	Conclusion	55
8	Future work	57
8.1	Improved leak detection	57
8.2	Parallelisation	57
8.3	Repetitive approach	57
A	Algorithms	58
A.1	Normal and feature size algorithms	58
A.2	Training algorithm	60

1. INTRODUCTION

The focus of this thesis will be on algorithms that are applicable for the problem of segmenting the human airway wall in Computerized Tomography (CT) images. We will describe some of the previous attempts at solving the problem and the practical results obtained, with a special emphasis on graph cut methods.

Our main contributions are the development of a novel three-dimensional graph cut method, whose graph columns follow greatest ascent and descent flow lines. We show that the employed globally optimal solution is a simplified solution to special cases of the optimal VCE-weight net surface problem,¹ which should result in faster running times. Columns following the electric lines of force were introduced in reference 2 as a way to avoid the graph self-intersecting. We show that the calculation of such columns can be done using simple convolutions and that the Gaussian kernel could be a better choice for the job than a kernel based on the electric field strength. We demonstrate how such constructed *curved* columns are superior to straight columns following the initial surface normal direction as employed by previous methods³⁻⁸ and a two-dimensional method previously published by the author.⁹ The methods are compared within a framework of 649 manually annotated two-dimensional cross-sections extracted from the airways of 15 different subjects. Additionally we investigate the developed method's use as a tool to diagnose Chronic Obstructive Pulmonary Disease (COPD) in large scale tests involving 2,512 CT images from the Danish Lung Cancer Screening Trial (DLCST).¹⁰ Airway abnormality measurements obtained using the methods are evaluated, using repeated images taken roughly one year apart, for reproducibility, correlation with lung function and COPD diagnostic ability.

Additionally we investigate how airway abnormality measures obtained using the proposed method are associated with subjects' perception of their general health status, disease symptoms, ability to participate in activities and psychological and social functioning as quantified by the St. George's Respiratory Questionnaire (SGRQ).

1.1 Motivation

The motivation for writing this thesis initially came from the cooperation between the DLCST, run by Gentofte University Hospital, and the Institute of Computer Science at the University of Copenhagen (DIKU), where I am a master student. This cooperation meant that DIKU had access to a large database of CT scans and lung function measurements in a population at high risk of developing COPD. Since lung function tests are the current gold standard of COPD diagnosis,¹¹ a unique chance to investigate the effects of COPD on the airways and algorithms for doing so, arose. This initially resulted in the development of the mentioned two-dimensional method.⁹ From the experience gained and results obtained it was clear that expanding this method to make full use of the three-dimensional nature of CT images could result in more accurate segmentations and measurements.

The reason airway wall segmentation methods are relevant for the prediction of COPD is because the disease is associated with airway abnormalities such as narrowing of the smaller airways and thickening of the airway walls. Outside the airways the disease is known to be responsible for the destruction of lung tissue, known as emphysema. These changes cause shortness of air, leading to loss of mobility, sickness and death.^{11,12} It is known to be caused, amongst other things, by smoking, which is why it is also known as smoker's lungs.

CT is becoming an increasingly popular tool to investigate COPD pathologies, however manual examination of the generated images is an enormous task because

of their size and three-dimensional nature and the complexity of the human airway system. Even though developments in computer power have made complete segmentations of the inner and outer airway wall border possible, the use of CT for airway measurements is still in the research phase and algorithms are still dealing with basic problems such as getting reproducible measurements in such a dynamic environment.

It should be mentioned that while we focus on COPD a range of diseases affect the airways and the methods described and developed might be just as relevant for those. For instance: bronchiectasis,¹³ cystic fibrosis^{14,15} and asthma¹⁶ are all associated with airway remodelling and associations with lung cancer have also been found.¹⁷ Possible applications could be within diagnosis, development of new drugs, scientific investigation of causes etcetera.

1.2 Problem domain

The human airway system is a rather complicated tree-like structure, which varies in size, shape and morphology. Starting with the trachea, which branches out into the bronchi, which continue branching and getting smaller in size until around generation 15 on average, counting the trachea as zero, they reach the bronchioles. The bronchioles branch further, while the airway wall gets more and more plastered with alveoli, this is where the gas exchange takes place. The alveolar ducts terminate in grape-like clusters known as alveolar sacs.¹⁸ Most of this system is below the resolution of the CT scanner, however even the visible parts are so complex and biologically and dynamically varied, such as the changes caused by inspiration, that segmentation methods must be fairly general and not make too many assumptions on its nature.

The airway wall is defined by the volume between the inner and outer airway wall border. The volume inside the airway wall is what conducts the air in and out of the lungs and is therefore usually air filled and called the lumen, whereas the volume outside the airway wall and inside the lungs mostly consists of lung parenchyma, but also vessels, connective tissue, muscles etcetera.

As CT is based on radiodensity, larger and denser structures such as the trachea and the main bronchi are easier to segment than the smaller airways. This also means that the inner airway wall border is easier to segment than the outer airway wall border, because of the greater contrast differences that exist between the air filled lumen and the airway wall, which is often cartilage, smooth muscles etc. with a density slightly higher than water, whereas the volume outside the airway has a density somewhere between air and water. The problem is further complicated by the fact that most airways are close to other structures, such as vessels and other airways with similar densities.

1.3 Document outline

We begin by describing previous relevant methods in section 2. Starting the chapter with a short introduction to algorithms that can segment the airways, as most of the airway wall segmentation methods, including our own rely on some kind of initial segmentation of the airways. A detailed explanation of the optimal net type of methods is also given within this section, as these are closely related to our own approach. Our method and details of the other methods included in the analysis are described in section 3. Included are also documentation of the training and evaluation procedures using manually annotated cross-sections and descriptions of the statistical investigation of the methods' use as COPD diagnostic tools. A short description of the DLCST and the Computer tomography, Biomarkers and Quality of life (CBQ) material, which the analysis is based on, is provided in section 4.

Results of the evaluations and the statistical investigations are described in section 5 and the implications of these are discussed in section 6. A conclusion is provided in section 7 and we round the thesis off with a discussion of possible future work in section 8.

1.4 Requirements on the reader

This thesis is intended to be a scientific contribution to the knowledge within the field of (medical) image analysis and segmentation and deals with many complicated topics from other fields, such as knowledge of the human respiratory system, COPD, medical imaging and statistics. Explaining all these topics in detail is infeasible and so a basic knowledge of at least some of them are required.

1.5 Acknowledgements

I would like to thank the following people for their support and guidance throughout the work period.

Mads Nielsen, my supervisor, for his guidance, many great ideas and suggestions.

Pechin Chien Pau Lo for letting me use his airway segmentation and tree extraction algorithms, his input and thoughts, and for letting me ask him all kinds of questions, about the software, the image cluster, the data, etcetera.

Marleen de Bruijne and Francois Lauze for their suggestions and input.

Asger Dirksen, Zaigham Saghir and Haseem Ashraf for medical input and help with access to DLCST and CBQ material.

In addition I would like to thank the people and organizations involved in the DLCST and the CBQ. The data from these studies made this thesis possible.

1.6 Used software

The following freely available software were used as a basis for this work.

- The CImg Library
C++ Template Image Processing Toolkit
<http://cimg.sourceforge.net>
- Insight Segmentation and Registration Toolkit (ITK)
<http://www.itk.org>
- MAXFLOW
Software for computing mincut/maxflow in a graph
<http://www.cs.ucl.ac.uk/staff/V.Kolmogorov/software.html>
- NormFet
Software for approximating normals and feature sizes from noisy point clouds
(available upon request)
<http://www.cse.ohio-state.edu/~tamaldey/normfet.html>
- Computational Geometry Algorithms Library (CGAL)
<http://www.cgal.org>
- Boost C++ libraries
<http://www.boost.org>

2. PREVIOUS METHODS

This section contains a description of previous relevant methods. We will begin with a short introduction to airway segmentation methods, because practically all airway wall segmentation methods rely on the existence of some initial segmentation of the airways.^{4,19–27} This initial segmentation, sometimes called the pre-segmentation, provides a good estimate of the airway surface position and orientation, but does not have to be locally accurate. The inner and outer airway wall borders are then found by searching in a neighbourhood of this surface. We use the term *airway segmentation*, because we do not care whether the method segments the lumen and the wall, or just the lumen. However it is perhaps a bit of a misnomer given that most of these methods actually use the clearly defined border between the lumen and airway wall to segment only the lumen.

2.1 Airway segmentation

Note that this section should be thought of as a short introductory overview of airway segmentation methods. We will refrain from going into detail and just mention some of the previously employed methods as the process is a complicated subject in itself.

A good airway segmentation algorithm should be able to find a large amount of the airway tree branches and have few false positives. See for instance the Exact'09 study,²⁸ in which multiple algorithms were compared on tree length, branch count and amount of false positives. Another requirement, which is especially relevant if the initial segmentation is to be used as a basis for airway wall segmentation and pathology measurements, is a reasonably accurate and precise surface representation.

The majority of airway segmentation methods are based on region growing processes.^{29–34} These algorithms use a seed point, usually in the trachea, and then continuously add voxels to the boundary of the segmentation using some selection criteria. A recurring problem with them, is that they often leak into the lung parenchyma, in which case the segmented airway volume grows explosively. These leaks need to be controlled and most of the mentioned methods implement various techniques for doing so.

Other types of methods are worth noting, such as front propagation methods,³⁵ in which the segmentation is done along some front moving through the airway tree. This enables natural restrictions on the growth, such as for instance limiting the radius of the front to some fraction of the local airway diameter thereby stopping uncontrollable leaks. Fuzzy connectivity methods,²² in which voxels are classified by some measure of connectedness. Or morphologically based methods^{36–38} that use grey-scale mathematical morphology operations. These methods look at the lung in its entirety, making it possible to detect airways, that for some reason, such as pathology, blockage etc. are not connected to the seed point.

Some interesting methods are difficult to classify, such as reference 39, which use a combination of a conservative region growing method to capture the large branches, airway section filtering which looks for smaller branches and graph searching to clean up false positives.

2.2 Airway wall segmentation

The problem of segmenting the human airway wall can be greatly simplified if analysis can be confined to two-dimensional slice images oriented perpendicular to the specific airway branch direction and extracted outside branch-points. In such

images the airway wall borders resemble two concentric circles. The centre and direction of the airway branches can be calculated from an initial airway segmentation by using skeletonisation,^{22,40} three-dimensional thinning^{41,42} or front propagation algorithms.^{35,43}

Many simpler approaches consist of simply casting rays from the airway centre and out within such cross-sectional images,^{19,21,44-47} reducing the problem to a simple one-dimensional edge-detection problem. The Full Width at Half Maximum (FWHM) has been a popular principle to find the location of the airway wall borders within such rays.⁴⁴⁻⁴⁸ Using this principle the inner airway wall border is defined to be at the position having the average of the minimum and maximum intensity, on the line segment of the ray going from the lumen centre to the wall intensity maximum. The outer wall border is found similarly on the ray from the wall intensity maximum to the point where the intensity values reach an outer minimum. The method has been shown to be strongly influenced by imaging parameters such as choice of reconstruction kernel and the airways' size and shape,⁴⁴ resulting in a consistent under- or overestimation of the measured structure's dimensions. We investigated the method in reference⁹ and found that it works reasonably well despite its simplicity, however more accurate methods exist.

Phase congruency as a feature descriptor was originally suggested by Morrone and Owens,⁴⁹ defining feature positions as the places where the harmonic components of the image are maximally in phase. This was used in reference²¹ to derive a *bronze standard* for airway wall location by using the maximum phase congruency of the signals obtained from multiple reconstruction kernels to locate the airway wall border positions. The method was compared with FWHM with very promising results. However the need for data reconstructed with multiple reconstruction kernels limits its usage.

Some of the problems with the FWHM method can be removed by modelling the scanning process. This was done in reference⁵⁰ by using a non-linear optimisation technique to match the observed ray profile with an ideal ray profile. The set of parameters obtained yield estimates of the inner and outer airway wall border locations. The method was demonstrated to give a reduction in measurement bias for thin-walled airways when compared with the FWHM method on phantom scans. Reference 19 used an integral based closed-form solution and a calibration parameter to model the scanner point spread function obtaining very accurate measurements when compared with phantom data. The method has been applied in a series of large scale tests⁵¹ showing good correlation with lung function tests.

All these methods can be said to be one-dimensional given that they independently sample the airway wall borders in each ray. This is problematic in areas where the airway borders are weakly defined. Methods that look for circular,⁵² elliptic^{53,54} or tubular⁵⁵ structures can be used in an attempt to overcome this. Elliptic long and short axes can also be used to adjust for the bias introduced by the cross-sectional image not slicing the airway perpendicularly. Not all cross-sectionally cut airways are circular or elliptic though, and so these methods are likely too rigid.

More sophisticated methods have been developed that allow one to put constraints or penalties on the solutions, such as for instance forcing them to have some degree of smoothness. For instance Saragaglia et al.²⁶ developed a method using mathematical morphology operators, energy-controlled propagation and reliable wall-based smoothing to find the wall borders. The method was demonstrated using phantoms to be able to estimate the location of the airway wall borders even when the airway was abutted by a vessel.

Graph searching was used in reference 20, 22 by transforming the cross-sections to polar coordinates. In such images the airway wall borders become horizontal (or vertical) lines. Using this fact a minimum-path graph was constructed with the use of cost functions, indicating the inverse likelihood of the airway wall border being in any specific position. The methods were shown to achieve sub-voxel accuracy on phantoms.

We used the same technique of transforming the cross-sections to polar coordinates in reference 9, however instead a maximum-flow algorithm was used to solve the problem. The method is also related to the optimal VCE-weight net surface problem described in section 2.2.1 and the edge penalties of reference 56, described in section 2.2.4, in that it uses smoothness and surface separation penalty edges to prioritise wall border smoothness and closeness. It was evaluated using manually annotated cross-sections and found to result in more accurate segmentations than the FWHM method and a similar method without separation penalties. The method's worth as a COPD diagnostic tool was evaluated in a large scale test involving more than 723 CT scans and found, for the majority of the investigated measures, to result in higher reproducibility, more correlation with lung function and better diagnostic ability, quantified as Area Under the receiver operating Characteristic (AUC), when compared with results of measures obtained with the FWHM method. The method was also used in a large scale quantitative analysis of airway abnormality measures in reference 57. The airway measures were found to correlate more with lung function and be less influenced by covariates, such as total lung volume, total lung weight, gender and age, than the investigated emphysema measures.

Recently true three-dimensional airway wall segmentation methods have been developed. These methods can use the three-dimensional structure of the airway wall border to infer its position in difficult areas. For instance, a cross-sectional image might only reveal one structure when an artery abuts the airway, whereas looking at the image in three dimensions, could reveal that the artery is separated from the airway further along the branch.

Saragaglia et al. extended their two-dimensional method to three dimensions using a deformable mesh, constructed from an initial segmentation of the lumen.²⁷ The deformable mesh was evolved using forces attracting it toward areas with high intensity and gradient magnitude values, an elastic force, which penalised local wall thickness variations, and a regularisation force, which locally smoothed the result. The results were not found to be significantly better than their two-dimensional method, when compared using a three-dimensional image model, simulating a cylindrical bronchus with branching areas.

The following sections details the recent growth in the use of graph cut methods, and especially within the optimal net type of methods, for three-dimensional segmentation problems within medical image analysis. These methods are related to our own approach and so we describe them thoroughly. It should be noted that some people use the term, *graph cut method*, exclusively to describe the methods of Greig et al.,⁵⁸ however we have adopted the common more loose definition, of any method that employ maximum-flow/minimum cut optimisation to solve the problem.

2.2.1 Optimal net

The use of graph cut methods in computer vision was first introduced by Greig et al.⁵⁸ They showed how the maximum a posteriori probability of a binary image, could be evaluated using maximum-flow/minimum cut algorithms. The problem is formulated in terms of a flow network, a graph, consisting of a source node, which the flow stems from, and a sink node, which the flow moves towards. The edges in

the graph have flow capacities and finding the maximum flow the network allows, yields the optimal solution to the original problem.

We will start with a few definitions. A *cut* is a disjoint partitioning of the vertices in the graph into a source and a sink set of nodes. Edges going from a vertex in the source set to a vertex in the sink set are said to be *cut* or to be part of the *cut-set*. The cost of the cut is given by the total capacity of the edges in the cut-set. The *minimum cut problem* is the problem of minimizing the cost of the cut. The cost of a minimum cut can be shown to be equivalent to the maximum flow the network allows,⁵⁹ which is why it is also termed, a *maximum flow problem*.

A wide range of algorithms exists for solving maximum flow problems, such as for instance the Ford-Fulkerson algorithm,⁵⁹ the Edmonds-Karp algorithm,⁶⁰ variants based on the push-relabel algorithm⁶¹ or the Boykov-Kolmogorov algorithm.⁶² The last has been shown to be very efficient at calculating the maximum flow in the types of graphs dealt with in this section.^{24,62}

No generalization of the method of Greig et al.⁵⁸ exists for images of more than 2 colors. However Wu and Chen¹ showed that if the problem can be formulated in terms of what they call proper ordered multi-column graphs and either one of two combinatorial optimisation problems called the optimal V-weight and the optimal VCE-weight net surface problem on such graphs, then optimal solutions can be found in polynomial time using minimum cut algorithms. We will explain the concepts and method in the following, starting with a brief explanation of the notation.

Let (i, j) denote the undirected edge between vertex i and j . Similarly let $(i \xrightarrow{w} j)$ and $(i \xleftrightarrow{w} j)$, denote the directed edge from i to j and the bidirectional edge between i and j respectively with weight w .

Assuming the existence of some $(D-1)$ -dimensional base graph $G_B = (V_B, E_B)$, the space of all possible solutions is a D -dimensional undirected graph $G = (V, E)$, with vertices V and edges E , generated from G_B by associating each vertex $i \in V_B$ with a column V_i of vertices in V of length K . A net surface is a subgraph of G defined by a function $N : V_B \rightarrow \{0, 1, \dots, K-1\}$, such that N intersects each column of vertices in G exactly once, see figure 1(a). That is, the following two statements hold for all i and j in V_B :

$$(i, j) \in E_B \Rightarrow (i_{N(i)}, j_{N(j)}) \in E \quad (1)$$

$$N(i) \neq N(j) \Rightarrow i \neq j \quad (2)$$

It can be thought as a topology preserving mapping of G_B within the solution space defined by the graph G .

Note that G and G_B are purely mathematical constructs. They are used to prove certain qualities of the solution, given by N . Later we will specify how the maximum flow graph \tilde{G} is constructed from G , such that these properties are upheld.

In the following chapters we will let V_i and V_j denote columns of vertices in V generated from connected vertices i and j in the base graph G_B . Single vertices in these columns are denoted i_k and $j_{k'}$, where $k, k' \in \{0, 1, \dots, K-1\}$. Additionally let the edge interval of i_k on V_j , $I(i_k, j)$ be defined as the vertices in V_j connected to vertex i_k :

$$I(i_k, j) = \{j_{k'} \mid j_{k'} \in V_j, (i_k, j_{k'}) \in E\} \quad (3)$$

Define $L(i_k, j) = \min(\{k' \mid j_{k'} \in I(i_k, j)\})$ and $U(i_k, j) = \max(\{k' \mid j_{k'} \in I(i_k, j)\})$, that is, the minimum and maximum index respectively in the neighbouring column j connected to i_k . Then the *proper ordering* implies the two conditions given by

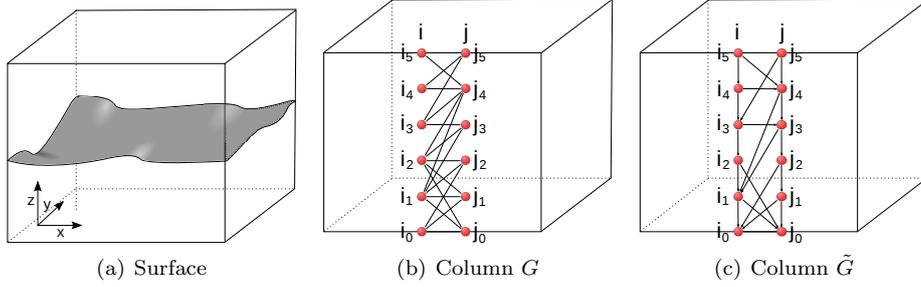


Figure 1. A net surface oriented along the x-y axis within a three-dimensional graph, symbolised by the cube, is illustrated in figure 1(a) and an example of the proper ordered columns of Wu and Chen¹ is shown in figure 1(b). The directed graph \tilde{G} build from G in figure 1(c) minus the source and sink edges.

equation 4 and 5. That each vertex in V_i is connected to a non-empty consecutive sequence of vertices in V_j :

$$(i, j) \in E_B \Rightarrow \forall i_k \in V_i \exists k' \exists s \geq 0 : \{(i_k, j_{k'}), (i_k, j_{k'+1}), \dots, (i_k, j_{k'+s})\} \in E \quad (4)$$

and for two consecutive vertices i_k and i_{k+1} in the column V_i the following must hold: i_k cannot be connected to vertices in the V_j column with higher indices, or further up, than i_{k+1} is and similarly i_{k+1} cannot be connected to vertices in the V_j column with lower indices, or further down, than i_k is:

$$(i, j) \in E_B \Rightarrow \forall \{i_k, i_{k+1}\} \in V_i : L(i_k, j) \leq L(i_{k+1}, j) \wedge U(i_k, j) \leq U(i_{k+1}, j) \quad (5)$$

An example of this graph construction is shown in figure 1(b).

Assuming each $v \in V$ is associated with some weight $w(v)$ inversely related to the likelihood that v is part of the surface of some object and let E_N and V_N denote the edges and vertices in the net surface N . *The optimal V-weight net surface problem*, then becomes the problem of finding N such that the total weight $\alpha(N)$ is minimized:

$$\alpha(N) = \sum_{v \in V_N} w(v) \quad (6)$$

Assuming additionally that the cost of each edge $(i_k, j_{k'}) \in E$ can be calculated by a convex non-decreasing function $f_{i,j}(|k - k'|)$ associated with each edge $(i, j) \in E_B$ then *the optimal VCE-weight net surface problem* becomes the problem of finding N such that the total vertex weight and the total edge weight $\beta(N)$ is minimized:

$$\beta(N) = \sum_{v \in V_N} w(v) + \sum_{(i_k, j_{k'}) \in E_N} f_{i,j}(|k - k'|) \quad (7)$$

We will first explain how the optimal V-weight net surface problem can be solved using maximum-flow algorithms. A directed graph $\tilde{G} = (\tilde{V}, \tilde{E})$ is build from G in the following way. The vertices in \tilde{G} directly corresponds to the vertices in G , that is, $\tilde{V} = \{\tilde{i}_k \mid i_k \in V\}$ and the weight of each vertex is found from w as follows:

$$\tilde{w}(\tilde{i}_k) = \begin{cases} w(i_0) & \text{if } k = 0 \\ w(i_k) - w(i_{k-1}) & \text{if } k > 0 \end{cases} \quad (8)$$

The intra column edges \tilde{E}_{intra} are infinite cost directed edges connecting each vertex in a column to the vertex below it:

$$\tilde{E}_{intra} = \left\{ (\tilde{i}_k \xrightarrow{\infty} \tilde{i}_{k-1}) \mid i \in V_B, \{i_k, i_{k-1}\} \in V_i \right\} \quad (9)$$

Given a vertex i_k in one column and the bottom-most of all the connected vertices in an adjacent column j , $j_L(i_k, j)$ then the corresponding vertices in \tilde{G} , \tilde{i}_k and $\tilde{j}_L(i_k, j)$ are connected with a directed edge of infinite cost. We call these the inter column edges \tilde{E}_{inter} :

$$\tilde{E}_{inter} = \left\{ (\tilde{i}_k \xrightarrow{\infty} \tilde{j}_L(i_k, j)) \mid (i, j) \in E_B, i_k \in V_i \right\} \quad (10)$$

Next we need to connect the source node s and the sink node t . Let \tilde{V}^+ and \tilde{V}^- denote the set of vertices in \tilde{G} with non-negative and negative weights respectively:

$$\tilde{V}^+ = \{v \mid \tilde{w}(v) \geq 0\} \quad (11)$$

$$\tilde{V}^- = \{v \mid \tilde{w}(v) < 0\} \quad (12)$$

Then the source and sink node edges, denoted \tilde{E}_s and \tilde{E}_t respectively, are given by:

$$\tilde{E}_s = \left\{ (s \xrightarrow{-\tilde{w}(v)} v) \mid v \in \tilde{V}^- \right\} \quad (13)$$

$$\tilde{E}_t = \left\{ (v \xrightarrow{\tilde{w}(v)} t) \mid v \in \tilde{V}^+ \right\} \quad (14)$$

The intra and inter column edges together with the source and sink node edges define the complete edge set of \tilde{G} , $\tilde{E} = \tilde{E}_{intra} \cup \tilde{E}_{inter} \cup \tilde{E}_s \cup \tilde{E}_t$. An example of such a graph minus the source and sink edges is given in figure 1(c). The minimum cut in \tilde{G} , can be thought of as a segmentation of the object and background vertices, as the vertices in the source and sink sets respectively. The surface N is given by the top-most vertex in each column that is part of the source set.

Obviously the cut cannot include any edges with infinite capacity, and so the found surface must be flat, in the sense that if some vertex is in the source set of the cut, then every vertex below it, in its column, must also be in the source set. It can be argued similarly that it can only follow edges in the graph G .

The solution can be shown to be minimal in $\alpha(N)$ by observing that a net in \tilde{G} defines a closed set of all the vertices "below" it in \tilde{G} . And reversely any non-empty closed set in \tilde{G} similarly define a net in \tilde{G} . The net and the corresponding closed set have the same weight, because of the way the weights in \tilde{G} were assigned. This means that finding the minimum closed set in \tilde{G} yields the solution to the optimal V-weight net surface problem. There is one caveat though, the minimum closed set can be empty, which means that every non-empty closed set must have positive weight. In these cases a simple, so called, *translation operation* can be performed to change the weight of any bottom vertex, say $\tilde{i}_0 \in \tilde{V}$ to ensure that a closed set has negative weight, $\tilde{w}(\tilde{i}_0) := \tilde{w}(\tilde{i}_0) - 1 - \sum_{j \in V_B} j_0$. This can be done, because these vertices are part of any non-empty closed set, and so the operation does not change any minimal non-empty closed set. Next we need to show that a cut in \tilde{G} correspond to a minimum closed set. Let S and T denote the source and sink sets respectively in a finite cost cut in \tilde{G} . Notice that the edges in the cut are given by $\{(v \xrightarrow{\tilde{w}(v)} t) \mid v \in S \cap \tilde{V}^+\} \cup \{(s \xrightarrow{-\tilde{w}(v)} v) \mid v \in T \cap \tilde{V}^-\}$, which leads us to a total cost of the cut of $\sum_{u \in S} \tilde{w}(u) - \sum_{v \in \tilde{V}^-} w(v)$. $\sum_{v \in \tilde{V}^-} w(v)$ is a constant and S is the source set, defining a closed set. This means the cost of the cut is only a constant from the cost of the corresponding closed set and so it must be the minimum closed set.

Next we will describe how the optimal VCE-weight net surface problem can be solved using maximum-flow algorithms. Assuming the directed graph \tilde{G} has been constructed as in the optimal V-weight net surface problem. Let $I(i_k, j)$ be defined

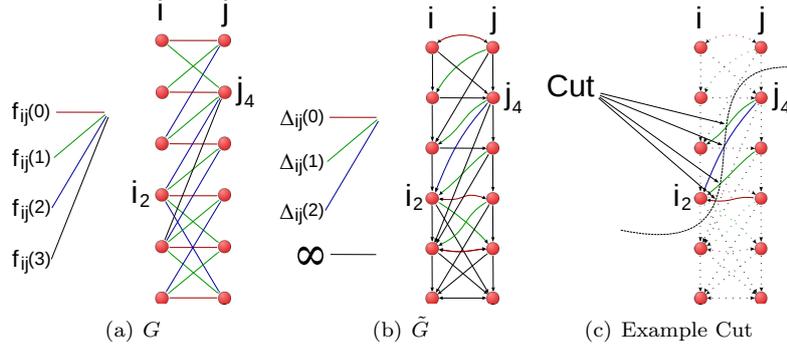


Figure 2. The edge weights of Wu and Chen¹ are illustrated in figure 2(a), how they are transformed to edge costs in the directed graph in figure 2(b) and finally how some of them would be cut if j_4 was found to be part of the source set and i_2 part of the sink set. Edges contributing costs to the cut are coloured, the rest are punctuated.

as before, and assume $I(i_k, j) = (j_p, j_{p+1}, \dots, j_{p+s})$, then define the set of indices in the j -column connected to i_k and on the same level or further down as:

$$P(i_k, j) = \{p \mid p \geq k, j_p \in I(i_k, j)\} \quad (15)$$

then the following defines the cost-penalty edges, E_{cost} :

$$\tilde{E}_{cost} = \left\{ (\tilde{i}_k \xrightarrow{\Delta_{i,j}(p-k)} \tilde{j}_p) \mid (i, j) \in E_B, i_k \in V_i, p \in P(i_k, j) \right\} \quad (16)$$

The capacity of these edges are defined by the following function:

$$\Delta_{i,j}(x) = \begin{cases} f_{i,j}(1) - f_{i,j}(0) & \text{if } x = 0 \\ (f_{i,j}(x+1) - f_{i,j}(x)) - (f_{i,j}(x) - f_{i,j}(x-1)) & \text{if } x > 0 \end{cases} \quad (17)$$

Figure 2 contains an illustration of these edges and how they are contributing to the cost of the cut. The complete edge set in an optimal VCE-weight net surface problem is then given by $\tilde{E} = \tilde{E}_{intra} \cup \tilde{E}_{inter} \cup \tilde{E}_{cost} \cup \tilde{E}_s \cup \tilde{E}_t$. The solution is easily seen to be minimal in $\beta(N)$ from the optimality proof the optimal V-weight net surface problem and the capacity of the cost-penalty edges.

While not all surfaces are flat, many can be transformed into a flat state using a suitable reversible transform. This holds for two-dimensional cross-sections of the airway branches as we explored in reference 9, and even whole branches as was described in references 25 and 24, however not for the complete airway tree as the bifurcation points cause problems.

2.2.2 Smoothness constraints

This section describes the, as far as we know, first practical use of the optimal V-weight net surface problem. We should note, before we go any further, a difference in wording: constraint and penalty are used through out this section and the following, and they carry different meanings. Constraints are used to describe conditions that the solution *must* satisfy, whereas penalties describe conditions that are penalised or less preferred and thus are associated with costs. Constraint edges and penalty edges, implement these conditions and costs.

Smoothness constraints were used reference 24 to force the surface to be $\Delta_x \geq 0$ smooth, by carefully constructing G such that each vertex in one column V_i was (only) connected to the vertices in neighbouring columns at most Δ_x vertical indices above and below it. Let the set of allowed indices be given by $A_k = \{\max(0, k -$

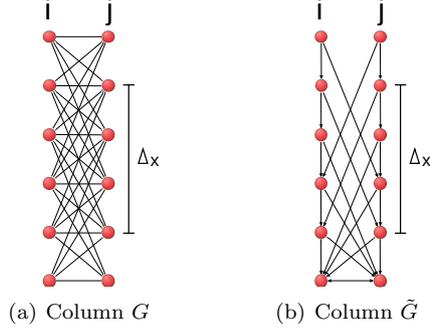


Figure 3. Column in a graph enforcing the smoothness constraints introduced in reference 24, seen in figure 3(b). Feasible surfaces can at most change $\Delta_x = 3$ vertical levels between two columns. The directed graph constructed from G minus the source and sink edges shown in figure 3(b).

Δ_x), $\max(0, k - \Delta_x) + 1, \dots, \min(K - 1, \Delta_x + k)\}$, then these smoothness constraint edges, E_{inter} , are given by:

$$E_{inter} = \left\{ (i_k, j_{k'}) \mid (i, j) \in E_B, i_k \in V_i, k' \in A_k \right\} \quad (18)$$

Example seen in figure 3. This forces any feasible surface to at most change Δ_x vertical indices between two connected columns:

$$(i_k, j_{k'}) \in E_N \Rightarrow |k - k'| \leq \Delta_x \quad (19)$$

The method was used to segment the cylindrical surfaces of human airway segments from pulmonary volumetric CT images by 'unfolding' them into flat surfaces. With the surfaces oriented along the $(x-y)$ -axis, see figure 1(a), x and y axis smoothness constraints, $\Delta_x \geq 0$ and $\Delta_y \geq 0$ were used to guarantee surface continuity in three dimensions. An expert compared the segmentation results in 100 randomly extracted perpendicular cross-sectional images from 317 airway segments with a strictly two-dimensional method and found the vast majority to be superior.

2.2.3 Coupled surfaces

Li et al.²⁵ were the first to show how multiple flat surfaces could be coupled into one maximum flow graph, by combining the intra surface edges of Wu and Chen,¹ with inter surface edges to constrain the surface separation. This coupled surface graph cut algorithm could combine the information of the image at the region of the multiple surfaces and thereby use clues from one surface to help place the other. This is extremely useful for airway wall segmentation, as the two borders are inherently connected.

n proper ordered multi-column graphs G^m , $m \in \{0, 1, \dots, n - 1\}$ are needed to segment n surfaces using this method, all the same size, but they can be associated with different vertex and edge weights. These graphs are then combined:

$$G = \{G^m = (V^m, E^m) \mid m \in \{0, 1, \dots, n - 1\}\} \cup \{E_{sep}^m \mid m \in \{0, 1, \dots, n - 2\}\} \quad (20)$$

Assuming G^m and G^{m+1} are the graphs of two surfaces, whose separation needs to be constrained, let $V_i^m \in V^m$ and $V_i^{m+1} \in V^{m+1}$ denote corresponding columns in each of these graphs respectively, let $i_k^m \in V_i^m$ and $i_{k'}^{m+1} \in V_i^{m+1}$ be vertices in each of the columns and let $\delta_l^{m,m+1}$ and $\delta_u^{m,m+1}$ define a lower and an upper bound respectively on the signed surface separation:

$$(i_k^m, i_{k'}^{m+1}) \in E_N \Rightarrow \delta_u^{m,m+1} \geq k - k' \geq \delta_l^{m,m+1} \quad (21)$$

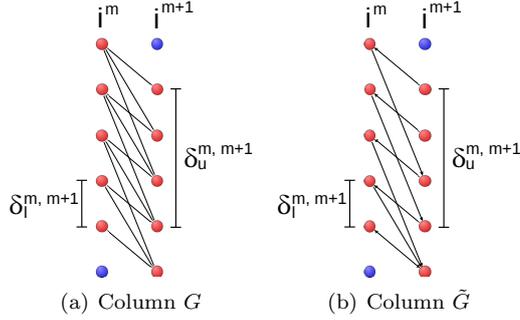


Figure 4. Figure 4(a) shows an example column implementing the separation constraints of reference 25 with $\delta_l^{m,m+1} = 1$ and $\delta_u^{m,m+1} = 3$. Figure 4(b) illustrates the directed graph constructed from G minus the source and sink edges. The blue vertices are deficient nodes.

We will assume without loss of generality that $\delta_u^{m,m+1} \geq 0$, because if this is not the case, then either the constraints represent an infeasible solution or we may simply switch the indices of G^m and G^{m+1} . Let $T_m = \{\delta_{m,m+1}^l, \delta_{m,m+1}^l + 1, \dots, \delta_{m,m+1}^u\}$, and define the set of allowed indices as $A_k^{m,m+1} = \{\max(0, k - \delta_u^{m,m+1}), \max(0, k - \delta_u^{m,m+1}) + 1, \dots, \min(K - 1, k - \delta_l^{m,m+1})\}$ then the surface separation constraints are implemented by the E_{sep} edges:

$$E_{sep} = \left\{ (i_k^m, i_{k'}^{m+1}) \mid m \in \{0, 1, \dots, n - 2\}, t \in T_m, i_k^m \in V_i^m, k' \in A_m^{m,m+1} \right\} \quad (22)$$

The directed graph is then constructed from G as described in section 2.2.1. Note that some vertices in the top and bottom of these columns are not strongly connected and they cannot appear in any feasible solution. Reference 25 calls such vertices deficient nodes. In order to make the graph a proper ordered multi-column graph, such vertices and their incident edges need to be removed. Figure 4 illustrates this graph construction method.

The original paper²⁵ verified the method on phantoms and 3D medical images from CT, magnetic resonance (MR) and ultrasound scanners. Amongst these test were segmentations of the inner and outer airway wall surfaces in 12 in vivo CT scans of 6 humans subjects. The segmentations were conducted on unfolded airway segments. The results were found to be statistically more accurate when compared with previous two-dimensional methods.

The method has later been used for various segmentation problems, such as the segmentation of six⁶³ and seven⁶⁴ intraretinal layer surfaces in 24 three-dimensional macular optical coherence tomography images from 12 subjects also with promising results.

2.2.4 Edge penalties

Even though the optimal V-weight net surface problem has found multiple uses since it was first described, only one paper (to our knowledge) contain a practical description of the usage of the optimal VCE-weight net surface problem. It occurred in reference 56, where it was used in combination with smoothness and separation constraints to segment liver lesions with necrosis or calcification and various other tumors in CT images. The implemented smoothness and separation penalties make it possible to penalise surfaces further from being smooth and close together. We will not describe the graph construction method, as it is the logical extension of the specific optimal V-weight net surface problems, already described in section 2.2.2 and section 2.2.3, to the optimal VCE-weight net surface problem, described in section 2.2.1, instead we included an illustration of the technique in figure 5. The

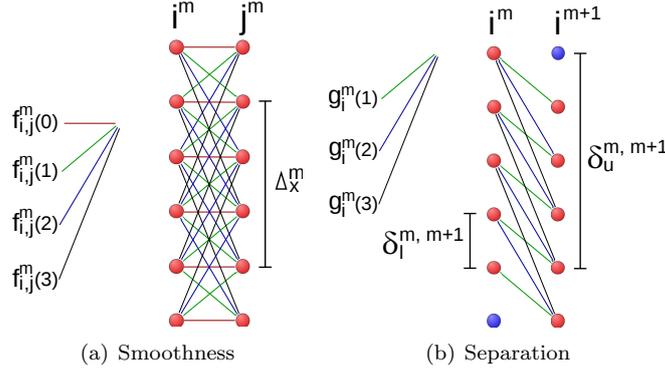


Figure 5. The edge penalties of reference 56 here shown as non-directed graphs with $\Delta_x^m = 3$, $\delta_l^{m,m+1} = 1$ and $\delta_u^{m,m+1} = 3$. Figure 5(a) shows an example of the smoothness penalty with edge costs given by the function: $f_{i,j}^m(x)$ and figure 5(b) shows an example of the separation penalty with edge costs given by the function: $g_i^m(x)$. Deficient nodes are shown in blue.

practical experiment conducted made use of smoothness penalties only and not separation penalties. The edge cost function was chosen as:

$$f_{i,j}^m(x) = x^2, m \in \{0, 1\}, (i, j) \in E_B \quad (23)$$

The results on the synthetic data show that the smoothness penalty enables the method to better cope with noise in the data, and according to the authors the method's capability to simultaneously identify liver tumor and necrosis boundaries is unprecedented.

2.2.5 Segmentation of non-flat surfaces

Li et al.³ expanded the coupled surface methods to be able to handle multiple closed surfaces by constructing each column at surface points and orienting it in the same direction as the surface normal. The method was demonstrated by segmenting a three-dimensional magnetic resonance image of an human ankle and validated against manual segmentations.

The method has later been used to segment the liver in 54 CT images, using a statistical shape model along with an evolutionary algorithm to detect the initial segmentation.⁵ In reference 6 a multi-object extension of the method was used to segment the knee-joint bone and cartilage in 17 three-dimensional MR data sets. In reference 7 the method was used to segment the femoral head, ilium, distal femur and proximal tibia in CT data.

The method is vulnerable to self-intersections due to its reliance on surface normals.

2.2.6 Medial axes columns

Liu et al.⁴ used the medial axes to find the normal direction and column lengths. The medial axis of the initial segmentation is a set of points, each of which has at least two nearest points on the surface. The distance from a surface point to the nearest medial axis point is a conservative bound on the distance one can travel along the normal direction, while keeping the surface point the nearest surface point. This is also known as the local feature size or local thickness. By using this distance as column length, the columns can be guaranteed not to be intersecting. The medial axis is complicated to compute, but approximations can be found using Delaunay/Voronoi diagrams of the surface points.⁶⁵

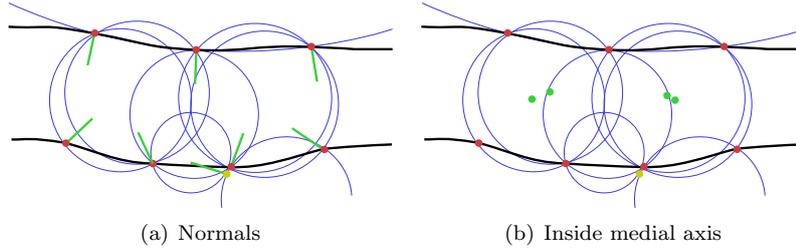


Figure 6. Figure 6(a) show how the largest of the Delaunay balls (blue) incident to surface vertices (red) can approximate normal directions, even in the presence of noise, as illustrated by the yellow vertex. The black lines indicate what the real surface might look like. Figure 6(b) illustrates the inside medial axis points (green) formed using a neighbourhood of 1 nearest neighbours. The distance from each vertex to the nearest of these defines the inner feature size at that specific vertex.

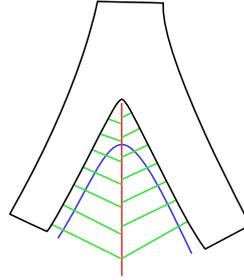


Figure 7. Illustrating the problems of using the medial axis (red) to determine column directions (green) in areas with extreme curvature, such as the branching areas. The black figure represents the initial segmentation, the blue curve represents the sought surface border. The four inner-most columns do not cross the real surface border, which means that the segmented surface will be wrong in these positions.

The direction from the surface point to the largest Delaunay ball incident to it, approximates the normal direction well, if the radius of the ball is large compared to the average distance to neighbouring surface points. Such balls are called polar balls, the concept is illustrated in figure 6(a). Reference 66 explore the error of this approximation, and gives a conservative bound of its size as a function of the sampling density.

The local feature size is more sensitive to the size of the Delaunay balls. Reference 66 show that the medial axis point closest to the surface point can be approximated by the centre of the largest Delaunay ball incident to any of the surface point's nearest neighbours. A conservative bound on the error of this approximation is given. The inner and outer local feature sizes are then approximated at each surface point by calculating the shortest distance to the inner and outer medial axis point sets.

The method was validated on 6 CT scans, by segmenting the airway wall from an initial segmentation of the inner airway border and by segmenting the pulmonary vascular trees. It succeeded in improving on the initial segmentation according to Liu et al.⁴ The method was evaluated further in reference 4 by using two double-wall bifurcating phantoms derived from human in vivo data. It achieved sub-voxel accuracy in the phantoms in both bifurcation/non-bifurcation and carina/non-carina areas.

The method is not without problems. The medial axis can be arbitrarily close to the surface points. This is especially a problem at bifurcations, see figure 7, where the resulting columns can become too short. A minimum column length was used to remedy this problem in reference 4. However implementing this work-around just means that columns gain some length at the risk of self-intersection, negating the idea behind the method in the first place.

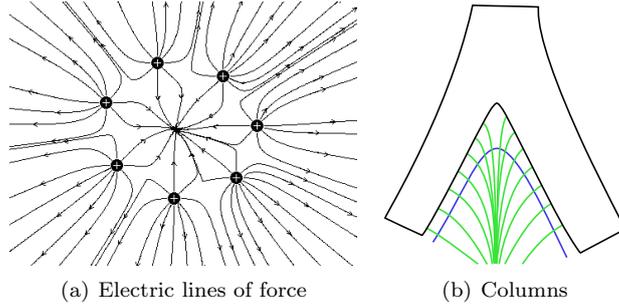


Figure 8. Figure 8(a) show two-dimensional simulated electric lines of force from 7 point charges. Figure 8(b) illustrates the idea behind the use of the non-intersection property of the electric lines of force (green), when used to construct columns. Notice that they all cross the sought surface (blue) and none intersect.

2.2.7 Electric lines of force inspired columns

Recently a paper was published on a novel graph construction technique, which was inspired by the non-intersection property of the electric lines of force.² Instead of using straight columns, the columns follow the electric field direction. Remember Coulomb’s law, says that the magnitude, E_i , of the electric field created by a single point charge, q_i , at a certain distance, r_i , is:

$$E_i = \frac{1}{4\pi\epsilon_0} \frac{q_i}{r_i^2},$$

where ϵ_0 is the electric constant. For a system of discrete charges, E_i , where $i \in \{1, 2, \dots, n\}$, the magnitude E is given by the sum of the individual charges:

$$E = \sum_{i=1}^n E_i$$

The electric lines of force follow the direction of the electric field, the direction within the field, which has the greatest rate of change, see figure 8(a).

Reference 2 assign unit charges to the surface points of the initial model and trace the columns using an electric field with magnitude $E'_i = 1/r_i^4$. This was done to reduce the effect of distant charges and does not compromise the non-intersection property. Figure 8(b) illustrates how the non-intersection property might help for instance in bifurcation areas.

The method was demonstrated using iterative graph searching of a tibial bone-cartilage interface and in a multi-surface graph search segmentation of mutually interacting femoral and tibial cartilage with promising results. The authors commented that the method’s most significant limitation is the relative expensive computation of the electric lines of force.

2.2.8 Cost functions

An integral part of most optimisation problems, is the ability to assign a cost to any given solution. A cost function fills this role by mapping some set of input variables to a cost. Within the subject at hand, this cost should reflect the inverse likelihood that any given input point belongs to the airway wall border. Specifically we need one cost function to so pinpoint the inner airway wall border and another to so specify the outer airway wall border.

Note that some of the specific methods explained in this section, are actually reward functions, in the sense that their returned values are highest at likely edge positions.

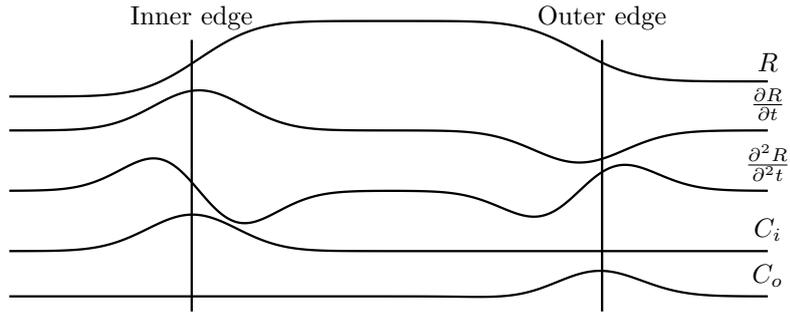


Figure 9. Ray intensity values R , ray derivatives $\partial R/\partial t$ and $\partial^2 R/\partial^2 t$ and the inner and outer cost functions C_i and C_o from reference 9. Notice that the extrema of the first and second order derivatives occur at different sides of the edges.

The matter of constructing such a cost function is inherently an edge detection problem, and many of the methods already described, such as FWHM, phase congruency²¹ or the approaches using models and parameters to approximate scanner specifics^{19,50} could be used with modifications. Classic edge detectors such as Canny⁶⁷ or Deriche⁶⁸ are also good options.

Many methods take the approach of using derivative filters to highlight the edges. For instance reference 3 used the negated gradient magnitude, $-|\Delta I|$, of a Gaussian smoothed image as a bone specific cost function and weighted combinations of the eigenvalues of the Hessian matrix of the image intensity as a cartilage specific cost function. Reference 56 used $-|\Delta I|$ to detect the tumor layer and the first derivative in the column direction to detect the necrosis layer.

Reference⁶ defined the cost function based on the magnitude of the image gradient as $e^{-|\Delta I|/2}$ and modified the values by adding or subtracting constants depending on whether the intensity values were increasing or falling and whether the search was toward cartilage or away from cartilage.

Reference⁷ derived a cost function based on simple thresholding of the image intensities. Costs were simply defined to be low where intensities changed from above to below the threshold.

A common approach is to use the fact that the first and second order derivative edge detectors place their maximums on different sides of the sought edge. Weightings of these can therefore be used to adjust the position of the found edge, in order to match some form of ground truth data.^{4,9,20,22,24,25} For instance reference 4, 22 and 24 used weightings of the Marr-Hildreth and Sobel masks. Reference 20 introduced boundary specific cost functions based on this concept, designed to find either the inner or the outer airway wall border. However the second order derivative used, added positive responses at both borders, leading to situations where the borders could be found on top of each other. We therefore modified the cost functions slightly in reference 9 to remove the second order derivative component at the wrong border. These inner and outer cost functions, C_i and C_o respectively, are given by equations 26 and 27. But first denote the positive and negative parts of the first order derivative in the column direction by $P(t)$ and $N(t)$:

$$P(t) = \frac{\partial R}{\partial t}(t) H\left(\frac{\partial R}{\partial t}(t)\right) \quad (24)$$

$$N(t) = \frac{\partial R}{\partial t}(t) \left(1 - H\left(\frac{\partial R}{\partial t}(t)\right)\right), \quad (25)$$

where $R(t)$ is the intensity of a ray going from the lumen through the wall and into the parenchyma, parameterised by t . H is some continuous approximation of the Heaviside step function. Then the inner and outer cost functions were defined by C_i and C_o using weights $\gamma_i \in [0, 1]$ and $\gamma_o \in [0, 1]$ in the following manner:

$$C_i(t) = \gamma_i P(t) + (1 - \gamma_i) \frac{\partial P}{\partial t}(t) \quad (26)$$

$$C_o(t) = (1 - \gamma_o) \frac{\partial N}{\partial t}(t) - \gamma_o N(t) \quad (27)$$

Figure 9 illustrates the concept.

A cost function based on a combination of first and second order derivative weightings, edge orientation preference and a position penalty term was suggested in reference 25. The orientation and position penalties may incorporate a priori knowledge about the rotation and position of the edge. A cost function tailored to optimal net type columns and based on the piecewise constant minimal variance criterion of Chan and Vese,⁶⁹ was also proposed:

$$C(i_k) = \sum_{j=0}^k (I(i_j) - a_1)^2 + \sum_{j=k+1}^{K-1} (I(i_j) - a_2)^2, i \in E_B, i_k \in V_i, \quad (28)$$

where I is the image intensity function. a_1 and a_2 should ideally be the mean intensity inside and outside the object respectively, but since the object is unknown before segmentation, approximations must be used:

$$a_1(i_k) = \text{mean}(\{i_{k'} \mid 0 \leq k' \leq k\}) \text{ and} \quad (29)$$

$$a_2(i_k) = \text{mean}(\{i_{k'} \mid k < k' < K\}), i \in V_B, i_k \in V_i \quad (30)$$

This cost function definition attempts to minimize the image intensity variance inside and outside the object.

The many coupled surfaces of reference 63, 64 required a lot of different cost function definitions. Combinations of Sobel masks and pixel intensity summation in limited regions were implemented to favour dark-to-light or light-to-dark transitions with dark or light regions above or below the surfaces. A cumulative summation of pixel intensities along the columns were used to discourage finding bright pixels above or below the surface. The Chan and Vese minimal variance term, explained above, was also used for a single surface.

2.3 Airway abnormality measures

This section contains a description of previous airway abnormality measures used to quantify COPD development. We are not aware of any three-dimensional airway abnormality measures, probably due to the fact that practically all previously published studies investigating COPD have been using one- or two-dimensional automatic, semi-manual or manual segmentation methods.^{9, 45, 45, 47, 48, 51, 52, 54, 55, 70, 70-74} The following metrics thus all require the existence of some segmented two-dimensional airway cross-section ideally oriented perpendicular to the airway centreline.

The most commonly reported metrics are the Inner Area (IA) also called the lumen area and the Wall Area percentage (WA%). WA% is the percentage of the airway area that is wall: $\text{WA}\% = 100 \times \text{WA} / (\text{IA} + \text{WA})\%$. Both measures have been consistently correlated with lung function measurements in the past.^{46-48, 51, 71-73} Generally smaller measures of IA and larger measures of WA% are associated with a worse lung function. IA is a problematic measure, given that it contains no normalisation. This means it is likely more sensitive to covariates such as weight,

height and sampling positions. A problem with WA% is that the changes to the measure caused by COPD are relatively small. For instance the difference between the sickest and healthiest are on the order of a few percentage points.^{9,46} This combined with the fact that the analysed airways are close in size to the resolution of the images, means that segmentation errors or even discretisation effects very quickly obscure any change caused by COPD.

Intensity or density based measures introduced recently^{9,47} should be more robust to such errors. This is because they may be more sensitive to the size of imaged structures, when such structures are small compared to the imaging resolution, due to partial volume effects. It has also been theorised that the average density of the airways increase with disease progression, due to mural calcification and fibrosis,⁴⁷ which would further increase their usefulness. Peak Wall Attenuation (PWAt)⁴⁷ and Normalised Wall Intensity Sum (NWIS)⁹ are measures introduced to try to capture such effects. PWAt is defined as the mean of the maximum intensity found within the wall along rays cast 360 degrees around the centre and out. NWIS is a normalised measure of mass changes within the airway wall and is defined as the sum of the intensities plus 1000 within the wall area normalised by the total airway area: $NWIS = \sum_{\mathbf{x} \in \mathbf{WA}} (I(\mathbf{x}) + 1000) / (\mathbf{WA} + \mathbf{IA})$, where \mathbf{WA} denotes the set of points inside the wall. PWAt and NWIS have been found previously^{9,47} to correlate equally or more with lung function, to be more reproducible and have better diagnostic ability than size based measures such as IA and WA%. Higher values of PWAt and NWIS were found to be associated with poorer lung function.

The broncho-arterial ratio is based on the fact that most airways are accompanied by an artery. These two structures have similar diameters in healthy individuals, but in COPD subjects, the lumen narrows. It is therefore possible to normalise the lumen diameter by dividing with the arterial diameter.^{54,55,74} Increasing values of the ratio have been shown to be associated with lower lung function.⁷⁴ A problem with the measure, is that it can be very difficult to automatically identify the accompanying artery.

Having a single number for each CT scan is attractive for diagnostic purposes and it can be achieved by simply averaging all the measurements sampled in some branch generation range.^{9,51} However this introduces errors due to differences in the branches found by the segmentation. PI10 is a measure which was designed to be consistent in the face of such differences. The measure is based on the assumption that there is a linear relationship between the square root of WA and the internal perimeter (PI) of the airway.^{45,70} Using linear regression of the sampled values of WA and PI it is thus possible to calculate what WA would be, if it was sampled at any specific position in the airway. PI10 is the WA so sampled at a PI of 10 mm. We investigated the measure in reference 9 and found the results disappointing. It proved to be the least reproducible of all the examined measures, likely because the use of linear regression makes the measure very sensitive to outliers.

Note that this list is by no means complete, the size changes have been quantified in many more ways. However it does represent some of the most common metrics.

3. METHOD

In this section our approach to airway wall segmentation and measurements will be described. Starting with a short description of the employed airway segmentation method and how additional information, such as the position of the airway centre-lines, the position of the bifurcation regions and the branch generation numbers were generated. Then we will describe specifics of the investigated graph construction methods, including our novel column construction technique with columns following

flow lines. We will describe the used cost functions, the training and evaluation of the methods and the statistical investigations performed.

3.1 Airway segmentation

No specific algorithm can be said to be the best airway segmentation algorithm at the moment. Judging by the Exact'09 study some, such as reference 30 and 38 stand out as being very exploratory and get larger tree lengths and branch counts. Whereas other such as reference 34 are very conservative and thus get very few false-positives, but also very short tree lengths and low branch counts.

The airway wall segmentation will assume some degree of correctness in the initial segmentation, provided by the airway segmentation method, in the sense that only minor offsets are needed to find the correct segmentation. This means that if leaks occur, which usually completely change the segmentation, the resulting airway wall segmentation will likely also be wrong.

If the goal of the combined algorithm is to diagnose a disease such as COPD, whose effects are widespread within the airways but very small in relative size, then a small false positive rate is probably more important than a small false negative rate. For instance, as described, the very commonly reported airway abnormality measure, WA%, only varies a few percent between sick and healthy, but a single leak might change the complete airway volume by many percentage points, completely invalidating the measurement.

Another important point is consistency. Airway abnormality measures are unfortunately not independent from the physical position in the airway in which they are measured and no good method exist, which can compensate or normalise for this dependency. So airways should ideally be resolved to a roughly similar depth and including roughly the same branches in each segmentation, in order to get reproducible measurements.

We used the method of reference 33 to segment the airways. It has a very low false-positive rate, just slightly worse than the most conservative method³⁴ participating in the Exact'09 study.²⁸ Yet the branch detection rate and the tree length are not far from the most exploratory methods.

The method incorporates trained local airway appearance models and uses the fact that an airway is always accompanied by an artery. The orientation of both structures are used as a criterion for region growing. This works well because arteries are easier to detect than airways.

3.2 Tree extraction

There is general agreement¹² that COPD mostly affects the smaller airways, so we need some way of splitting the airways into regions based on their local size. We want this split to be independent of disease progression, in order to not bias the results, so it cannot for instance be based on just the size of the measured structures, as these are known to be affected. An anatomical consistent way to do it, is to simply do it based on the generations of the airway branches.

We therefore extracted the centrelines, branches and generation numbers from the airway tree with a front propagation method, as described in reference 28. Using the initial segmentation and starting in the trachea, the centroid of the front is stored as the branch centreline as it moves down into the bronchi. Bifurcations are detected as the wavefront becomes disconnected upon hitting bifurcation points, see figure 10.

The computed centrelines were also used to extract cross-sections used to build the training and test set mentioned in section 3.5.

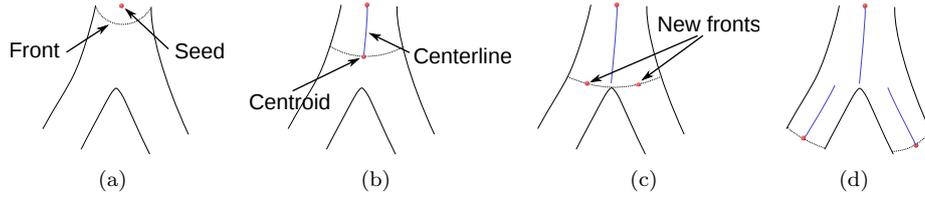


Figure 10. Illustrating how the wavefront moves from the seed in figure 10(a), storing the centreline as the centroid of the wavefront as it moves forward in figure 10(b) until it hits a branching point in figure 10(c), splits and continue in figure 10(d).

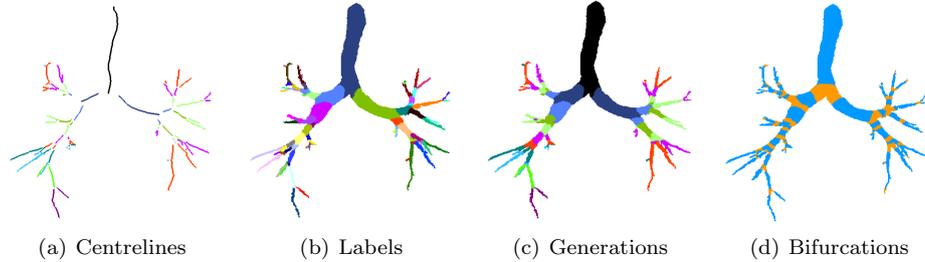


Figure 11. Figure 11(a) shows the computed centrelines, figure 11(b) shows the identified branches, figure 11(c) shows the generations and figure 11(d) shows the bifurcation regions in orange and the non-bifurcation regions in blue.

The measurements were grouped in branches and generations, by simply classifying each voxel in the resulting segmentation by the nearest centreline point's branch and generation information, see figure 11(b) and figure 11(c).

An advantage of the proposed method compared to our previous slice based two-dimensional method is that it enables measurements within bifurcation areas. Some have hypothesized that these areas might be more susceptible to diseases because, as experiments show, a larger amount of inhaled particles are deposited here.^{75, 76} In order to investigate this, the bifurcation regions were defined similarly by the 10% outermost centreline points in each branch, excluding the top most points in the trachea. The rest of the points belong to the non-bifurcation regions, see figure 11(d) for an illustration of this. It should be mentioned that the concept of doing measurements specifically in bifurcation areas, was suggested in reference 8, however we have not seen it used in practice before.

3.3 Graph construction

It is our opinion that smoothness and separation penalties are needed to accurately reflect the continuous and coherent nature of the airway wall borders. Having smoothness and separation penalties makes it possible to find a very non-smooth and well resolved surface in areas where the airway wall borders are clearly defined and still fall back to a smoother surface when they are weakly defined. This is because the optimal surface is found as a trade-off between edge and vertex costs and vertex costs are in most cases inversely proportional to how well the surface is defined.

Smoothness and separation constraints on the other hand seem very unnatural and not very fitting parameters when used to model a soft elastic surface. One could argue, that having a lower bound on the surface separation is useful, because after all, the inner wall border should never cross the outer. However our experience from working with the two-dimensional method,⁹ was that this happened mostly in cases where the initial segmentation was completely wrong, for instance when the initial segmentation had leaked into the parenchyma. Leaks in the initial segmentation often results in errors that are too big for the airway wall segmentation method to

correct, the best one can do therefore is to detect them, and not include them in the measurements. Not having a lower bound on the surface separation enables us to discard a lot of these areas as the places where the inner surface overlaps the outer. Section 3.6 contains a description of how it was used in practice.

Unfortunately the performance of the described optimal net based surface separation penalty is related to $\delta_u^{0,1}$ and $\delta_l^{0,1}$ in such a way that performance gets worse if the surface is constrained less and the performance of the smoothness penalties are similarly related to the smoothness constraints Δ_x^0 and Δ_x^1 . To see why notice that the amount of edges added in two corresponding columns as a consequence of the separation penalty and the separation constraint is given by:

$$\sum_{k=0}^{K-1} (\min(K-1, k - \delta_l^{0,1}) - \max(0, k - \delta_u^{0,1}) + 1) \quad (31)$$

As an example, assume that $K = 10$, $\delta_u^{0,1} = 3$ and $\delta_l^{0,1} = 1$, which results in 24 edges, the same case with $\delta_u^{0,1} = 6$ results in 39 edges, and the completely unconstrained case results 100 edges. The amount of edges added in two neighbouring columns due to the smoothness penalty and the smoothness constraint is given by:

$$\sum_{k=0}^{K-1} (\min(K-1, k + \Delta_x^m) - \max(0, k - \Delta_x^m) + 1), m \in \{0, 1\} \quad (32)$$

This adds up to $3 \times K^2$ edges with two surfaces, no constraints and only penalties. So the conclusion is that constraints are probably needed, if for nothing else than performance reasons. If the edge cost function $f_{i,j}(|k - k'|)$ is linear and does not depend on the indices i and j , then our graph construction technique from reference 9, which we will describe and extend to three dimensions in the following, offers a much simpler solution. It has at most $3 * (K_{mean} - 2)$ edges as a consequence of its smoothness and separation penalties, where K_{mean} is the average number of vertices in each column. It also only has $|V_B| * 2$ source and sink edges compared to the $|V_B| * K$ used in the optimal net type of graph. Both graphs have a similar amount of intra-column edges. It should be noted that many of the optimal net edges do not have to be stored explicitly, because they have infinite cost, however in loosely constrained graphs the non-infinite cost edges will still dominate.

We will specify the graph construction for a single surface, denoted m first. This will then be expanded to multiple surfaces, using a coupled surface graph later. We will use the same technique of specifying the problem in terms of a base graph $G_B^m = (V_B^m, E_B^m)$, which can be mapped into an undirected graph, $G^m = (V^m, E^m)$, with vertices, V^m , and edges E^m , by associating each vertex i in the base graph with a column of vertices in V^m , denoted V_i^m . Lets assume that each i is associated with I_i and O_i points inside and outside the initial segmentation respectively. Combined with the initial surface point, i_0 , these points represents the column V_i^m :

$$V_i^m = \{i_{-I_i}, i_{1-I_i}, \dots, i_0, \dots, i_{O_i}\}, \text{ for } i \in V_B^m. \quad (33)$$

Just as in the Optimal Net type of methods, a column represent the space of all possible solutions, a sought surface point can take. But unlike the Optimal Net type of columns, we will not require that they have equal length. We will deal with different ways of constructing columns in section 3.3.2 and 3.3.3. The columns are connected in a neighbourhood, reflecting the topology of the surface, see section 3.3.1. The vertices, V^m , in G^m are given by:

$$V^m = \{V_i^m \mid i \in V_B^m\} \quad (34)$$

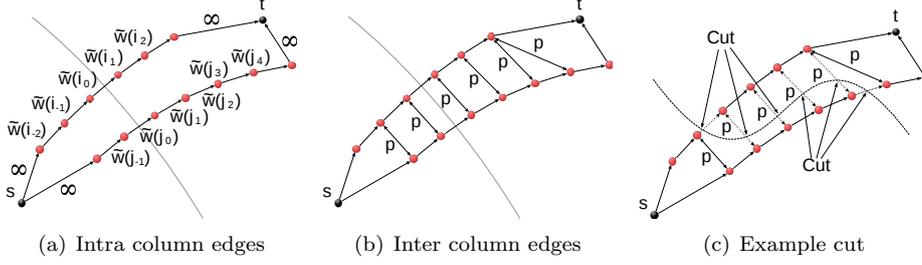


Figure 12. Figure illustrating two connected columns, with edges in black, vertices in red and source s and sink node t and the initial segmentation surface represented by the grey line. The columns illustrated have inside lengths of $I_i = 2$ and $I_j = 1$ and outside lengths of $O_i = 3$ and $O_j = 5$. Figure 12(a) shows the intra column directed edges with capacities given by a translated cost function \tilde{w} . Figure 12(b) illustrates the graph when the inter column edges are added, implementing the smoothness penalty p . Figure 12(c) illustrates how a cut might separate the source and sink nodes. Edges contributing to the cost of the cut are punctuated.

The undirected graph edges, reflecting every possible solution, are given by:

$$E^m = \{(i_k, j_{k'}) \mid (i, j) \in E_B^m, i_k \in V_i^m, j_{k'} \in V_j^m\} \quad (35)$$

The vertices in the graph are associated with weights given by a cost function w^m whose values reflect the inverse likelihood that the sought surface exist at any given position, see section 3.4, and the edges have costs implementing a smoothness penalty, given by the parameter $p_m \geq 0$, note that this corresponds to the optimal net edge cost function $f_{i,j}^m(|k - k'|) = p_m |k - k'|$. Later we will explain how this construction is expanded to the dual surface problem at hand.

G^m is transformed into the directed maximum flow graph by adding $I_i + O_i + 1$ vertices associated with each column V_i^m as follows:

$$\tilde{V}^m = \{\tilde{i}_k \mid i \in V_B^m, i_k \in V_i^m\} \quad (36)$$

Directed edges with infinite capacity are then added from the source s to the inner-most vertices, and from the outer-most vertices to the sink t , lets denote them \tilde{E}_s^m and \tilde{E}_t^m respectively, as follows:

$$\tilde{E}_s^m = \{(s \xrightarrow{\infty} \tilde{i}_{-I_i}) \mid i \in V_B^m\} \quad (37)$$

$$\tilde{E}_t^m = \{(\tilde{i}_{O_i} \xrightarrow{\infty} t) \mid i \in V_B^m\} \quad (38)$$

The intra column edges \tilde{E}_{intra}^m , associated with the cost function, are given by:

$$\tilde{E}_{intra}^m = \{(\tilde{i}_k \xrightarrow{\tilde{w}^m(i_k)} \tilde{i}_{k+1}) \mid i \in V_B^m, \{i_k, i_{k+1}\} \in V_i^m\} \quad (39)$$

Notice that the intra edge costs in the directed graph are equal to the corresponding vertex cost in the undirected graph plus a constant, see equation 50. This is illustrated in figure 12(a). Next we will show how the inter-column edges, implementing the smoothness penalty, are added. First we define $O(i, j) = \max(O_i, O_j) - 1$, $I(i, j) = \max(I_i, I_j) - 1$, for any neighbouring column j and the bounded vertex index $B(i, k)$:

$$B(i, k) = \begin{cases} -I_i & \text{if } k < -I_i \\ O_i & \text{if } k > O_i \\ k & \text{else} \end{cases} \quad (40)$$

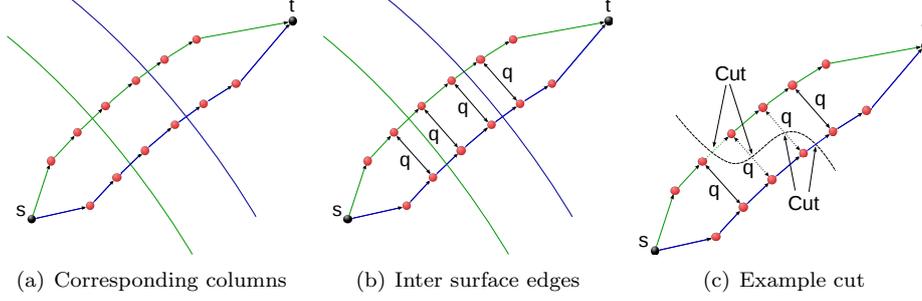


Figure 13. The figures show two corresponding columns. The column marked with green arrows represents the inner column and the blue arrowed column represents the outer column. Figure 13(a) shows the graph before the separation penalty edges are added, the inner and outer surface positions are marked with a green and blue line respectively. Figure 13(b) shows the graph after. The black bidirectional edges are the separation penalty edges. Edges contribution to the total cost of the solution are marked with arrows and punctuated lines in figure 13(c), assuming the solution (cut) is found at the illustrated surface positions.

The inter-column edges \tilde{E}_{inter}^m are then given by:

$$\tilde{E}_{inter}^m = \left\{ (\tilde{i}_{B(i,k)} \overset{p_m}{\leftrightarrow} \tilde{j}_{B(j,k)}) \mid (i,j) \in E_B^m, \right. \\ \left. k \in \{-I(i,j), 1-I(i,j), \dots, O(i,j)\} \right\} \quad (41)$$

Figure 12(b) contains an illustration of these edges. Notice that the weight of any cut in the graph is a sum of the cut intra and inter column edges. Increasing the value of p_m will make the intra column edges relatively cheaper to cut, and will thus favor straighter and smoother surfaces.

Next we will explain how the graphs of the surfaces are coupled, implementing separation penalties $q_m \geq 0$, associated with the separation between surface m and $m+1$. Note that these correspond to the optimal net edge cost functions $g_i^m(|k-k'|) = q_m|k-k'|$. The graphs are similar in that they are generated from the identical base graphs G_B^m and their columns have the same length. We then combine the graphs as previously described:

$$G = \{G^m \mid m \in \{0, 1, \dots, n-1\}\} \cup \{E_{sep}^m \mid m \in \{0, 1, \dots, n-2\}\} \quad (42)$$

where n is the number of surfaces, $n=2$ in the application within this thesis. As before let I_i and O_i denote the inside and outside column lengths. Then the separation penalty edges E_{sep}^m , reflecting every possible surface separation, are given by:

$$E_{sep}^m = \{(i_k^m, i_{k'}^{m+1}) \mid i_k^m \in V_B^m, i_{k'}^{m+1} \in V_B^{m+1}, i_k^m \in V_i^m, i_{k'}^{m+1} \in V_i^{m+1}\} \quad (43)$$

The separation edges in the directed graph \tilde{E}_{sep}^m are given by:

$$\tilde{E}_{sep}^m = \left\{ (\tilde{i}_k \overset{q_m}{\leftrightarrow} \tilde{j}_k) \mid i \in V_B^m, j \in V_B^{m+1}, k \in \{-I_i, 1-I_i, \dots, O_i\} \right\} \quad (44)$$

Figure 13 show the inner and outer columns before and after the inter surface edges are added. Notice that increasing the value of q_m , will make cuts that are separated in the inner and outer columns, progressively more expensive, and will thus cause the method to favour closer surfaces.

The complete directed graph \tilde{G} is given by:

$$\tilde{V} = \left\{ \tilde{V}^m \mid m \in \{0, 1, \dots, n-1\} \right\} \quad (45)$$

$$\tilde{E} = \left\{ \tilde{E}_s^m \cup \tilde{E}_t^m \cup \tilde{E}_{intra}^m \cup \tilde{E}_{inter}^m \mid m \in \{0, 1, \dots, n-1\} \right\} \cup \left\{ \tilde{E}_{sep}^m \mid m \in \{0, 1, \dots, n-2\} \right\} \quad (46)$$

$$\tilde{G} = (\tilde{V}, \tilde{E}) \quad (47)$$

Maximum flow algorithms always find the global optimal solution, so the solution is naturally going to be optimal in the sum of the cost of the cut intra, inter and separation penalty edges. However we want to prove that it is also an optimal solution to the instances of the optimal VCE-weight net surface problem defined by equation 42. Lets first redefine N to include variable length columns:

$$N : \{V_i \mid i \in V_B\} \rightarrow \{i_{-I_i}, i_{1-I_i}, \dots, i_{O_i}\} \quad (48)$$

Note that the original problem set is a subset of this, and so it does not affect the proof.

Notice that every column in \tilde{G} must be cut, since they define a direct line of flow from the source to the sink. In some degenerate cases columns might be cut multiple times, that is the cut surface contains folds (or loops). This can be avoided, by observing, that such a cut would always include more intra column edges than a non-folded cut. It is easy to see that increasing the cost of every intra column edge with a constant, does not affect the choice of non-folded minimum cut, however possible minimum cuts with folds gets progressively more expensive. So all we need to know is how much to increase the cost of each intra column edge to avoid folds in the solution. The reason a minimum cut might contain folds in the first place, is because it could potentially decrease the total cost of the cut inter column and separation edges. A simple bound on this cost is given by:

$$\varphi = \sum_{m=0}^{n-2} |\tilde{E}_{sep}^m| * q_m + \sum_{m=0}^{n-1} |\tilde{E}_{inter}^m| * p_m \quad (49)$$

Note that smaller numbers will do. This can be used to perform an operation, similar to the translation operation described in section 2.2.1:

$$\tilde{w}^m(\tilde{v}) = w^m(v) + \varphi, \text{ for } v \in V^m \quad (50)$$

Now because of the way the graph is constructed, if a vertex is part of the sink set, in a minimum cut, then every vertex following it must be as well:

$$\tilde{i}_k \in T \Rightarrow \tilde{i}_{k+1} \in T, \text{ for } m \in \{0, 1, \dots, n-1\}, i \in E_B^m, \{i_k, i_{k+1}\} \in \tilde{V}_i^m, \quad (51)$$

The conclusion is, that a minimum cut in \tilde{G} is a net surface in G . It can also be seen that any possible net surface in G also defines a cut in \tilde{G} , since it can be thought of as partitioning of the vertices into two disjoint subsets, with the source in one set and the sink in the other. The cost of such a cut in \tilde{G} is given by the sum of the costs of the cut intra and inter column and separation edges. The cost of the cut intra column edges is easily seen to be equal to the vertex cost in the corresponding net surface $\sum_{v \in V_N} w(v)$, plus a constant. To explain why the inter column edge construction corresponds to an edge cost function of $f_{i,j}^m(|k-k'| = p_m|k-k'|)$, lets assume without loss of generality that $k \leq k'$, and notice that if $i_k \in V_N^m$ and $j_{k'} \in V_N^m$ then the following defines the set of inter surface edges $\{(\tilde{i}_{k+1} \xrightarrow{p_m} \tilde{j}_{k+1}), (\tilde{i}_{k+2} \xrightarrow{p_m} \tilde{j}_{k+2}), \dots, (\tilde{i}_{k'} \xrightarrow{p_m} \tilde{j}_{k'})\}$ between i and j that are cut, which obviously leads to $p_m|k-k'|$ as a total cost. The same argument can be given for the separation edges and so the solution is minimal in $\beta(N)$.

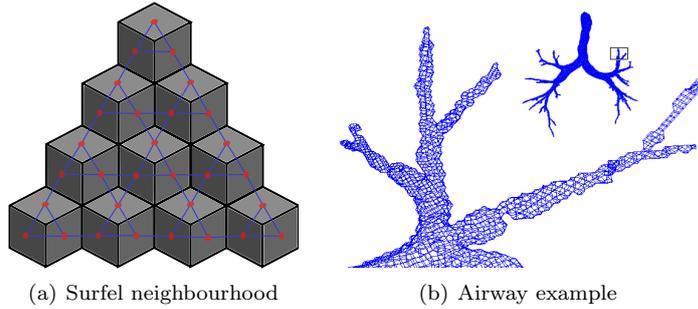


Figure 14. Initial voxel segmentation in shades of grey, surfel centres in red and neighbourhood shown as blue lines in figure 14(a). Figure 14(b) shows an example of an initial airway surface mesh with this neighbourhood, in the area around the left upper lobe bronchus, insert shows position (view is rotated slightly).

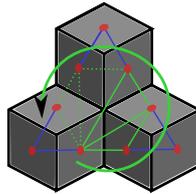


Figure 15. Illustrating how the mesh can be triangulated around a surfel corner.

3.3.1 Surface mesh construction

We need to know where on the surface to sample the columns and how the columns should be linked in a neighbourhood. This is usually done by transforming the initial segmentation to some mesh based representation and using the mesh vertex positions as column positions.^{3,4} For instance reference 3 used an isosurface algorithm to convert an implicit surface into a triangulated mesh, but unfortunately the details of this algorithm are vague. Such a mesh should be as evenly sampled as possible and the neighbourhood should be fairly regular. This is because if the mesh is not evenly sampled, then the algorithm is going to resolve some parts of the surface in more detail than others. If we for instance used the centres of the voxels on the initial surface as our mesh vertices, then we would resolve the larger airways in more detail than the smaller, because the surface voxel count relative to the surface area is less in the smaller airways. The neighbourhood is used to connect the vertices in the graph. These edges are used to constrain the solution, so if it is not regular, then some parts of the surface will be differently constrained than others. The triangulation performed by the popular marching cubes⁷⁷ for instance is not regular and triangulation in general can be very hard, if not impossible, to do without producing irregular meshes, with for instance very thin triangles or a significant degree of variation in the number of edges connected to each vertex.

Our solution to these problems is to use a mesh with vertices at the centre of each surface element, in short: *surfel*, and the neighbourhood is given by the surfel edge neighbours, that is surfels sharing an edge. In this way each surfel has 4 neighbours, see figure 14(a). The surfel mesh also represents an evenly distributed sampling of the discrete surface. In practice, the voxelised continuous objects, we are dealing with, will be slightly oversampled in areas with high curvature or along surfaces, which are far from parallel with the grid axes. Note that this is not a regular mesh, as it was defined in reference 78, as it does not have an equal amount of sides to each face. An improvement, that could make it more evenly constrained, would be to use surfel vertex neighbours as well and weight the edges according to a scheme, which evens out these inequalities, see reference 79. However we decided not to, as the resulting graph would be more complex.

The mesh can be converted to polygons for visualisation or as an intermediate step toward voxelisation, by noticing that each surfel corner is at the centre of some non-planar surface polygon. Each polygon can therefore be obtained by following a clockwise or anti-clockwise shortest path along the neighbourhood edges. A triangle mesh can be constructed from such polygons, by simple triangulation, see figure 15.

There are many approaches to converting polygon meshes to voxel representations.^{80–84} We used a variant of the parity count method of reference 84, in which rays are cast from the centre of each voxel to be classified to some point known to be outside the model. The voxel is classified as inside if there are an unequal amount of triangle intersections along the ray. The advantage of the method is that the mesh triangles can be projected onto a plane defined by the ray direction and candidate triangles can be found with a simple two-dimensional table look-up. It can be further sped up by using graphical hardware acceleration.⁸⁴ The method was crucial in speeding up training, as we did not have to voxelise the complete airway tree in each iteration. Instead it was possible to project the columns to the ray-plane as a pre-computational step and only classify the voxels in the cross-sections. Note that in practice multiple rays are needed as the test does not work when the ray hits a triangle vertex or edge.

3.3.2 Medial axes columns

The medial axes column method⁴ will be included as a reference method in the analysis, because, unlike the electric lines of force inspired method of reference 2, it has been used to segment airways before and therefore provides an interesting baseline method. This section contains details on its implementation.

As mentioned in section 2.2.6 reference 66 describes an algorithm, which uses Delaunay balls to approximate the normals and feature sizes of a noisy point cloud. We have changed the algorithm a bit, since we need the normals pointing in the same direction, as a convention inward was chosen as opposed to outward, and a clear distinction between the inside and outside medial axis. This is very easy, since our initial segmentation is not a point cloud, but a closed surface, with well defined inside and outside volumes. Pseudocode for the algorithms are given in appendix A.1.

The following explains the normal calculation algorithm in words. Given a mesh and some error tolerance level $\tau \geq 0$, the algorithm works by choosing the largest Delaunay ball incident to each vertex in the mesh and using the direction from the vertex to the centre of the ball as the normal direction, flipping its direction if it is pointing outside the mesh. This normal qualifies if the radius of the ball is larger than τ times the distance to the nearest neighbour. If the normal does not qualify, the normal of the nearest neighbour which does, is used.

The feature sizes are calculated from the mesh and some neighbour search count k . The outside medial axis point set is approximated by the centres of the largest outside Delaunay ball amongst each group of k neighbours to each vertex in the mesh. The inside medial axis point set is similarly given from the centres of the largest inside balls. The inside and outside feature sizes at each vertex are then given by the shortest distance to each of these point sets respectively. Choosing larger k will have the effect of reducing noise and leads to larger feature sizes, but reduces the effect of small details.

3.3.3 Flow line columns

The electric lines of force inspired columns of reference 2 are interesting in that they represent a simple and elegant solution to the self-intersection problems of using straight columns. The major disadvantage is that to calculate the electric

field in a specific position, one has to sum up the contributions of all the point charges. This becomes a huge task since these are placed in every surface point. Secondly the initial segmentation could be noisy and it is given in voxel form, which means that some kind of regularisation might be needed. We will attempt to solve this in the following.

Reformulating the field strength as an integral:

$$E(\hat{\mathbf{x}}) = \frac{1}{4\pi\epsilon_0} \int \frac{\rho(\mathbf{x})}{|\mathbf{x} - \hat{\mathbf{x}}|^2} d\mathbf{x} \quad (52)$$

$$= \int Q(\mathbf{x}) \frac{1}{|\mathbf{x} - \hat{\mathbf{x}}|^2} d\mathbf{x} \quad (53)$$

$$= \int Q(\mathbf{x}) R(\mathbf{x} - \hat{\mathbf{x}}) d\mathbf{x}, \quad (54)$$

where $\hat{\mathbf{x}}$ is the position in space to be evaluated and $Q(\mathbf{x})$ the contraction of all the constant terms and the charge density. Meaning the field strength can be formulated as the convolution: $(Q * R)(\hat{\mathbf{x}})$.

Unfortunately the convolution is not defined for $\mathbf{x} = \hat{\mathbf{x}}$ when $Q(\mathbf{x}) \neq 0$, because $R(\mathbf{x}) \rightarrow \infty$ for $\mathbf{x} \rightarrow 0$. A simple way to get around this problem is to introduce a regularisation term α :

$$R'(\mathbf{x}) = \frac{1}{\alpha + |\mathbf{x}|^\beta} \quad (55)$$

, where $\beta = 2$. This has essentially the same values when \mathbf{x} is large, but is defined for all \mathbf{x} when $\alpha > 0$. Increasing the value of α has the added effect of smoothing the result, which is useful if the initial segmentation contains a lot of noise. β can be set to higher values than 2 to decrease the effect of far away charges.

Next we define $Q(\mathbf{x})$ inside and outside the initial segmentation S :

$$Q(\mathbf{x}) = \begin{cases} 1 & \text{if } \mathbf{x} \in S \\ 0 & \text{if } \mathbf{x} \notin S \end{cases}$$

This is different from the surface point only placed charges of reference 2 and it allows us to trace the flow lines consistently through the surface, whereas the point charges of reference 2 introduces extrema, which can cause the flow lines to behave erratically, see figure 8(a). Numerically this also has the advantage that we can simply start following the flow lines from the mesh vertices given from the initial segmentation.

The idea is to trace the flow lines from the mesh vertices inward in the gradient direction and outward in the negative gradient direction. To do this we need the gradient of the regularised electric field strength magnitude:

$$\nabla E'(\hat{\mathbf{x}}) = \nabla \int Q(\mathbf{x}) R'(\mathbf{x} - \hat{\mathbf{x}}) d\mathbf{x} \quad (56)$$

$$= \int Q(\mathbf{x}) \nabla R'(\mathbf{x} - \hat{\mathbf{x}}) d\mathbf{x} \quad (57)$$

which is yet another convolution. We can calculate the gradient of R' analytically:

$$\nabla R'(\mathbf{x}) = \nabla R'(x_0, x_1, \dots, x_n) = \left(\frac{\partial R'}{\partial x_0}, \frac{\partial R'}{\partial x_1}, \dots, \frac{\partial R'}{\partial x_n} \right) \quad (58)$$

$$\frac{\partial R'}{\partial x_i} = -\frac{\beta x_i^{\beta-1}}{(\alpha + x_0^\beta + x_1^\beta + \dots + x_n^\beta)^2} \quad (59)$$

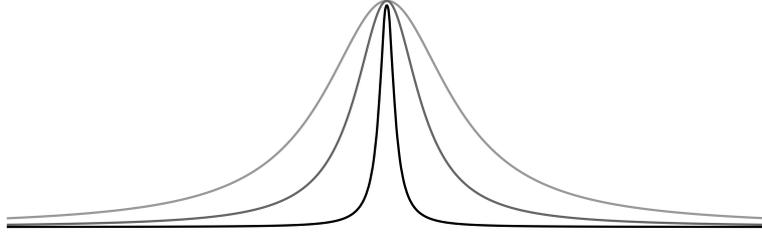


Figure 16. A one-dimensional version of the filter $R'(x)$ with $\beta = 2$ different values of α . Filter values rescaled to have the same maximum size.

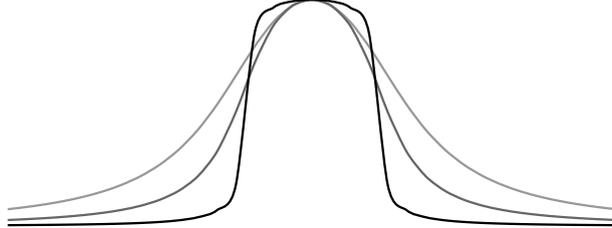


Figure 17. Convolution of the filters shown in figure 16 with $Q(x) = 1$ for $x \in [-1, 1]$ and $Q(x) = 0$ for $x \in]\infty, -1[\cup]1, \infty[$. Filter responses rescaled to have the same maximum size.

To get a sense of what the filter looks like, figure 16 contains plots of R' with different values of α . It is interesting to look at the convolution of the filter with $Q(\mathbf{x})$ as we defined it. For clarity reasons we do this in one dimension, but it can be thought of as what the filter response looks like when sampled as a ray through a branch of the segmentation. Increasing values of α will smooth and regularise the response, which decreases the effect of noise and small errors in the initial segmentation but also removes the fine detail, which can cause topology to change, for instance making two small and close to each other branches merge.

There is no particular reason why the flow lines should be calculated from the electric field strength. Any greatest ascent or descent lines on a continuous function will have the same non-intersection property. An interesting case is the Gaussian $G(\mathbf{x}) = ce^{-|\mathbf{x}|^2/(2\sigma^2)}$ as it is separable, unlike the filter presented in equation 55. Meaning the convolution operation becomes much less expensive, as it can be done by three one-dimensional filters. Another interesting kernel is the Canny-Deriche filter,⁶⁸ given by $D(\mathbf{x}) = -cxe^{-\alpha|x|}$, because it can be implemented in a fast recursive form. The first order derivatives of the Gaussian is what Canny⁶⁷ used in his optimal edge detector, the latter was used by Deriche⁶⁸ for the same purpose. They have been proven to be very useful for edge detection and thus are useful for reducing noise while keeping relevant detail. We have conducted experiments with the Gaussian and the electric field strength kernels within this work.

3.3.4 Numerical integration of columns

We need to find the columns in the inward and outward directions from each mesh vertex, i_0 , by following the flow lines. Parametrising the unknown flow line f by t , we know that:

$$f(0) = i_0 \quad (60)$$

$$\frac{\partial f}{\partial t}(t) = \nabla E(f(t)) \quad (61)$$

This is what is known as an initial value problem or an ordinary differential equation with a specified value. Note that the solution is also called an integral curve, a field line, a trajectory or an orbit depending on the field. We know the initial

condition, the mesh vertex, i_0 , and the first order derivative of the unknown flow line function f at any point in the domain, which is the gradient ∇E and we need to determine $f(t)$ at regular intervals along the arc. So lets say $s(f(t))$ measures the arc length relative to $t = 0$, then we need to determine the values of $f(t)$ where $s(f(t)) \in \{-I_i\Delta s, -(I_i - 1)\Delta s, \dots, 0, \Delta s, \dots, O_i\Delta s\}$. Here I_i and O_i correspond to the number of inner and outer column points in the column V_i . Note the subscript is included, because these are not constant from column to column, as the flow lines can reach local extrema and stop.

Solving the ordinary differential equations exactly is impractical if not impossible, but a range of numerical methods have been developed for these kinds of problems, such as the Runge-Kutta family of methods or linear multistep methods.⁸⁵ They work by approximating the value of the function at any point by its Taylor Series expansion at some point a :

$$f(t) = f(a) + \frac{\partial f}{\partial t}(a)(t - a) + \frac{\partial^2 f}{\partial^2 t}(a)(t - a)^2/2! + \dots + \frac{\partial^k f}{\partial^k t}(a)(t - a)^k/k! + \dots \quad (62)$$

This expansion is truncated up to some order of derivatives k , that is $k = 1$ in the well known Euler method. So assuming we know the value of the function and its k first derivatives at some point a , we can then evaluate the function forward in time to t using a step size of $h = t - a$. In general, smaller step sizes leads to smaller truncation errors. The values of the higher order derivatives ($k > 1$) are unknown, but these can be approximated by taking smaller intermediate Euler method steps or using the values of previous steps.

The regularisation, as we will see later, needs to be very small in order not to destroy small details in the initial segmentation. This leads to numerical integration problems that are inherently unstable, so called stiff problems, with areas where the gradient changes very rapidly. We experimented with different methods such as a simple explicit Euler, midpoint method, explicit and implicit versions of the 2nd and 4th order Runge-Kutta method and controlled the step size by testing whether the flow line ran 'up hill' and limiting the resulting sampling positions to some fraction of the initial segmentation voxel size. However we found that the results were either poor or the running times were slow. Instead we settled on an adaptive Runge-Kutta-Fehlberg (4, 5) (RKF45) method.⁸⁶ It calculates an approximation for the step size error by comparing the results obtained from a fourth and fifth order Runge-Kutta step. The fifth order Runge-Kutta step reuses the calculations from the fourth order step and is thus very efficient. If the error is larger than some tolerance, for instance we used an absolute error tolerance of 0.0001 mm, the step size is reduced and the step is repeated. If the error is less than the tolerance then the step size is increased.

Sampling f at regular arc length intervals can be done in various ways. The method we used in this work was to linearly interpolate the values from samples obtained at the steps naturally taken by the RKF45 method. While questions can be posed to the accuracy of this approach, it is fast and works well in practice because the step sizes are naturally decreased in areas with higher curvature, decreasing the inaccuracies of using linear interpolation. Alternatively the interval search can be incorporated into the step size choice of the RKF45 method, for instance halving the step size when an interval is over-stepped by some tolerance. However we found that this did not lead to significant improvements in accuracy and only made the method much slower.

3.4 Cost functions

We expect the isosurfaces in the electric field strength or Gaussian smoothed images to be good candidates for the inner and outer airway wall border curvature and position, in that they are slightly regularised and displaced compared to the initial segmentation. The integrated columns will run perpendicular to these isosurfaces, meaning we can construct our cost function to look for surfaces running orthogonal to the column direction.

The cost functions of reference 9, described in section 2.2.8, were designed to detect such oriented edges and we decided to use them, since they had worked very well previously. It should be noted, that we could have explored many different options, for instance we found the piecewise constant minimal variance cost function of reference 25, also described in section 2.2.8, interesting. However we postponed such experiments due to the time constraints of this work.

We transformed the values according to the following formulas, since we needed strictly positive values and because C_i and C_o are actually reward functions:

$$w^i(v) = \frac{\max(C_i) - C_i(v)}{\text{mean}(C_i)} \quad (63)$$

$$w^o(v) = \frac{\max(C_o) - C_o(v)}{\text{mean}(C_o)} \quad (64)$$

The derivatives involved, were obtained using simple central differences from cubic interpolated values.

3.5 Training and evaluation

In this chapter the training and evaluation of the examined methods are described. These are the previously developed two-dimensional method,⁹ the method using straight columns, described in section 3.3.2, calculated using k neighbours and an error tolerance of τ , the methods using flow line columns calculated from the potential kernel with regularisation parameters given by α and β , and from the Gaussian kernel with standard deviation of σ , denoted $2D$, $S_{k,\tau}$, $P_{\alpha,\beta}$ and G_σ respectively.

Besides the mentioned parameters: k , τ , α , β and σ all of the methods have inner and outer smoothness penalties, inner and outer cost function derivative weightings and separation penalties, denoted p_i , p_o , γ_i , γ_o and q respectively. Optimal values of each of these parameters need to be estimated, which is a difficult task since no ground truth information is available.

A way to obtain some measure of ground truth is to use phantom or manually annotated data. A phantom is a synthetically constructed object, whose measurements are known. Phantoms used to simulate airways are usually constructed from plexiglass tubes within a medium of similar density to the lungs. By scanning phantoms and using the images, it is possible to calibrate method parameters in order to obtain segmentations more accurately reflecting real-life dimensions.

Unfortunately the phantom data available consisted of straight tubes with very little resemblance to actual human airways. We felt this was unacceptable given that the parameters amongst other things control smoothness. A possible option was to train some of the parameters on the phantom and train the smoothness penalty parameters on manually annotated data. However this was ruled out for complexity reasons and because it was not known how mutually independent the parameters would be.

Instead manually annotated data was chosen, as it would more accurately reflect the natural variations present in the human lungs. It should be noted that other

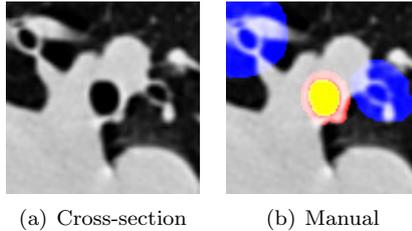


Figure 18. Figure 18(a) shows an extracted cross-section and Figure 18(b) shows the lumen area (yellow), the airway area (red+yellow) and the excluded parts near other airway branches (blue) manually marked.

studies have examined the accuracy of manual segmentations and found them to be problematic.⁵² This is probably true, if one only considers the raw accuracy of a single segmentation. For instance humans will likely overestimate the size of small scale structures because of partial volume effects, just as the FWHM method does. There are however differences in the way that automatic and manual segmentation processes work that make them relevant. For instance looking at segmentation results of our previously published two-dimensional method⁹ made us realise that, while it worked very well in most cases, sometimes it broke down completely. For instance segmenting vessels, parts of other airways or a combination of other nearby structures as part of the airway. Knowledgeable humans tend to not make these mistakes. It is our opinion that limiting such errors is still more important than getting the segmentation just right. This is probably even more true, when the application is diagnosis of a disease such as COPD, which affects more or less the complete lung, but mostly with very subtle local changes to the airways. In such cases a consistent over- or underestimation will likely still catch the changes, but complete segmentation errors can drown the signal in noise.

3.5.1 Manually segmented cross-sections

361 manually segmented cross-sections were available to us from an earlier project,⁹ which we then extended to 649 images taken at random places in the airways from 15 subjects in total. The cross-sections are extracted as explained in reference 9, which means that the cross-section is centred on some airway and lies perpendicular to the airway centreline, see Figure 18(a). Sometimes these images contain more than one airway branch. Such other branches were marked and excluded from the analysis. We did this for two reasons, firstly because manually segmenting branches that are not viewed at a perpendicular angle can be very difficult, and secondly we wanted the results to be comparable with that of the strictly two-dimensional method of reference 9, which would only segment the centred airway, see figure 18(b).

So let M_l , M_a and X denote the manually marked lumen, airway and excluded parts respectively. Similarly let A_l and A_a denote the parts segmented by the algorithms. We then evaluate the correctness of a segmentation using the relative area of overlap outside the excluded area, Φ as follows:

$$\Phi(M_l, M_a, A_l, A_a, X) = \frac{1}{2} \left(\frac{|(M_l \cap A_l)/X|}{|(M_l \cup A_l)/X|} + \frac{|(M_a \cap A_a)/X|}{|(M_a \cup A_a)/X|} \right) \quad (65)$$

The 649 images were split randomly into a training and a test data set, consisting of 329 and 319 images from 8 and 7 subjects.

3.5.2 Algorithm

The quick and dirty binary search algorithm described in reference 9 was re-used to search the parameter space for optimal choice of parameters on the training data

set. Given initially randomised parameters, error tolerances for each parameter and left and right search interval limits. The algorithm iterates through the parameters continuously halving the search intervals around some current best guess for a given parameter until the difference is less than the error tolerance for that parameter or a better guess is found. It then tries the next parameter, until no parameters get updated with better guesses. The algorithm was run many times to make sure that it did not just find a local minimum. The pseudocode is given in appendix A.2.

3.5.3 Practical issues

Training is a very time consuming task. For this reason it was necessary to move the computation of the mesh and columns outside the iterating loop. Meaning it was not time-wise acceptable within the framework to train the values of k, τ, α, β and σ . Instead the values of these were chosen largely based on trial and error by running the training algorithm with different values, combined with random initial values of $p_i, p_o, \gamma_i, \gamma_o$ and q and by using visualisations of the generated columns and normals. This was repeated until we were confident that optimal choices were found. β was fixed to 2 and 4 because we wanted to investigate columns based on the electric field strength and columns similar to those in reference 2. Section 5.2 describe some of the consequences and advantages of these parameter choices.

3.6 Leak detection

False positives within the initial segmentation, meaning structures segmented as airway, but which in reality are something else, are a big problem for the method. It is not designed to correct major errors, mainly because the graph construction method assumes the sought surfaces to have roughly the same orientation as the isosurfaces in the convoluted images. Leaks in the initial segmentation, as described in section 2.1, are the most frequent cause of such errors.

A simple leak detection method was implemented, inspired by our work in reference 9, which showed that the inner airway wall border being segmented outside outer airway wall border, was a good indication that the initial surface was wrong. We looked at segmentation results of our new methods on a few scans, which had leaks, and came up with the following leak measure Ψ :

$$\Psi(A_l, A_a) = \frac{|A_l/A_a|}{|A_a|} \quad (66)$$

We found higher values of Ψ to often be associated with leaks or wrong segmentations and so it was used to remove suspect branches, see section 5.6.1. It should be mentioned that this, was not intended as being anything more than a quick fix. Future studies should attempt to quantify the problem and perhaps come up with some measure of segmentation reliability in specific branches or areas, see section 8.1.

3.7 Measurements

This section contains a description of the airway abnormality measurements conducted with the use of the airway wall segmentation methods.

As described in the introduction, many different airway abnormality measures have been developed. Unfortunately all of them are two-dimensional, meaning they assume the existence of some two-dimensional cross-sectional slice of the airway, ideally lying perpendicular to the airway centreline. This is a very natural presupposition given simpler previous segmentation methods, which also often did the segmentation in such cross-sections. However this is not as straightforward to do when

the airway has been segmented in its entirety, because as explained, these cross-sections will then often include other segmented airways. Such airways often touch the central airway and it may not be easy to automatically decide which is which, leading to measurement errors. Of course there are workarounds to such problems, as for instance being very careful not to extract cross-sections near bifurcations or implementing various schemes to detect other airways in the cross-section. However it is our opinion that a better option would be to simply develop new measures, which can take advantage of the three-dimensional nature of current segmentation methods. Having the complete three-dimensional segmentation of the wall, also allows one to do more advanced measurements such as measurements on curvature, changes in thickness along the surface etcetera.

A very simple option, and one we chose to follow within this thesis, is to simply extend the existing measures to three-dimensions. IA and WA% can be very easily extended to corresponding volume measures, we will refer to the three-dimensional equivalents as the Inner volume (IV) and the Wall Volume percentage (WV%). The airway volume, or the sum of IV and the Wall Volume WV, IV+WV were included as well, as it gives us information on the reproducibility of the segmentation of the outer airway wall border. We also included two simple density based measures, the Mean Inner area Density (MID) and the Mean Airway Density (MAD), computed as the average density of the lumen and of the lumen plus wall respectively. They have the advantage that they can be computed from a volume as well as a cross-section. MAD is very similar to NWIS used in reference 9, but can be measured in simple physical units, such as HU or g/L. MID is an experiment, which because it measures density of what is inside the airway, might pick up the increased mucus production associated with COPD airways.

For diagnostic purposes having a single number for each CT scan is attractive. In reference 9 this was achieved by simply averaging all the cross-sectional measurements within generation 3 – 11. This is obviously not possible with the three-dimensional equivalent measures. Instead one might average the measures over the branches they were collected in. However this results in smaller branches being weighted similar to larger branches, which is problematic as leaks from the initial segmentation often results in many smaller incorrect branches. Instead we computed the measures based on volume. So let W denote the part of the airway tree to be investigated, this could for instance be generation 3–11 or the bifurcation regions determined as described in section 3.2. Then the measures were determined as:

$$IV = |W \cap A_l| \quad (67)$$

$$WV = |W \cap A_a/A_l| \quad (68)$$

$$WV\% = 100 \times WV/(IV + WV)\% \quad (69)$$

$$MAD = \sum_{\mathbf{x} \in W \cap A_a} I(\mathbf{x})/|W \cap A_a| \quad (70)$$

$$MID = \sum_{\mathbf{x} \in W \cap A_l} I(\mathbf{x})/|W \cap A_l|, \quad (71)$$

where I is the image intensity function.

3.8 Statistical investigations

The optimal set of parameters of each of the investigated methods were evaluated using the test data set and equation 65. Results of these were compared using a two-sample Student's t -test for equal means with unequal variances. The variances were compared using a two-sample F -test for equal variances.

In order to establish that the developed method actually is a better algorithm for diagnosing COPD than the previously developed two-dimensional algorithm, a series of large scale tests were performed involving the data from the DLCST. The goals were to investigate whether the use of it resulted in more reproducible measurements, measurements that correlated more with lung function and whether the diagnostic ability of the measurements were better.

Reproducibility was quantified using Pearson product-moment correlation coefficients of repeated measures. Fisher's z transformation was used to compare these correlation coefficients for any statistical difference.

Spearman's rank correlation coefficient was used to investigate the association between the measures and lung function, quantified as FEV1 (% pred) and the SGRQ components. These correlation coefficients were not compared statistically, as we are not aware of any such test.

The AUC values, used to investigate diagnostic ability of the measures in each disease stage, were statistically compared using a DeLong test. Note that only the three first stages were included, as there was only a few stage 4 subjects in the data set.

Only statistical significant correlation coefficients are shown ($p < 0.05$).

4. MATERIAL

The data used in this thesis come from the DLCST and the CBQ study.

4.1 DLCST

DLCST is a Danish randomized lung cancer CT screening trial, involving 4104 smokers and previous smokers with a history of at least 20 pack years and of age 50 to 70 years. We only use data from the screening arm of the trial, which at the baseline meant a total of 2052 subjects. These subjects, with the exception of those that later dropped out, had CT scans and lung function tests performed yearly within a 5 year period. At the time of writing this thesis, the trial database was still being finalised and so it is not based on the complete DLCST data set, but instead a subset of what was available, see section 4.3.

The trial images were obtained using a Multi Detector CT (MDCT) scanner (16 rows Philips Mx 8000). The scans were taken after full inspiration and including the entire ribcage and upper abdomen with a low dose (120 kV and 40 mAs). 1mm and 3mm thick slice images were taken and reconstructed using hard and soft kernels respectively. Only the thin slice images were used in this work. The in-slice resolution of the thin slice images varied slightly, but was most commonly $0.78125 \text{ mm} \times 0.78126 \text{ mm}$

Spirometry or lung function test were performed according to the recommendations of the European Respiratory Society.¹⁰ In this work values of Forced Expiratory Volume in one second (FEV1) expressed as a percentage of the predicted value (FEV1 (% pred)), calculated from subject age, sex and height⁸⁷ and Forced Vital Capacity (FVC) are used. Briefly, FVC is the volume of air that can forcibly be blown out after full inspiration, whereas FEV1 is the amount of air that can be forced out of the lungs in one second. The standard way of diagnosing COPD uses FEV1 (% pred) and the ratio of FEV1 to FVC.¹¹ Table 4.1 shows how the 4 stages of the disease are determined.

Severity of COPD	Lung function test values
Stage I: mild	FEV1/FVC < 0.70 and FEV1 (% pred) \geq 80 %
Stage II: moderate	FEV1/FVC < 0.70 and $50\% \leq$ FEV1 (% pred) < 80%
Stage III: severe	FEV1/FVC < 0.70 and $30\% \leq$ FEV1 (% pred) < 50 %
Stage IV: very severe	FEV1/FVC < 0.70 and FEV1 (% pred) < 30 % or FEV1 (% pred) < 50% and chronic respiratory failure.

Table 1. The standard way of determining COPD disease stages.¹¹ For a definition of chronic respiratory failure see reference 11.

4.2 CBQ

The CBQ study was launched as a longitudinal follow-up study of COPD disease progression within the CT screening arm of DLCST trial.⁸⁸ As a result a large amount of the subjects within the DLCST trial have completed the SGRQ. The SGRQ is a standardized self-completed questionnaire for assessing impairment of health and perceived well-being in subjects suffering from airway diseases. Answers to the questions are weighted and three component scores are calculated: symptoms - reflecting subjects' perception of their respiratory problems, activity - measuring disturbances to subjects' daily activity, and impacts - which covers various disturbances of psychological and social function. Additionally a total score can be produced.⁸⁹ SGRQ was designed to work with asthma and COPD but has been validated for use with multiple other respiratory diseases.⁹⁰⁻⁹² Significant association between the SGRQ scores and spirometry measures FEV1 and FVC, the 6 minute walking test, anxiety, depression and dyspnoea have been found.⁸⁹ Additionally each subject was asked to rate their current state of health, as either: very good, good, adequate, bad or very bad. This question is not part of the SGRQ, but was asked each subject when they completed the questionnaire. The SGRQ data provide a much needed independent assessment of airway abnormality measures. That is, it is possible to examine whether they provide additional information not present within simple spirometry measures.

4.3 Incomplete data

Processing all the available data was infeasible within the given time frame. As mentioned the resolution of images in the data set varied, different scanner intensity cut off values were used to sample the data, for instance some used -1000 HU and others -1024 HU, and while these issues are minor we felt it was safer to work with a more uniform data set until the effects of such differences could be completely determined. We therefore based our data on the images with a resolution of $0.78125 \text{ mm} \times 0.78126 \text{ mm} \times 1 \text{ mm}$ and a cut off value of -1000 HU. Repeated measurements were needed since we wanted to assess reproducibility, so we further limited the data set to those subjects having had repeated spirometry and scans taken within a delta of maximum two years. The closest in time were chosen if a subject had more than two. We will refer to the first of these as the baseline and the second as the followup. The total number of paired scans and lung function measurements was $2 \times 1,256$.

Because the CBQ study was launched later, even less data was available. The total number of paired scans, lung function measurements and SGRQ data was 690.

Note that a statistical investigation of the effects of this sub-sampling of the data should be performed. But this is not possible since we do not have access to the complete data set. We expect any possible publications of the results of the work within this thesis to either be based on the complete data or contain such an investigation.

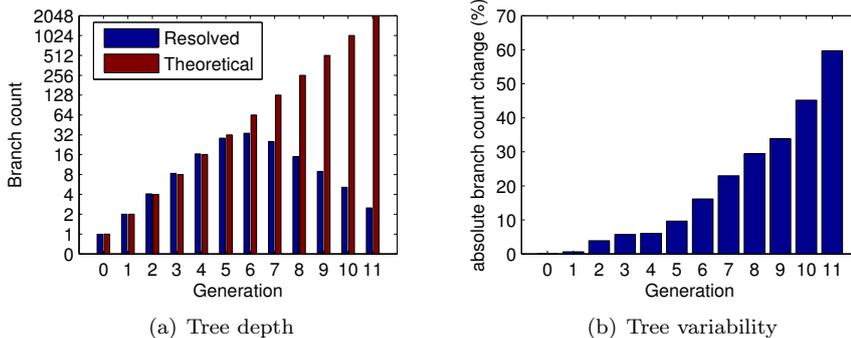


Figure 19. Figure 19(a) demonstrates the ability of the airway segmentation algorithm to segment the airway tree and the tree extraction algorithm’s ability to correctly assign generation numbers, by comparing the theoretical branch count in each generation with the found number of branches. Figure 19(b) shows the variability of the resulting trees, quantified as the average absolute percentage change in branch count in each generation from baseline to followup scan in each generation.

5. RESULTS

5.1 Airway segmentation and tree extraction

The airway segmentation method and the tree extraction algorithm are some of the most deciding factors in the quality of the airway measurements, since they determine which branches gets included in the analysis and which do not. We have reasonable confidence in the algorithms, because as mentioned the airway segmentation algorithm was evaluated as part of the Exact’09 study and the tree extraction algorithm was used in the same study on multiple airway segmentation algorithms.²⁸ However something that was not assessed was the ability of the algorithms to consistently segment, extract branches and assign generation numbers to a reasonable depth. We therefore chose to plot the average number of found branches in each generation against the theoretical maximum number of branches in each generation. Assuming the airway branch count in each generation can be described by 2^g , where g is the generation number. Note that correctness of this assumption hinges on the definition of what constitutes a branch and a bifurcation, however the comparison with a binary tree is not unreasonable.¹⁸ Figure 19(a) shows how the found branch number very closely follows the theoretical maximum number until generation 5, where a total of 96% of the branches are found. Then it drops off dramatically to 74% at generation 6, 47% at generation 7 etc. Figure 19(b) shows the average absolute percentage difference in branch count in each generation when comparing baseline scans with followup scans. The plot shows for instance that at generation 5 the followup scan has on average 9.6% more or less branches than the baseline scan. This number rises steadily to 60% at generation 11.

5.2 Parameters

The purpose of this section is to show the effects of different algorithm parameter values and in the case of the parameters that were not trained, argue why we used the values we did.

5.2.1 Graph resolution

We based our experiments on graphs with a resolution of 0.5 mm, meaning the initial segmentation surface and the columns were sampled using a 0.5 mm spacing. We did experiments with other resolutions and our conclusions were that lower resolutions lead to worse results. Higher resolutions also did seem to result in even better results, however computation times were found to be unacceptable. Future

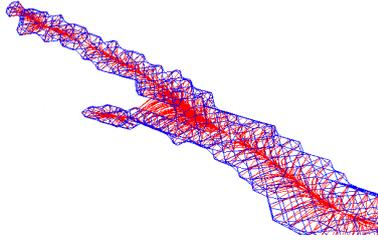


Figure 20. Illustrating why too much regularisation becomes a problem. The blue lines represent the mesh within a small branch of the airway tree. The red lines are the inside part of the columns. Notice how they move outside the initial surface where the branches are closest. This is because the regularisation has caused the extrema normally found inside the branches to start to merge.

k	Inner column length (mm)	Outer column length (mm)
2	0.6	1.1
4	0.7	1.1
8	1.0	1.5
16	1.5	2.4

Table 2. Average column lengths as a function of k .

studies could work on optimising the implementations in order to improve on this, see section 8.2.

5.2.2 Flow line regularisation

The initial results of running the training algorithm with different values of α and σ indicated that only minimal amounts of regularisation were needed. Figure 20 shows an example of what happens, when too much regularisation is used. In fact when we experimented with increasingly smaller values of σ numerical errors proved to be an issue before the lack of regularisation did. Numerical errors result in shorter columns. This is because when following the flow lines away from the surface points, the gradient flattens and at some point, when it approaches machine precision, the RKF45 method hits its error limit and stops the integration. We decided that in order to be able to properly segment every possible airway, the inside and outside column length should be able to reach at least 4 mm. For the Gaussian kernel this meant that we could not choose a σ smaller than 0.45 mm. The choice of α was not as influenced by column length, this is because, α becomes less of a regularisation factor than β for larger values of $|\mathbf{x}|$. We ended up using an α of 10^{-4} , as this would have practically no regularising influence on the filter response, but still meant the filter was defined for $|\mathbf{x}| = 0$.

5.2.3 Normal error tolerance and medial axes neighbours

τ was set to 2.5, which meant that roughly 3% of the normals did not qualify. We could see very little change in the quality of the normals by changing it, besides the negative effect caused by too few normals qualifying. Choosing the right k proved to be a trade-off between column length and self-intersections. This is because small scale medial axis detail is removed as k is increased, which causes the average distance to the medial axis to grow. The risk of self-intersections grows as well because the normals no longer necessarily point in the direction of the nearest medial axis point. Table 2 show the inner and outer column lengths for different values of k . In order to limit the number of self-intersections, a value of 16 was used as k in the rest of the experiments. We think it is the smallest k with an acceptable column length.

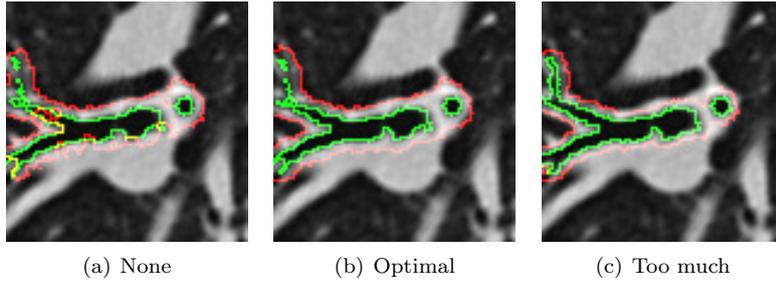


Figure 21. Figure illustrates the segmentation results near a bifurcation area with different amounts of smoothing as determined by the parameters p_i and p_o . Figure 21(a) corresponding to both being set to 0. Figure 21(b) shows the optimal amount as determined by the training algorithm. Figure 21(c) illustrates the problems with very high values.

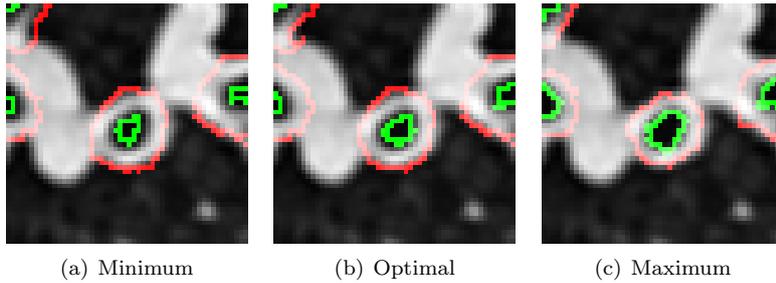


Figure 22. Figure illustrates the segmentation results with different values of γ_i and γ_o . In figure 22(a) the values are 0, in figure 22(b) they are optimal as determined by the training algorithm and in figure 22(c) they are 1.

5.2.4 Derivative weighting, smoothness and separation penalties

Figure 21 illustrates what happens with the segmentation result, when the smoothness parameters are changed. Zero smoothing is essentially the same as running a one-dimensional method and the wall borders become very inconsistent and jagged and they even switch places in some spots. Too much smoothing causes the borders to follow the isosurfaces defined by the initial segmentation too closely, which is a problem especially for the outer border. It can be seen, in figure 21(c), that it follows the curvature of the lumen much more than in the other two images.

Figure 22 shows a cross-section of the segmentation result obtained with different cost function derivative weightings. Notice how the wall shrinks as the first order derivative becomes dominant.

Figure 23 illustrates how the separation penalty can help by prioritising airway wall borders closer together. The specific example shows how the method cannot separate the artery from the outer airway wall border without using a separation

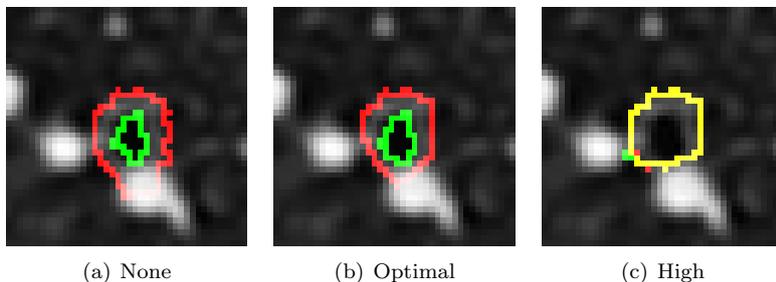


Figure 23. A cross-section in an area benefiting from the separation penalty. Figure 23(a) was obtained with $q = 0$, figure 23(b) with the optimal value as determined by the training algorithm, and figure 23(c) demonstrates the problems with too high values.

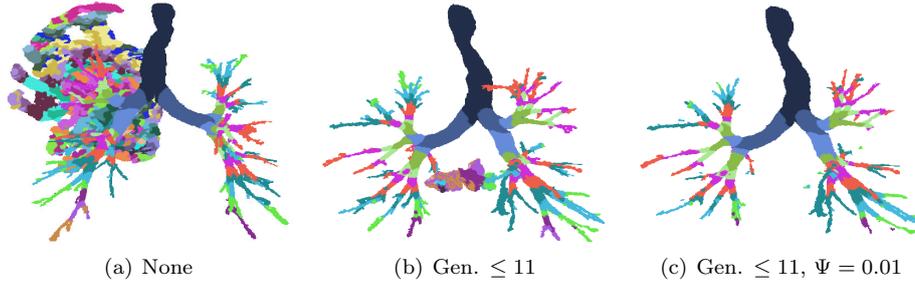


Figure 24. The leak detection method of section 3.6. Images show the lumen resulting from the $G_{\sigma=0.45}$ method coloured by generation numbers. Figure 24(a) contains all the branches obtained from the initial segmentation, figure 24(b) only generation 0-11 and figure 24(c) only the branches in generation 0-11 satisfying $\Psi < 0.01$. Note that the viewing-angle of the left-most image is different from the centre and right-most, as the leak would otherwise block most of the airways.

Method	p_i	p_o	γ_i	γ_o	q
$2D$	0.218	0.267	0.785	0.311	0.00569
$S_{k=16, \tau=2.5}$	0.119	0.458	0.996	0.555	0.000
$P_{\alpha=10^{-4}, \beta=2}$	0.111	0.732	0.991	0.379	0.0394
$P_{\alpha=10^{-4}, \beta=4}$	0.111	0.892	1.00	0.264	0.0342
$G_{\sigma=0.45}$	0.119	0.141	0.992	0.594	0.0364

Table 3. Optimal parameters obtained during training of each of the methods. Note that the values of the $2D$ method are not directly comparable to the methods implemented as part of this thesis, because of implementation differences.

penalty. It also demonstrates how the value must be determined just right, as a too high setting can cause the airway wall borders to be found on top of each other.

5.2.5 Leak detection

Figure 24 illustrates the effect of the simple leak detection method, described in section 3.6. Notice how the initial segmentation has leaked into the parenchyma in figure 24(a). The leak follows the side of the lung causing its dome-like appearance. Figure 24(b) reveals that just limiting the generations to less than 12, results in a much better looking airway tree, but a leak can still be spotted bottom-centre. Using the leak detection method with a Ψ of 0.01 results in the airway tree in figure 24(c). The bottom-centre leak is gone, but so are some of the smallest branches. $\Psi = 0.01$ was chosen based on such examinations of a few segmentations with leaks, and it was used in the rest of the experiments in this thesis.

5.3 Training

Table 3 show the optimal parameters obtained in the training of each of the methods. Most of them are relatively similar, which is an indication that this way of training the parameters, makes sense. There are however some notable differences as well. For instance $S_{k=16, \tau=2.5}$ have 0 separation penalty and the outer smoothness penalty varies relatively much from method to method. To make sure that these choices were not the result of some local minima and associated with real differences in the methods, every method was trained again with the optimal parameters of the other methods as initial values, and this did not lead to large changes in the choice of parameters. We have no explanation for why the $S_{k=16, \tau=2.5}$ method apparently works the best with no separation penalty, however we think that the outer smoothness penalty and the regularisation used to calculate the columns are related, in such a way that more regularisation seem to be associated with a choice of higher smoothness penalties. To understand why this might be, it is important to understand that the initial segmentation is actually a segmentation of the lumen.

Part	$2D$	$S_{k=16,\tau=2.5}$	$P_{\alpha=10^{-4},\beta=2}$	$G_{\sigma=0.45}$
Columns (s)	-	11.7	529	109
Minimum cut (s)	-	4.11	27.2	9.89
Total (s)	149	73	639	183
Memory (GB)	0.31	1.2	2.9	1.8
Graph vertices ($\times 10^6$)	-	3.6	9.0	5.4
Graph edges ($\times 10^6$)	-	25	63	38

Table 4. Time to do the column construction process. For instance in the $S_{k=16,\tau=2.5}$ case this would be the calculation of the medial axes and normals. In the $P_{\alpha=10^{-4},\beta=2}$ and $G_{\sigma=0.45}$ cases this would be the convolution with the kernels and flow line computations. Listed are also the time to do the minimum cut calculation and the total, which includes construction of meshes and voxelisation of the segmentation results.

$2D$	$S_{k=16,\tau=2.5}$	$P_{\alpha=10^{-4},\beta=2}$	$P_{\alpha=10^{-4},\beta=4}$	$G_{\sigma=0.45}$
0.124 ± 0.091	0.142 ± 0.077	0.120 ± 0.079	0.123 ± 0.080	0.111 ± 0.059

Table 5. The results of different methods and kernels on the test data set. Mean \pm standard deviation as given by equation 65.

The orientation and curvature of the lumen surface is different from the outer airway wall border near bifurcations. At bifurcations, the lumen becomes separated into two new branches before the airway walls do, but increasing the regularisation, causes such lumen to 'melt' together, making the orientation and curvature of these surfaces more like the outer airway border, which increases the usefulness of the outer smoothness penalty.

5.4 Performance

The programs were run on an Intel I7-920 at 2.67 GHz with 12 GB of ram. It should be noted that implementations were not particularly optimised and the computations were not parallelised. The numbers in table 5.4 are averages obtained on the 7 airways in the test data set with the optimal parameters. The differences between the $S_{k=16,\tau=2.5}$, $P_{\alpha=10^{-4},\beta=2}$ and $G_{\sigma=0.45}$ methods in time to calculate the minimum cut and memory usage are mostly due to differences in column length. Columns can probably be shortened, which would likely equalise the numbers without changing the segmentation results much.

The most interesting conclusion is that the three-dimensional methods are not much slower than the $2D$ method. This is likely due to complexities in resampling the images into a lot of two-dimensional slices. These two-dimensional slices are often much bigger than they need to be because of uncertainties in calculating branch diameters from the initial segmentation. Also note the much faster column construction of the method using the Gaussian kernel compared to the kernel based on the electric field strength, partly due to its separability.

5.5 Comparison with manual segmentations

Table 5.5 shows the results of the comparisons with the manual segmentations in the test data set, for each of the investigated methods using the optimal parameters. The flow line method using a Gaussian kernel with $\sigma = 0.45$ achieves the best result and the method using medial axes columns with $k = 16$ neighbours to define the medial axes performs the worst. It is interesting to note that $G_{\sigma=0.45}$ actually performs statistically better $p < 0.05$ than any of the other methods, and also has a significantly lower variance in the quality of the results ($p < 0.0001$). This is an indication, that not only is it the best performing method, it is also less likely to break down, as described in section 3.5.1.

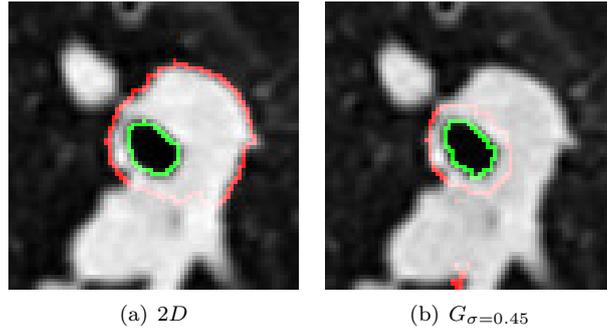


Figure 25. Figure 25(a) is an example cross-section of a segmentation *break down* of the 2D method, caused by only using two-dimensional information to infer the outer airway wall border position. Figure 25(b) illustrates how the $G_{\sigma=0.45}$ method more correctly segments it.

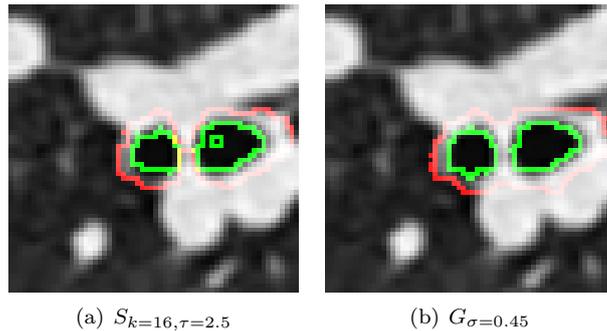


Figure 26. Cross-sections extracted near a branching area with segmentations obtained from the $S_{k=16, \tau=2.5}$ method in figure 26(a) and from the $G_{\sigma=0.45}$ method in figure 26(b)

The biggest advantage of the implemented three-dimensional methods compared to the 2D method, is their ability to use all three dimensions of information to infer the position of the surfaces. This was visible in the extracted two-dimensional cross-sections, in which the 2D method's tendency to include other nearby structures was apparent. Figure 25 is such an example, notice how the large abutting structure gets included in the 2D method's segmentation, whereas $G_{\sigma=0.45}$ is able to more correctly place the outer airway wall border. The tendency also likely explains the higher degree of variation observed in the segmentation results of this method.

We investigated the cross-sections segmented by the straight normal method $S_{k=16, \tau=2.5}$ and it seems as if the method's poor performance can be explained by two major problems. Poor quality normals and an inability to correctly handle the high curvature in bifurcation regions, as also predicted in section 2.2.6. The first problem results in some very erratic looking segmentations with spikes and even holes, which cannot be explained by the image's intensity function. The second makes the method very unlikely to correctly segment the area between two very close branches. We found these places to be full of holes, probably caused by self intersections or too short normals. See figure 26 for an example of both of these issues. Notice how the outer airway wall border segmentation obtained with $S_{k=16, \tau=2.5}$ collapses onto the inner airway wall border in the area between the two branches.

The reason the electric field strength inspired kernels $P_{\alpha=10^{-4}, \beta \in \{2,4\}}$ performed worse than the tested Gaussian kernel was less clear. In fact manually investigating the generated cross-sectional images revealed only small consistent differences in areas between two very close airways. Here the $P_{\alpha=10^{-4}, \beta \in \{2,4\}}$ methods often had a few small holes where the $G_{\sigma=0.45}$ method had none. We think this is an indication

Method	0	1	2	3	4	5
$2D$	643	1059	1232	1250	1253	1254
$G_{\sigma=0.45}$	1256	1256	1256	1256	1256	1256
Method	6	7	8	9	10	11
$2D$	1255	1250	1227	1132	937	577
$G_{\sigma=0.45}$	1256	1255	1234	1144	953	604

Table 6. Amount of samples with branches analysed in each of the generations, for the two methods.

of too much regularisation. Even with a β of 4 remote 'charges' have more effect on the local gradient than with the investigated Gaussian kernel. This causes the flow line columns to bend away from the other nearby initial segmentation too early leaving voxels in between the segmentations. A way to solve this is obviously to introduce less regularisation by increasing β and lowering α , but the effect of the so obtained kernels are likely to be similar to just using a Gaussian, but without the performance advantages of its separability.

5.6 Generation based comparison

The goal of this section is to evaluate the methods' and measures' reproducibility and ability to measure airway abnormalities related to lung function. A comparison of the best of the proposed methods ($G_{\sigma=0.45}$) and our previously developed two-dimensional method ($2D$) will be performed. The two methods might have different strengths and weaknesses in larger and smaller airways and so the comparison was first done on a generation basis. Results of measurements in generation ranges will be shown later.

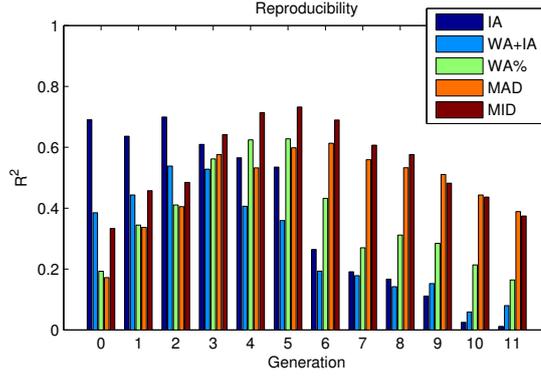
5.6.1 Differences in leak detection methods

The results presented in the following sections are meant to illustrate the capabilities of the actual measures being produced by the methods, so leaks being detected by our new three-dimensional method should not influence the results of the $2D$ method and vice-versa. One consequence of this is that the number of branches sampled in each generation are not equal. Table 6 illustrates the differences in the amount of generations that have branches analysed as a consequence of differences in the leak detection methods and extracted branches. It is clear that the leak detection method of $2D$ is much more aggressive. It is perhaps surprising that it also affects the larger airways, but it is consistent with the results of section 5.6.2 and 5.6.3, which seem to indicate that it has more problems with the larger airways.

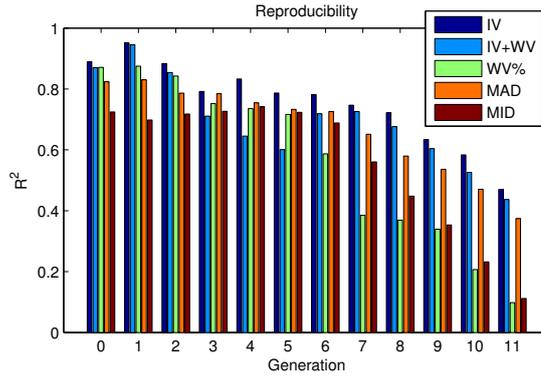
5.6.2 Reproducibility

Figure 27 shows the reproducibility of each of the measures using the $2D$ and the $G_{\sigma=0.45}$ methods, quantified as the coefficient of determination. That is, it is the total amount of variance in the measures at followup, which can be explained linearly by the measures at baseline. It is important to note that the graphs illustrate the combined reproducibility of either using the $2D$ method with the two-dimensional measures or the $G_{\sigma=0.45}$ method with the three-dimensional measures. So strictly speaking it cannot for instance be used to say anything about the $G_{\sigma=0.45}$ method alone. It should also be noted that no large outliers were observed in the data, as Pearson product moment correlation coefficients are sensitive to them.

Comparing the two-dimensional measures obtained with the $2D$ method with their three-dimensional counterparts obtained with $G_{\sigma=0.45}$ generation by generation revealed that IV and WV+IV were statistically more reproducible than IA and WA+IA in all generations ($p < 0.0001$), WV% and MAD were more reproducible



(a) 2D

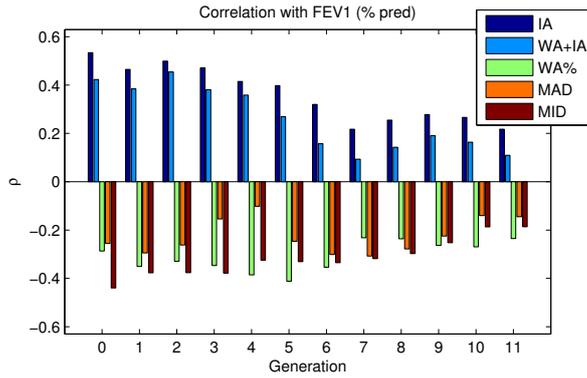


(b) $G_{\sigma=0.45}$

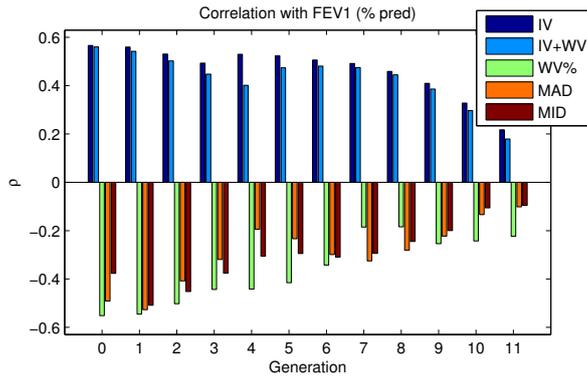
Figure 27. Reproducibility as a function of generation of the two-dimensional method in figure 27(a) and of our new three-dimensional method using the Gaussian kernel in figure 27(b). Here quantified as R^2 values are calculated from Pearson correlation coefficients of the measures at baseline and first year followup.

in generations 0 – 7 ($p < 0.001$) and no significant difference could be observed for the generations 8 – 11. MID were the only measure to be statistically less reproducible when obtained with the new three-dimensional method in any generation. This was observed in generations 8-11. However it was still more reproducible in generations 0 – 3 ($p < 0.05$). Figure 27 also indicates that the site of the biggest improvement in terms of reproducibility from using $G_{\sigma=0.45}$ and the three-dimensional measures as opposed to 2D and the two-dimensional measures is in the larger airway branches. In fact most of the two-dimensional measures, IA being the only exception, have a statistical significant peak around generation 2-6 where they are most reproducible ($p < 0.0001$). This is not the case for their three-dimensional counterparts. We think the reason for this is that the larger airways are closer to other structures, such as for instance the heart and major blood vessels, which make the segmentation of the outer airway wall border more difficult. The 2D method has difficulties differentiating between these structures, where as $G_{\sigma=0.45}$ is able to use three-dimensional information to infer the position of the border in such difficult areas. This was also indicated by investigations of cross-sections extracted in the larger airways, example seen in figure 25(a). Note though that some of the three-dimensional measures are less reproducible in the trachea than in generation 1, this is likely just a side-effect of it extending outside the image and thus being cut differently from scan to scan.

The improvement (if any) of the proposed method and measures seem to be least in the smallest analysed airways. We investigated some of the cases where the



(a) 2D



(b) $G_{\sigma=0.45}$

Figure 28. Spearman's correlation coefficients ρ of the measures in each generation with FEV1 (% pred) at baseline. Measures computed using the two-dimensional method in figure 28(a) and with our new three-dimensional method using the Gaussian kernel in figure 28(b).

measures fluctuated the most and it seem to be related to leakage within the initial segmentation. As explained earlier, the 2D method discards any cross-sections where the inner airway wall border is found outside the outer at just a single pixel. This is a very aggressive leakage detection method, which removes a lot of cross-sections from the analysis and a lot of them are not actual leaks. This is not a big problem, because there are so many of them. But the corresponding leakage detection method employed in the three-dimensional method removes whole branches at a time and for this reason the limit was set less aggressively at 1%, see section 3.6. However this also seem to have resulted in more undetected leaks, which probably influences the measures.

5.6.3 Correlation with lung function

Figure 28 shows Spearman's correlation coefficients of the measures with lung function quantified as FEV1 (% pred) in each generation. The correlation coefficients of IV and WV+IV were found to be higher than IA and WA+IA in every analysed generation. WV% and MAD correlated more than their two-dimensional counterparts in the larger airways and slightly less in the smaller airways. The three-dimensional MID was found to correlate less than the two-dimensional version in most generations, generation 1 and 2 being the only exceptions.

In general it is a bad idea to make conclusions about the accuracy of the method from correlations with lung function, because so little is known about how the measures are connected. A measure might for instance pick up some pathology within parenchyma, that is wrongly included in the segmentations, which causes

the correlation coefficients to rise. Even more likely, wrong segmentations could be caused by pathology, which could make a less stable algorithm more sensitive to disease progression. When that is said, it is important to establish that we are not just measuring something irrelevant very consistently and the conclusion from the data displayed in figure 28 is that we are not. The lung function correlation coefficients seem to tell a similar story as the reproducibility measurements in the previous section. That is, that the newly proposed three-dimensional method and measures for the most part result in measurements, that are more correlated with lung function than our previously developed two-dimensional method⁹ and that this correlation increase is most apparent in the larger airways.

Another interesting observation, which we will discuss in section 6.2.1, is that especially IV and WV+IV seem to gain the most when compared with the two-dimensional versions.

5.7 Per scan measures

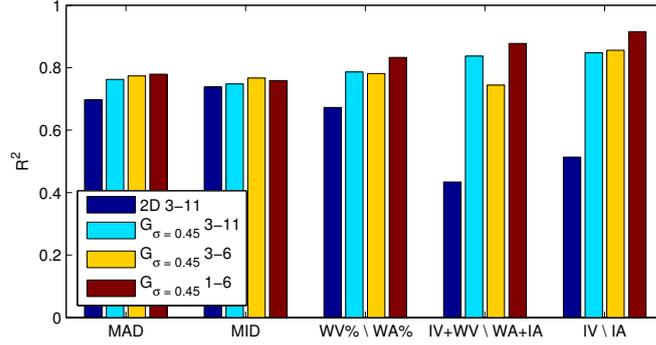
The purpose of this section is to show results of the measures obtained in branch generation ranges, and to illustrate that this way of combining measurements taken throughout the airway tree into one single measure makes sense. Additionally a comparison of the new three-dimensional algorithm with the previous two-dimensional method will be performed with the so obtained measures.

It is important to note, before we go any further with this analysis, that a choice of generation range is not just a matter of getting the highest correlation with lung function or reproducibility, as the relationship between airway measures and lung function is fairly complex with many covariates.⁵⁷ As such, the method of combining the measures from different sites in the airway should ideally be based on a detailed statistical analysis, covariates should be included, and the results should be verified on an independent data set. A future study could look into this.

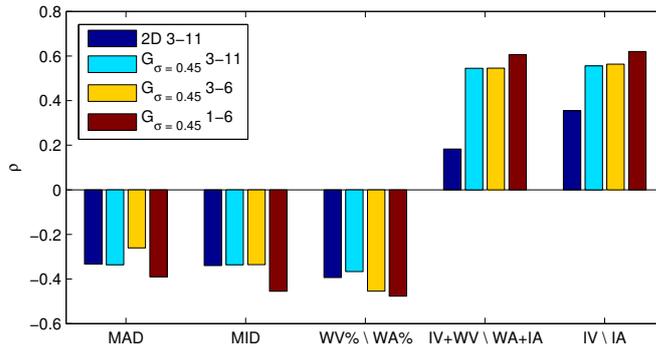
5.7.1 Reproducibility and correlation with lung function

Measures obtained with the 2D method have previously been established to correlate most with lung function and be most reproducible in the generation range 3-11.⁵⁷ This does not seem to be the case with the new method, judging by results of section 5.6.2 and 5.6.3. Because the new method is able to segment the larger airways even better than the airways of medium size, maybe the choice of generation range should be based just as much on the capability of the airway segmentation and tree extraction algorithm to consistently extract the branches. A generation range of 1-6 for instance would mean that roughly 80% of the branches would be correctly found and included in the analysis and there would only be an average variation in this number of at most 16% in any specific generation, see figure 19. The trachea should be ignored because it was not included in its entirety in the images.

Figure 29 show the reproducibility and lung function correlation coefficients of the 2D method over the previously established optimal range of 3-11 and the proposed method in the ranges 3 – 11, 3 – 6 and 1 – 6. It is interesting that even in the range of 3 – 11 where the 2D method works the best the newly proposed method produces measures of MAD, WV%, IV+WV and IV that are more ($p < 0.001$) and in the case of MID equally reproducible to the two-dimensional measures. The three-dimensional measures also correlated more with lung function, the only exceptions being MID and WV%, which correlate just slightly less than their two-dimensional counterparts. Changing the generation range interval over 3 – 6 to 1 – 6 increases the reproducibility of the three-dimensional measures further and the absolute values of the lung function correlation coefficients of every measure become larger than



(a) Reproducibility



(b) Lung function

Figure 29. Reproducibility of the 2D method measures sampled in range 3-11 and the $G_{\sigma=0.45}$ method measures sampled in ranges 3-11, 3-6 and 1-6 shown in figure 29(a). Correlation with lung function for the same methods, measures and generation ranges shown in figure 29(b).

the corresponding absolute values of the correlation coefficients obtained with the two-dimensional measures.

The reproducibility of FEV1 (% pred) within this data set was 0.87, which means that is statistically higher than all the measures, except IV ($p < 0.05$) in all the sampled generation ranges and IV+WV in 1 – 6. IV is actually statistically more reproducible in the generation range 1 – 6 than FEV1 (% pred).

We note that the reproducibility of these *per scan measures* generally do not become better than what is seen in generation 1 alone for instance. No general trend of larger lung function correlation coefficients can be observed either. However they may capture changes to a larger part of the airway without being much less reproducible than measures obtained in any single generation. It also seems likely given that we know COPD mostly affects the smaller airways that they better represent the true abnormalities caused by COPD and not just covariates.

5.7.2 COPD Diagnostic ability

Figure 30 shows the diagnostic ability of the methods and measures calculated in a generation range of 3–11. The AUC values of IV at stage 1 and 2, IV+WV at stage 1, 2 and 3 and MAD at stage 1 and 2 were all statistically larger than their two-dimensional counterparts ($p < 0.05$), whereas WV% was found to be statistically smaller in all stages ($p < 0.05$). No statistically significant difference was found within the rest of the measures. Although the results are not completely clear, a majority of the measures does show larger AUC values with the new method

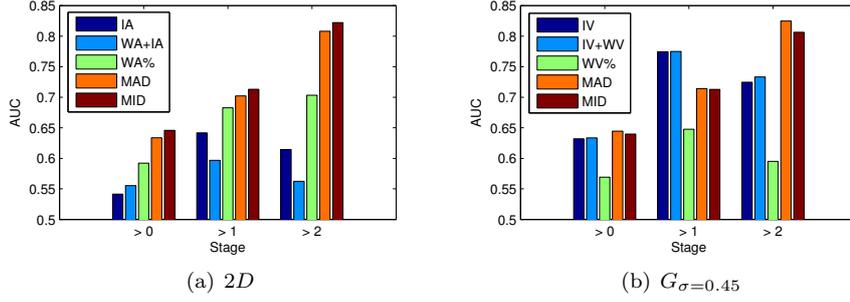


Figure 30. Diagnostic ability of the methods and measures quantified as AUC. Differentiating between subjects suffering from the disease and not (> 0), subjects with stage 2 and worse (> 1) and subjects with stage 3 and worse (> 2). Measures computed with the 2D and $G_{\sigma=0.45}$ methods are shown in figure 30(a) and figure 30(b) respectively.

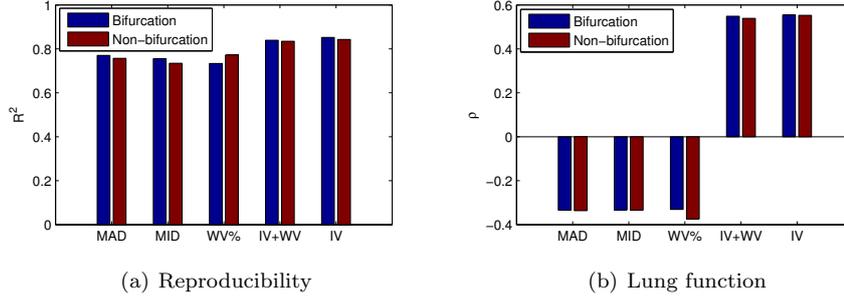


Figure 31. Reproducibility and correlation with lung function inside and outside bifurcation regions.

than with the previous method. The increases are also numerically much larger than the decreases. Again the largest improvements seem to be with the measures quantifying lumen and total airway size, IV and IV+WV.

We investigated the AUC values of the measures obtained with $G_{\sigma=0.45}$ in the generation ranges 3-6 and 1-6 as well (results not shown), however the larger lung function correlation coefficients and improved reproducibility observed in these ranges, does not seem to equate to better diagnostic ability. Perhaps this is an indication that the smaller airways are important, despite the lower accuracy and precision associated with measuring them.

5.7.3 Bifurcation regions

Figure 31 shows reproducibility and Spearman's correlation coefficient of the measures with FEV1 (% pred) inside and outside bifurcation regions. No significant difference was observed in reproducibility between the bifurcation and non-bifurcation regions using the measures MAD, MID, WV%, IV+WV, IV and a significance level of 0.05. This indicates that the proposed method is equally reproducible inside and outside bifurcation regions. The lung function correlation coefficients are also very similar, which is evidence to suggest that airway abnormalities caused by COPD are not more or less prevalent in bifurcation regions. Further investigations will have to establish whether measurements in the bifurcation and non-bifurcation regions are independent and thus complement each other.

5.7.4 St. George's Respiratory Questionnaire

Figure 32 show Spearman's rho of FEV1 (% pred) and each of the three-dimensional airway abnormality measures with the SGRQ components. Most of the coefficients indicate moderate links between the measures and how questionnaires were completed. Especially how subjects perceive their symptoms and what impact the

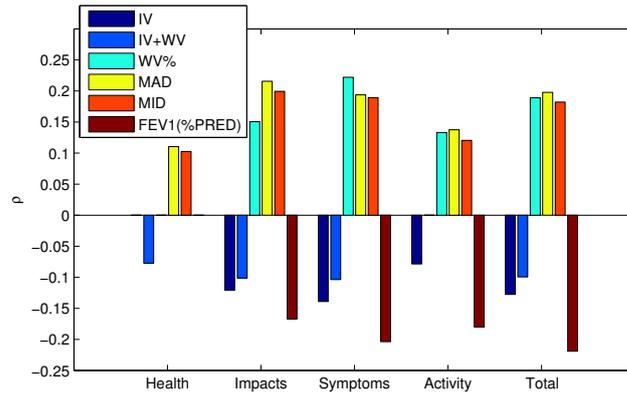


Figure 32. Significant Spearman’s rank correlation coefficients of each of the three-dimensional airway abnormality measures and FEV1 (% pred) with the SGRQ component scores (Impacts, Symptoms and Activity), total score (Total) and the extra posed question on subject health (Health).

disease has on their daily lives were found to be related to the analysed abnormality measures. WV%, MAD and MID were found to be positively correlated with the components, meaning that an increase in wall volume relative to the complete airway area, an increase in mean airway density or mean inner volume density were associated with a decrease in health status and perceived well being. The opposite was true for IV and IV+WV, meaning that a smaller lumen volume or airway volume were found to be associated with negative health effects. IV+WV was however not found to be significantly correlated with the activity component. These are all results that fall well in line with current theory, saying that COPD leads to an increase in mass within the airway walls and narrowed airways, which causes health problems. The figure also shows that out of the analysed measures, WV%, MAD and MID are likely the best airway indicators of subject health and perceived well being as quantified by SGRQ.

Figure 32 also shows that even the answer to the simple question of how subjects describe their current health is significantly correlated with airway abnormalities, specifically IV+WV, MAD and MID. This reflects the relationship, already indicated by the SGRQ components, that higher values of mean airway density, inner volume density or a lower airway volume are associated with a worse health status.

6. DISCUSSION

6.1 Medial axes and normal directions from Delaunay balls

It should be noted, that the problem of obtaining normals to discrete representations of continuous surfaces can be approached in many different ways and we have not investigated the subject in detail. However from the experience and results obtained within this work it seems to us that the method of reference 66 is more applicable to actual noisy point clouds and not, as in this problem, connected meshes of vertices obtained from an already rather correct segmentation of the lumen.³³ The errors in our case are probably not as much noise as discretisation effects.

We think that our mesh choice is a likely contributing factor to these problems, as it directly reflects the voxels of the initial segmentation. Some Delaunay balls will undoubtedly be very small due to its step-like nature. We do not think triangulations, such as those obtained with the popular marching cubes algorithm⁷⁷ are solutions to these problems, as they still directly reflect the discrete surface. We have experienced a little with pre-smoothing the surface, as done in reference,⁷⁹

instead with better results. Surface smoothing however can change the topology of the surface, and the objects shrink as a result.

Another more complicated solution could be to use a level set method⁶⁹ to obtain an implicit representation of the initial segmentation surface, and then create a mesh from this. This implicit surface representation could be used to estimate the medial axes⁹³ and normals more precisely.

Another interesting work uses parabolic partial differential equations to extract the medial axes from mesh representations, the method is robust to noise because of the diffusion equation used, which smooths the result. Unfortunately it seems slow judging from the reported results.⁹⁴

6.2 Two versus three-dimensional methods

It is interesting that the two-dimensional method work so well compared to the analysed three-dimensional methods, see Table 5.5. It can probably in part be explained by the comparison being done with two-dimensional data, which is likely to be biased toward a two-dimensional method. There is also a bias favouring the two-dimensional method, in that any other airway within the cross-section not correctly marked 'to be excluded', see the blue areas in figure 18(b), will count as errors in the three-dimensional methods, but not in the two-dimensional method. There is also a quality in the three-dimensional methods, not reflected in the comparison, in that they do not require an accurate extraction of centred perpendicular cross-sections to work properly, but two-dimensional methods do. However both types of methods suffer from this in the comparison, as the training set is build on such cross-sections.

When all this is said, we still think there is evidence to say that the two-dimensional method in accuracy is at least comparable to $S_{k=16, \tau=2.5}$. One might consider why this is so, given that three-dimensional methods should be able to utilise one more dimension of coherency in the data, and thus better be able to handle areas where the wall borders are weakly defined. We think that one reason might be that the smoothness penalties implemented by the two-dimensional method better captures the structure of the human airways. That is, it might be better to penalise deviance from some idealised circular/tubular structure than from the surface curvature of some voxelised initial segmentation. The first probably captures some general knowledge of the human airways, whereas the last, because it is generated from the same volume of data, is probably too specific and thus becomes sensitive to noise, mucus and other artefacts not reflecting the true nature of the airway.

This lead us to the conclusion, that a graph construction scheme based on connecting the inter column edges using the euclidean distance between column vertices, rather than how far they are from the initial surface, could improve our segmentation results. Such a scheme would prefer straighter surfaces in the euclidean sense, rather than surfaces following the isosurfaces in the convoluted images.

6.2.1 Measure differences

IV and IV+WV are the measures, which benefit the most from the proposed $G_{\sigma=0.45}$ method compared to the two dimensional versions, both in terms of reproducibility, increased correlation with lung function and diagnostic ability. This indicate an improvement in the segmentation method, given that these two measures directly quantify the segmented volumes and areas. There are however also other explanations for this. The branch volume could remain more constant than branch width during inspiration and expiration effects. For instance if these effects vary from

elongation and thinning during inspiration to shortening and widening during expiration. We also know, from previous investigations,⁵⁷ that the airway segmentation method is sensitive to pathology in such a way that sicker subjects have fewer successfully segmented branches. This enhances the association between the measures and the disease, an enhancement, which is likely to be even more present in the volume measures.

7. CONCLUSION

The motivation for solving the problem of segmenting human airway walls in CT images was described with focus on applications within COPD analysis and diagnosis.

Previous methods were described, reflected on and evaluated. Starting with simple one-dimensional methods, such as FWHM, over more sophisticated two-dimensional methods, using constraints and penalties to enforce solution coherency, to the most advanced methods of today, which are able to accurately and precisely determine the position of the airway wall borders in three dimensions. The methods' theory and purpose in practice were accounted for with a special emphasis on graph cut methods.

We have suggested new methods for solving the problem, by combining the airway segmentation method of reference 33 and the tree extraction method described in reference 28 with a novel graph construction technique, whose non-intersecting columns follow greatest ascent and descent flow lines within a continuous function calculated from the airway segmentation. An optimal solution was obtained using the maximum flow graph of reference 9. The method was inspired by the electric lines of force columns initially described in reference 2. We have documented how such columns can be calculated using simple convolutions, and that the Gaussian kernel could be a better choice than a kernel based on the electric field strength. We have furthermore shown how solutions obtained using the maximum flow graph of reference 9 relates to the optimal net family of methods and that a subset of the optimal VCE-weight net problems specified on proper ordered multi-column graphs, can be solved more efficiently using it. Additionally, a method using straight columns pointing in the normal direction and calculated using an approximation to the medial axes, similarly to reference 4, was documented and implemented as a reference method.

The methods were applied on data from the DLCST and CBQ. 649 cross-sectional images were extracted randomly from the airways of 15 subjects and manually segmented. The data set was then split randomly in a training and an independent test data set. Optimal values of method parameters were found using the training data set and quantitatively and qualitatively evaluated on the test data set. The proposed method using non-intersecting flow line columns calculated with a Gaussian kernel was found to be statistically superior in terms of mean segmentation relative area of overlap to the same method using flow line columns calculated from kernels based on the electric field strength, the method using straight columns and a previously developed two-dimensional method.⁹

Additionally a large scale comparison of the proposed method and the previously developed two-dimensional methods' COPD diagnostic abilities were performed. Airway abnormality measures obtained with the two methods were evaluated on reproducibility, using the coefficient of determination, calculated from Pearson product moment correlation coefficients, correlation with lung function, using Spearman's rank correlation coefficient of the measures and FEV1 (% pred), and diagnostic ability in the form of AUC values of each measure in the three first

disease stages. IV and IV+WV were found to be more reproducible in all examined generations (0-11) than similar two-dimensional measures computed using the previously developed two-dimensional method ($p < 0.0001$). WV% and MAD were similarly found to be more reproducible in generations 0-7 ($p < 0.001$) and no statistical difference was observed in the range 8-11, whereas MID was found to be more reproducible in generations 0-3 and less reproducible from 8-11 ($p < 0.05$). The proposed method also enabled measures that as a majority correlated more with lung function than similar measures obtained with the two-dimensional method. IV and IV+WV were found to correlate more in every examined generation, WV% and MAD correlated more in the larger and slightly less in the smaller airways, and MID was found to correlate less in most generations. Measures obtained with the proposed method in the branch generation range 3-11, previously found to be the best range for the two-dimensional method, were in the case of MAD, WV%, IV+WV and IV found to be more reproducible ($p < 0.001$), whereas no significant change could be observed for MID. The three-dimensional measures were also found to correlate more with lung function, with the exception of MID and WV%, which correlated just slightly less. The diagnostic ability of the measures obtained with the proposed method were for the most part better. IV had significantly higher AUC values in stage 1 and 2, IV+WV in stage 1, 2 and 3 and MAD at stage 1 and 2 ($p < 0.05$). WV% was found to be statistically worse in every examined stage ($p < 0.05$).

We also investigated, whether the bifurcation regions, which the previous two-dimensional method for the most part could not segment, were more or less influenced by the disease. The results show that the method is able to segment these areas with just as high reproducibility as the non-bifurcation regions ($p < 0.05$). The lung function correlation coefficients inside and outside these regions were also found to be very similar, indicating that airway abnormalities caused by COPD are neither less nor more prevalent in bifurcation areas.

The relationship between airway abnormality measures and how subjects perceive their health was investigated using the SGRQ. WV%, MAD and MID were found to be positively correlated with the SGRQ components, whereas IV and IV+WV were found to be negatively correlated ($p < 0.05$). We found that especially symptoms and what impact the disease had on their social functioning were related to airway abnormalities and that WV%, MAD and MID are likely the best indicators of subject health and perceived well being as quantified by SGRQ answers.

In general the results obtained from the work in this thesis are well in line with current theory, saying that COPD leads to an increase in mass within the airway walls and narrowed airways, which causes poorer lung function and other health problems. Results indicate that this increase in mass, can be quantified using the three-dimensional measures: the wall volume percentage, the mean airway density and the mean inner volume density. The narrowing can similarly be quantified using the inner volume and the total airway volume.

Even though measurements for the most part do not match basic spirometry in reproducibility, they do enable localization of the disease within the lungs, making it possible for instance to link symptoms with disease phenotypes. The results presented demonstrate the use of CT and the proposed method for COPD investigation and diagnosis.

We therefore conclude that we have fulfilled the goals set out in the project description.

8. FUTURE WORK

8.1 Improved leak detection

It is our opinion that removing leaks or erroneous areas in the initial segmentation, is one of the most important issues for improving measurement accuracy and precision.

The issue could be approached, for instance by building a reference of known leaks and accepted branches. It is possible that the Exact'09²⁸ database could be used for this. Using this reference a leak classifier could be build, perhaps using airway wall measurements as feature descriptors. This classifier could be used to obtain weights describing how believable a segmentation of any given branch was.

8.2 Parallelisation

Within this work a graph resolution of 0.5 mm was used, meaning the initial surface and the columns were sampled at a 0.5 mm spacing. This was done because higher resolutions, simply lead to unacceptable computation times. Experiments with resolutions of 1 mm showed much worse results and thus extrapolating leads to the conclusion that higher resolutions than 0.5 mm, might equate to even better results. To properly investigate this, the computation time and memory usage of the algorithm will have to be improved.

The algorithmic running time of the method is dominated by the used maximum flow algorithm. As mentioned, many algorithms exist and their worst case running times differ. The Boykov-Kolmogorov⁶² algorithm we used in this thesis has a worst case running time of $O(|E||V|^2|C|)$, where $|V|$ is the number of vertices, $|E|$ is the number of edges and $|C|$ is the cost of the minimum cut. So there is a lot to be gained if the graph could somehow be made smaller.

It is not likely that a graph representing the complete airway tree is needed to correctly segment for instance the trachea. In fact, it should be possible to measure how much the influence of given vertices and edges on the minimum cut decreases with distance in the graph. Such measurements could be used to split the graph in smaller parts, which individually would be faster to solve and require less memory.

8.3 Repetitive approach

The proposed method is very sensitive to the accuracy of initial segmentation. Mostly because the columns, for accuracy and performance reasons, need to be rather well resolved and short. It is possible that a repetitive approach could improve on this sensitivity, by perhaps evolving the surfaces in iterations, starting with longer columns with more widely spaced vertices and working towards smaller columns with closer spaced vertices as the solution converged. It is likely that the convolution kernels would have to use more regularization in the first iterations and this could then be lowered with iterations, using common concepts from scale space theory.

Before this is possible, some work would have to be done, on how to connect the two surfaces between iterations. We would also need to establish that multiple iterations actually converge to a solution.

APPENDIX A. ALGORITHMS

A.1 Normal and feature size algorithms

Assuming the existence of the following functions:

Points (M)	Returns the vertices of a mesh M .
Delaunay (P)	Finds the Delaunay triangulation of a point set P .
Nearest (P, k, p)	Finds the k nearest neighbours in a point set P to p .
LargestBall (D, p)	Given a Delaunay triangulation D , returns the largest Delaunay ball incident to a point p .
LargestInsideBall (M, D, p)	Given a Delaunay triangulation D , returns the largest Delaunay ball incident to a point p inside the mesh M .
LargestOutsideBall (M, D, p)	Given a Delaunay triangulation D , returns the largest Delaunay ball incident to a point p outside the mesh M .
Radius (B)	Return the radius of a ball B .
Centre (B)	Return the centre of a ball B .
Distance (p_0, p_1)	Return the euclidean distance between two points p_0 and p_1 .
Distance (P, p)	Return the euclidean distance from the point p to the point set P .
Outside (M, p)	Returns true if the point p is outside the mesh M .

Given a mesh M and an error tolerance τ , the following defines the normal calculation algorithm in pseudocode:

```

Normals( $M, \tau$ )
   $P \leftarrow$  Points( $M$ )
   $N \leftarrow$  Nil
   $E \leftarrow F \leftarrow \{\}$ 
   $D \leftarrow$  Delaunay( $P$ )
  for each  $p \in P$ 
     $m \leftarrow$  Nearest( $P, 1, p$ )
     $B \leftarrow$  LargestBall( $D, p$ )
    if Radius( $B$ ) >  $\tau$  Distance( $m, p$ )
       $c \leftarrow$  Centre( $B$ )
       $n \leftarrow (c-p)/\mathbf{Distance}(c, p)$ 
      if Outside( $M, c$ )
         $n \leftarrow -n$ 
       $N[p] \leftarrow n$ 
       $E \leftarrow E \cup \{p\}$ 
    else
       $F \leftarrow F \cup \{p\}$ 
  for each  $p \in F$ 
     $q \leftarrow$  Nearest( $E, 1, p$ )
     $N[p] \leftarrow N[q]$ 
  return  $N$ 

```

Given a neighbour search count k , the following defines the feature size algorithm in pseudocode:

```

Features( $M, k$ )
   $P \leftarrow \mathbf{Points}(M)$ 
   $D \leftarrow \mathbf{Delaunay}(P)$ 
   $M_i \leftarrow M_o \leftarrow \{\}$ 
  for each  $p \in P$ 
     $N \leftarrow \mathbf{Nearest}(P, k, p)$ 
     $r_i \leftarrow r_o \leftarrow 0$ 
    for each  $n \in N$ 
       $B_i \leftarrow \mathbf{LargestInsideBall}(M, D, n)$ 
      if  $\mathbf{Radius}(B_i) > r_i$ 
         $r_i \leftarrow \mathbf{Radius}(B_i)$ 
         $p_i \leftarrow \mathbf{Centre}(B_i)$ 
       $B_o \leftarrow \mathbf{LargestOutsideBall}(M, D, n)$ 
      if  $\mathbf{Radius}(B_o) > r_o$ 
         $r_o \leftarrow \mathbf{Radius}(B_o)$ 
         $p_o \leftarrow \mathbf{Centre}(B_o)$ 
     $M_i \leftarrow M_i \cup \{p_i\}$ 
     $M_o \leftarrow M_o \cup \{p_o\}$ 
   $F_i \leftarrow F_o \leftarrow \mathbf{Nil}$ 
  for each  $p \in P$ 
     $F_i[p] \leftarrow \mathbf{Distance}(M_i, p)$ 
     $F_o[p] \leftarrow \mathbf{Distance}(M_o, p)$ 
  return  $F_i, F_o$ 

```

A.2 Training algorithm

Assuming the existence of some function **Evaluate**(P) which runs the airway wall segmentation method on the complete training set with the parameters P and returns the sum of the errors of each involved cross-section as calculated in equation 65, an initial set of randomized parameters P , error tolerances for each parameter ϵ , left L and right R search interval limits and a maximum number of iterations k . The pseudocode is given below:

```
Optimise( $k, P, \epsilon, L, R$ )
   $E_p \leftarrow \mathbf{Evaluate}(P)$ 
  for  $i \leftarrow 0$  to  $k$ 
    for  $j \leftarrow 0$  to length( $P$ )
       $p \leftarrow P[j]$ 
       $l \leftarrow L[j], r \leftarrow R[j]$ 
      while true
         $\text{backup} \leftarrow P[j]$ 
         $P[j] \leftarrow l$ 
         $E_l \leftarrow \mathbf{Evaluate}(P)$ 
         $P[j] \leftarrow r$ 
         $E_r \leftarrow \mathbf{Evaluate}(P)$ 
         $P[j] \leftarrow \text{backup}$ 
        if  $E_l < E_r$ 
          if  $E_l < E_p$ 
             $E_p \leftarrow E_l$ 
             $p \leftarrow l$ 
             $\text{moved} \leftarrow \mathbf{true}$ 
            break
          else if  $E_l > E_r$ 
            if  $E_r < E_p$ 
               $E_p \leftarrow E_r$ 
               $p \leftarrow r$ 
               $\text{moved} \leftarrow \mathbf{true}$ 
              break
        if  $r - P[j] \leq \epsilon[j]$  and  $P[j] - l \leq \epsilon[j]$ 
          break
         $l \leftarrow (P[j] + l) / 2$ 
         $r \leftarrow (P[j] + r) / 2$ 
       $P[j] \leftarrow p$ 
    if  $\text{!moved}$ 
      break
```

REFERENCES

- [1] Wu, X. and Chen, D. Z., “Optimal net surface problems with applications,” in [*ICALP '02: Proceedings of the 29th International Colloquium on Automata, Languages and Programming*], 1029–1042, Springer-Verlag, London, UK (2002).
- [2] Yin, Y., Song, Q., and Sonka, M., “Electric field theory motivated graph construction for optimal medical image segmentation,” in [*GbRPR '09: Proceedings of the 7th IAPR-TC-15 International Workshop on Graph-Based Representations in Pattern Recognition*], 334–342, Springer-Verlag, Berlin, Heidelberg (2009).
- [3] Li, K., Millington, S., Wu, X., Chen, D. Z., and Sonka, M., “Simultaneous segmentation of multiple closed surfaces using optimal graph searching,” in [*IPMI'05: Proceedings of the 19th International Conference on Information Processing in Medical Imaging*], 406–417 (2005).
- [4] Liu, X., Chen, D. Z., Wu, X., and Sonka, M., “Optimal graph-based segmentation of 3d pulmonary airway and vascular trees across bifurcations,” in [*First International Workshop on Pulmonary Image Processing*], 103–111 (2008).
- [5] Heimann, T., Münzing, S., Meinzer, H.-P., and Wolf, I., “A shape-guided deformable model with evolutionary algorithm initialization for 3d soft tissue segmentation,” in [*IPMI'07: Proceedings of the 20th international conference on Information processing in medical imaging*], 1–12, Springer-Verlag, Berlin, Heidelberg (2007).
- [6] Yin, Y., Zhang, X., and Sonka, M., “Optimal multi-object multi-surface graph search segmentation: Full-joint cartilage delineation in 3d,” in [*Proceedings of the 12th annual Conference on Medical Image Understanding and Analysis*], 104–108 (2008).
- [7] Kainmueller, D., Lamecker, H., and Zachow, S., “Multi-object segmentation with coupled deformable models,” *Annals of the British Machine Vision Association (BMVA)* **5**, 1–10 (2009).
- [8] Liu, X., Chen, D. Z., Tawhai, M., Hoffman, E., and Sonka, M., “Measurement, evaluation and analysis of wall thickness of 3D airway trees across bifurcations,” in [*The Second International Workshop on Pulmonary Image Processing*], 161–171 (2009).
- [9] Petersen, J., Lo, P., Nielsen, M., , Edula, G., Ashraf, H., Dirksen, A., and de Bruijne, M., “Quantitative analysis of airway abnormalities in CT,” in [*Medical Imaging 2010: Computer-Aided Diagnosis. Edited by N. Karssemeijer; R. M. Summers. Proceedings of SPIE*], **7624** (2010).
- [10] Pedersen, J. H., Ashraf, H., Dirksen, A., Bach, K., Hansen, H., Toennesen, P., Thorsen, H., Brodersen, J., Skov, B. G., Døssing, M., Mortensen, J., Richter, K., Clementsen, P., and Seersholm, N., “The danish randomized lung cancer CT screening trial—overall design and results of the prevalence round.,” *Journal of Thoracic Oncology* **4**, 608–614 (May 2009).
- [11] Rabe, K. F., Hurd, S., Anzueto, A., Barnes, P. J., Buist, S. A., Calverley, P., Fukuchi, Y., Jenkins, C., Rodriguez-Roisin, R., van Weel, C., and Zielinski, J., “Global strategy for the diagnosis, management, and prevention of chronic obstructive pulmonary disease: GOLD executive summary.,” *American Journal of Respiratory and Critical Care Medicine* **176**, 532–555 (September 2007).
- [12] Hogg, J. C., Chu, F., Utokaparch, S., Woods, R., Elliot, W. M., Buzatu, L., Cherniack, R. M., Rogers, R. M., Sciurba, F. C., Coxson, H. O., and Paré, P. D., “The nature of small-airway obstruction in chronic obstructive pulmonary disease,” *The New England Journal of Medicine* **350**(26), 2645–2653 (2004).

- [13] Ooi, G. C., Khong, P. L., Chan-Yeung, M., Ho, J. C. M., Chan, P. K. S., Lee, J. C. K., Lam, W. K., and Tsang, K. W. T., “High-resolution CT quantification of bronchiectasis: clinical and functional correlation,” *Radiology* **225**, 663–672 (2002).
- [14] Hillard, T. N., Regamey, N., Shute, J. K., Nicholson, A. G., Alton, E. W. F. W., Bush, A., and Davies, J. C., “Airway remodelling in children with cystic fibrosis,” *Thorax* **63**, 1074–1080 (2007).
- [15] Martínez, T. M., Llapur, C. J., Williams, T. H., Coates, C., Gunderman, R., Cohen, M. D., Howenstine, M. S., Saba, O., Coxson, H. O., and Tepper, R. S., “High-resolution computed tomography imaging of airway disease in infants with cystic fibrosis,” *American Journal of Respiratory and Critical Care Medicine* **172**, 1133–1138 (2005).
- [16] Bousquet, J., Chanez, P., Lacoste, J. Y., White, R., Vic, P., Godard, P., and Michel, F. B., “Asthma: a disease remodeling the airways,” *Allergy* **47**, 3–11 (1992).
- [17] Tockman, M. S., Anthonisen, N. R., Wright, E. C., and Donithan, M. G., “Airways obstruction and the risk for lung cancer,” *Annals of Internal Medicine* **106**(4), 512–518 (1987).
- [18] Weibel, E. R., “What makes a good lung?,” *Swiss Medical Weekly* **139**(27-28), 375–86 (2009).
- [19] Weinheimer, O., Achenbach, T., Bletz, C., Düber, C., Kauczor, H.-U., and Heussel, C. P., “About objective 3-D analysis of airway geometry in computerized tomography,” *IEEE Transactions on Medical Imaging* **27**(1), 64–74 (2008).
- [20] Kiraly, A. P., Odry, B. L., Naidich, D. P., and Novak, C. L., “Boundary-specific cost functions for quantitative airway analysis,” in [*MICCAI’07: Proceedings of the 10th international conference on medical image computing and Computer-Assisted Intervention-Part I*], 784–791, Springer-Verlag, Berlin, Heidelberg (2007).
- [21] Estépar, R. S. J., Washko, G. G., Silverman, E. K., Reilly, J. J., Kikinis, R., and Westin, C.-F., “Accurate airway wall estimation using phase congruency,” in [*MICCAI’06: Proceedings of the Ninth International Conference on Medical Image Computing and Computer-Assisted Intervention-Part II*], 125–134 (2006).
- [22] Tschirren, J., Hoffman, E. A., McLennan, G., and Sonka, M., “Intrathoracic airway trees: segmentation and airway morphology analysis from low-dose CT scans,” *IEEE Transactions on Medical Imaging* **24**(12), 1529–1539 (2005).
- [23] Tschirren, J., Hoffman, E. A., McLennan, G., and Sonka, M., “Segmentation and quantitative analysis of intrathoracic airway trees from computed tomography images,” *Proceedings of the American Thoracic Society* **2**, 484–487 (2005).
- [24] Li, K., Wu, X., Chen, D. Z., and Sonka, M., “Efficient optimal surface detection: theory, implementation, and experimental validation,” in [*Medical Imaging 2004: Image Processing. Edited by Fitzpatrick, J. Michael; Sonka, Milan. Proceedings of the SPIE.*], **5370**, 620–627 (2004).
- [25] Li, K., Wu, X., Chen, D. Z., and Sonka, M., “Optimal surface segmentation in volumetric images—a graph-theoretic approach,” *IEEE Transactions on Pattern Analysis and Machine Intelligence* **28**(1), 119–134 (2006).
- [26] Saragaglia, A., Fetita, C., Prêteux, F., Brillet, P. Y., and Grenier, P. A., “Accurate 3d quantification of the bronchial parameters in MDCT,” in [*Mathematical Methods in Pattern and Image Analysis. Proceedings of SPIE*], **5916** (2005).

- [27] Saragaglia, A., Fetita, C., and Prêteux, F., “Assessment of airway remodeling in asthma: Volumetric versus surface quantification approaches,” in [*Medical Image Computing and Computer-Assisted Intervention MICCAI 2006*], Larsen, R., Nielsen, M., and Sporring, J., eds., *Lecture Notes in Computer Science* **4191**, 413–420, Springer Berlin / Heidelberg (2006).
- [28] Lo, P., van Ginneken, B., Reinhardt, J. M., and de Bruijne, M., “Extraction of airways from CT (EXACT’09),” in [*The Second International Workshop on Pulmonary Image Analysis*], 175–189 (2009).
- [29] Wiemker, R., Bülow, T., and Lorenz, C., “A simple centricity-based region growing algorithm for the extraction of airways,” in [*The Second International Workshop on Pulmonary Image Analysis*], (2009).
- [30] Feuerstein, M., Kitasaka, T., and Mori, K., “Adaptive branch tracing and image sharpening for airway tree extraction in 3-d chest ct,” in [*The Second International Workshop on Pulmonary Image Analysis*], (2009).
- [31] Tschirren, J., Yavarna, T., and Reinhardt, J. M., “Airway segmentation framework for clinical environments,” in [*The Second International Workshop on Pulmonary Image Analysis*], (2009).
- [32] Weinheimer, O., Achenbach, T., and Düber, C., “Fully automated extraction of the airways from CT scans based on self-adapting region growing,” in [*The Second International Workshop on Pulmonary Image Analysis*], (2009).
- [33] Lo, P., Sporring, J., Ashraf, H., Pedersen, J. J. H., , and de Bruijne, M., “Vessel-guided airway segmentation based on voxel classification,” *The first International Workshop on Pulmonary Image Analysis* **1**, 113–122 (2008).
- [34] Lee, J. and Reeves, A. P., “Segmentation of the airway tree from chest CT using local volume of interest,” in [*The Second International Workshop on Pulmonary Image Analysis*], (2009).
- [35] Schlathöler, T., Lorenz, C., Carlsen, I. C., Renisch, S., and Deschamps, T., “Simultaneous segmentation and tree reconstruction of the airways for virtual bronchoscopy,” in [*Medical Imaging 2002: Image Processing. Proceedings of SPIE*], (feb 2002).
- [36] Aykac, D., Hoffman, E. A., McLennan, G., and Reinhardt, J. M., “Segmentation and analysis of the human airway tree from three-dimensional X-ray CT images,” *IEEE Transactions on Medical Imaging* **22**, 940–950 (August 2003).
- [37] Pisupati, C., Wolff, L., Zerhouni, E., and Mitzner, W., “Segmentation of 3d pulmonary trees using mathematical morphology,” in [*Mathematical morphology and its applications to image and signal processing*], 409–416 (1996).
- [38] Fetita, C., Ortner, M., Brillet, P.-Y., Prêteux, F., and Grenier, P., “A morphological-aggregate approach for 3d segmentation of pulmonary airways from generic msct acquisitions,” in [*The Second International Workshop on Pulmonary Image Analysis*], (2009).
- [39] Graham, M. W., Gibbs, J. D., and Higgins, W. E., “Robust system for human airway-tree segmentation,” in [*Medical Imaging 2008: Image Processing. Proceedings of SPIE*], **6914** (2008).
- [40] Tschirren, J., Palágyi, K., Reinhardt, J. M., Hoffman, E. A., and Sonka, M., “Segmentation, skeletonization, and branchpoint matching - a fully automated quantitative evaluation of human intrathoracic airway trees,” in [*MICCAI ’02: Proceedings of the 5th International Conference on Medical Image Computing and Computer-Assisted Intervention-Part II*], 12–19, Springer-Verlag, London, UK (2002).
- [41] Palágyi, K., Sorantin, E., Balogh, E., Kuba, A., Halmai, C., Erdohelyi, B., and Hausegger, K., “A sequential 3d thinning algorithm and its medical applications,” in [*IPMI ’01: Proceedings of the 17th International Conference on*

- Information Processing in Medical Imaging*], 409–415, Springer-Verlag, London, UK (2001).
- [42] Malandain, G., Bertrand, G., and Ayache, N., “Topological segmentation of discrete surfaces,” *International Journal of Computer Vision* **10**(2), 183–197 (1993).
 - [43] Perchet, D., Fetita, C. I., Vial, L., Prêteux, F., Caillibotte, G., Sbiërlea-Apiou, G., and Thiriet, M., “Virtual investigation of pulmonary airways in volumetric computed tomography,” *Computer Animation and Virtual Worlds* **15**(3-4), 361–376 (2004).
 - [44] Namkug, K., Beom, S. J., Sik, S. K., Jin, C. E., and Ho, K. S., “Semi-automatic measurement of the airway dimension by computed tomography using the full-width-half-maximum method: a study on the measurement accuracy according to the CT parameters and size of the airway,” *Korean journal of radiology : official journal of the Korean Radiological Society* **9**, 226–235 (2008).
 - [45] Nakano, Y., Wong, J. C., de Jong, P. A., Buzatu, L., Nagao, T., Coxson, H. O., Elliott, W. M., Hogg, J. C., and Paré, P. D., “The prediction of small airway dimensions using computed tomography,” *American Journal of Respiratory and Critical Care Medicine* **171**, 142–146 (2005).
 - [46] Nakano, Y., Muro, S., Sakai, H., Hirai, T., Chin, K., Tsukino, M., Nishimura, K., Itoh, H., Paré, P. D., Hogg, J. C., and Mishima, M., “Computed tomography measurements of airway dimensions and emphysema in smokers: correlation with lung function,” *American Journal of Respiratory and Critical Care Medicine* **162**, 1102–1108 (2000).
 - [47] Washko, G. R., Dransfield, M. T., Estépar, R. S. J., Diaz, A., Matsuoka, S., Yamashiro, T., Hatabu, H., Silverman, E. K., Bailey, W. C., and Reilly, J. J., “Airway wall attenuation: a biomarker of airway disease in subjects with COPD,” *Journal of Applied Physiology* **107**, 185–191 (April 2009).
 - [48] Lee, Y. K., Oh, Y.-M., Lee, J.-H., Kim, E. K., Lee, J. H., Kim, N., Seo, J. B., and Lee, S. D., “Quantitative assessment of emphysema, air trapping, and airway thickening on computed tomography,” *Lung* **186**, 277 (August 2008).
 - [49] Morrone, M. C. and Owens, R. A., “Feature detection from local energy,” *Pattern Recognition Letters* **6**(5), 303–313 (1987).
 - [50] Reinhardt, J., N.D.D’Souza, and E.A.Hoffman, “Accurate measurement of intrathoracic airways,” *IEEE Transactions on Medical Imaging* **16**, 820–827 (December 1997).
 - [51] Achenbach, T., Weinheimer, O., Biedermann, A., Schmitt, S., Freudenstein, D., Goutham, E., Kunz, R. P., Buhl, R., Dueber, C., and Heussel, C. P., “MDCT assessment of airway wall thickness in COPD patients using a new method: correlations with pulmonary function tests,” *European Radiology* **18**(12), 2731–2738 (2008).
 - [52] King, G. G., Müller, N. L., Whittal, K. P., Xiang, Q.-S., and Paré, P. D., “An analysis algorithm for measuring airway lumen and wall areas from high-resolution computed tomographic data,” *American Journal of Respiratory and Critical Care Medicine* **161**, 574–580 (2000).
 - [53] Saba, O. I., Hoffman, E. A., and Reinhardt, J. M., “Maximizing quantitative accuracy of lung airway lumen and wall measures obtained from X-ray CT imaging,” *Journal of Applied Physiology* **95**, 1063–1075 (2003).
 - [54] Chabat, F., Hu, X., Hansell, D., and Yang, G., “ERS transform for the automated detection of bronchial abnormalities on CT of the lungs,” *IEEE Transactions on Medical Imaging* **20**, 942–952 (September 2001).

- [55] Wiemker, R., Blaffert, T., Bülow, T., Renisch, S., and Lorenz, C., “Automated assessment of bronchial lumen, wall thickness and bronchoarterial diameter ratio of the tracheobronchial tree using high-resolution CT.,” in [*CARS*], Lemke, H. U., Inamura, K., Doi, K., Vannier, M. W., Farman, A. G., and Reiber, J. H. C., eds., *International Congress Series* **1268**, 967–972, Elsevier (2004).
- [56] Li, K. and Jolly, M.-P., “Simultaneous detection of multiple elastic surfaces with application to tumor segmentation in CT images,” in [*Medical Imaging 2008: Image Processing. Edited by Joseph M. Reinhardt; Josien P. W. Pluim. Proceedings of SPIE*], **6914**, 69143S–69143S–11 (2008).
- [57] Petersen, J., “Quantitative analysis of airway measures associated with chronic obstructive pulmonary disease.” Department of Computer Science, University of Copenhagen, Master project (2010).
- [58] Greig, D. M., Porteous, B. T., and Seheult, A. H., “Exact maximum a posteriori estimation for binary images,” *Journal of the Royal Statistical Society. Series B (Methodological)* **51**(2), 271–279 (1989).
- [59] Ford, L. R. and Fulkerson, D. R., “Maximal flow through a network,” *Canadian Journal of Mathematics* **8**, 399–404 (1956).
- [60] Edmonds, J. and Karp, R. M., “Theoretical improvements in algorithmic efficiency for network flow problems,” *Journal of the ACM* **19**, 248–264 (1972).
- [61] Goldberg, A. V. and Tarjan, R. E., “A new approach to the maximum flow problem,” in [*STOC ’86: Proceedings of the eighteenth annual ACM symposium on Theory of computing*], 136–146, ACM, New York, NY, USA (1986).
- [62] Boykov, Y. and Kolmogorov, V., “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 (2004).
- [63] Garvin, M., Abramoff, M., Kardon, R., Russell, S., Wu, X., and Sonka, M., “Intraretinal layer segmentation of macular optical coherence tomography images using optimal 3-d graph search,” *IEEE Transactions on Medical Imaging* **27**, 1495–1505 (October 2008).
- [64] Haeker, M., Wu, X., Abramoff, M., Kardon, R., and Sonka, M., “Incorporation of regional information in optimal 3-d graph search with application for intraretinal layer segmentation of optical coherence tomography images,” in [*IPMI’07: Proceedings of the 20th international conference on Information processing in medical imaging*], 607–618, Springer-Verlag, Berlin, Heidelberg (2007).
- [65] Attali, D., daniel Boissonnat, J., and Edelsbrunner, H., “Stability and computation of medial axes - a state-of-the-art report,” in [*Mathematical Foundations of Scientific Visualization, Computer Graphics, and Massive Data Exploration*], Springer (2007).
- [66] Dey, T. K. and Sun, J., “Normal and feature estimation from noisy point clouds,” in [*Proceedings of the 26th International Conference on Foundations of Software Technology and Theoretical Computer Science*], 21–32 (2006).
- [67] Canny, J., “A computational approach to edge detection,” *IEEE Transactions on Pattern Analysis and Machine Intelligence* **8**(6), 679–698 (1986).
- [68] Deriche, R., “Using canny’s criteria to derive a recursively implemented optimal edge detector,” *International Journal of Computer Vision* **1**(2), 167–187 (1987).
- [69] Chan, T. F. and Vese, L. A., “Active contours without edges,” *IEEE Transactions on Image Processing* **10**(2), 266–277 (2001).

- [70] Patel, B. D., Coxson, H. O., Pillai, S. G., Agusti, A. G. N., Calverley, P. M. A., Donner, C. F., Make, B. J., Müller, N. L., Rennard, S. I., Vestbo, J., Wouters, E. F. M., Hiorns, M. P., Nakano, Y., Camp, P. G., Fauerbach, P. V. N., Sreaton, N. J., Campbell, E. J., Anderson, W. H., Paré, P. D., Levy, R. D., Lake, S. L., Silverman, E. K., Lomas, D. A., and on behalf of the International COPD Genetics Network, “Airway wall thickening and emphysema show independent familial aggregation in chronic obstructive pulmonary disease,” *American Journal of Respiratory and Critical Care Medicine* **178**, 500–505 (2008).
- [71] Berger, P., Perot, V., Desbarats, P., de Lara, J. M. T., Marthan, R., and Laurent, F., “Airway wall thickness in cigarette smokers: Quantitative thin-section CT assessment,” *Radiology* **235**, 1055–1064 (2005).
- [72] Orlandi, I., Moroni, C., Camiciottoli, G., Bartolucci, M., Pistolesi, M., Villari, N., and Mascalchi, M., “Chronic obstructive pulmonary disease: Thin-section CT measurement of airway wall thickness and lung attenuation,” *Radiology* **234**, 604–610 (2005).
- [73] Coxson, H. O., Quiney, B., D., Xing, L., McWilliams, A. M., Mayo, J. R., and Lam, S., “Airway wall thickness assessed using computed tomography and optical coherence tomography,” *American Journal of Respiratory and Critical Care Medicine* **177**, 1201–1206 (2008).
- [74] Aziz, Z. A., Wells, A. U., Desai, S. R., Ellis, S. M., Walker, A. E., MacDonald, S., and Hansell, D. M., “Functional impairment in emphysema: Contribution of airway abnormalities and distribution of parenchymal disease,” *American Journal of Roentgenology* **185**(6), 1509–1524 (2005).
- [75] Balásházy, I., Hofmann, W., and Heistracher, T., “Local particle deposition patterns may play a key role in the development of lung cancer,” *Journal of Applied Physiology* **94**, 1719–1725 (January 2003).
- [76] Ramuzat, A. and Riethmuller, M. L., “PIV investigation of oscillating flows within a 3D lung multiple bifurcations model,” in [*11th International Symposium on Application of Laser Techniques to Fluid mechanics*], (2002).
- [77] Lorensen, W. E. and Cline, H. E., “Marching cubes: A high resolution 3D surface construction algorithm,” *SIGGRAPH Computer Graphics* **21**(4), 163–169 (1987).
- [78] Akleman, E. and Chen, J., “Regular meshes,” in [*SPM '05: Proceedings of the 2005 ACM symposium on Solid and physical modeling*], 213–219, ACM, New York, NY, USA (2005).
- [79] Fourey, S. and Malgouyres, R., “Normals estimation for digital surfaces based on convolutions,” *Computers & Graphics* **33**(1), 2–10 (2009).
- [80] Oomes, S., Snoeren, P., and Dijkstra, T., “3d shape representation: Transforming polygons into voxels,” in [*SCALE-SPACE '97: Proceedings of the First International Conference on Scale-Space Theory in Computer Vision*], 349–352, Springer-Verlag, London, UK (1997).
- [81] Huang, J., Yagel, R., Filippov, V., and Kurzion, Y., “An accurate method for voxelizing polygon meshes,” in [*VVS '98: Proceedings of the 1998 IEEE symposium on Volume visualization*], 119–126, ACM, New York, NY, USA (1998).
- [82] Feito, F. R. and Torres, J. C., “Inclusion test for general polyhedra,” *Computers & Graphics* **21**(1), 23–30 (1997).
- [83] Ogayar, C. J., Segura, R. J., and Feito, F. R., “Point in solid tests for triangle meshes. comparative study,” *Eurographics* (2003).
- [84] Nooruddin, F. S. and Turk, G., “Simplification and repair of polygonal models using volumetric techniques,” *IEEE Transactions on Visualization and Computer Graphics* **9**(2), 191–205 (2003).

- [85] Butcher, J. C., [*Numerical methods for ordinary differential equations*], Wiley, second ed. (2008).
- [86] Fehlberg, E., “Klassische runge-kutta-formeln vierter und niedrigerer ordnung mit schrittweiten-kontrolle und ihre anwendung auf wärmeleitungsprobleme,” *Computing (Arch. Elektron. Rechnen)* **6**, 61–71 (1970).
- [87] Pellegrino, R., Viegi, G., Brusasco, V., Crapo, R. O., Burgos, F., Casaburi, R., Coates, A., van der Grinten, C. P. M., Gustafsson, P., Hankinson, J., Jensen, R., Johnson, D. C., MacIntyre, N., McKay, R., Miller, M. R., Navajas, D., Pedersen, O. F., and Wanger, J., “Interpretative strategies for lung function tests,” *European Respiratory Journal* **26**, 948–968 (November 2005).
- [88] Ashraf, H., Edula, G., de Bruijne, M., Nielsen, M., Dirksen, A., Dahlback, M., and Pedersen, J. H., “A longitudinal follow-up of COPD disease progression in a large cohort of smokers screened for early lung cancer: CBQ study,” in [*American Thoracic Society International Conference*], (2009).
- [89] Jones, P. W., Quirk, F. H., and Baveystock, C. M., “The St. George’s Respiratory Questionnaire,” *Respiratory Medicine* **85**, 25–31 (September 1991).
- [90] Wilson, C. B., Jones, P. W., O’Leary, C. J., Cole, P. J., and Wilson, R., “Validation of the St. George’s Respiratory Questionnaire in bronchiectasis,” *American Journal of Respiratory Critical Care Medicine* **156**, 536–541 (August 1997).
- [91] Pasipanodya, J. G., Miller, T. L., Vecino, M., Munguia, G., Bae, S., Drweyer, G., and Weis, S. E., “Using the St. George Respiratory Questionnaire to ascertain health quality in persons with treated pulmonary tuberculosis,” *Chest* **132**, 1591–1598 (november 2007).
- [92] Beretta, L., Santaniello, A., Lemos, A., Masciocchi, M., and Scorza, R., “Validity of the Saint George’s Respiratory Questionnaire in the evaluation of the health-related quality of life in patients with interstitial lung disease secondary to systemic sclerosis,” *Rheumatology* **46**(2), 296–301 (2007).
- [93] Teixeira, R. C., “Medial axes and mean curvature motion i: Regular points,” *Journal of Visual Communication and Image Representation* **13**, 135–155 (march 2002).
- [94] Du, H. and Qin, H., “Medial axis extraction and shape manipulation of solid objects using parabolic pdes,” in [*SM ’04: Proceedings of the ninth ACM symposium on Solid modeling and applications*], 25–35, Eurographics Association, Aire-la-Ville, Switzerland, Switzerland (2004).